

Multicast and Routing Guide

8200zl

ProCurve Switches
K.12.XX

www.procurve.com



ProCurve Series 8200zl Switches

September 2007
K.12.xx

Multicast and Routing Guide

© Copyright 2007 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

5991-8586
September 2007

Applicable Products

ProCurve Switch 8212zl (J8715A)

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	xv
Printed Publications.....	xv
Electronic Publications.....	xv
Software Feature Index	xvi

1 Getting Started

Contents	1-1
Introduction	1-2
Conventions	1-2
Command Syntax Statements	1-2
Command Prompts	1-3
Screen Simulations	1-3
Configuration and Operation Examples	1-3
Keys	1-3
Sources for More Information	1-4
Getting Documentation From the Web	1-6
Online Help	1-6
Menu Interface	1-6
Command Line Interface	1-7
Web Browser Interface	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-8
Physical Installation	1-8

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
Number of Multicast Filters 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	2-19
CLI Commands for IGMP Proxy Configuration	2-21
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

3 PIM-DM (Dense Mode)

Contents	3-1
Overview	3-2
Introduction	3-3
Feature Overview	3-4
PIM-DM Operation	3-5

Multicast Flow Management	3-7
General Configuration Elements	3-9
Terminology	3-9
PIM-DM Operating Rules	3-10
Configuring PIM-DM	3-11
Global and PIM Configuration Contexts	3-12
PIM VLAN (Interface) Configuration Context	3-15
Displaying PIM Data and Configuration Settings	3-22
Displaying PIM Route Data	3-23
Displaying PIM Status	3-28
Operating Notes	3-35
Troubleshooting	3-37
Messages Related to PIM Operation	3-38
Applicable RFCs	3-41
Exceptions to Support for RFC 2932 - Multicast Routing MIB ..	3-42

4 PIM-SM (Sparse Mode)

Contents	4-1
Introduction	4-4
Feature Overview	4-5
Terminology	4-6
PIM-SM Operation and Router Types	4-9
PIM-SM Operation	4-9
Rendezvous-Point Tree (RPT)	4-9
Shortest-Path Tree (SPT)	4-10
Restricting Multicast Traffic to Rendezvous-Point Trees (RPTs)	4-11
Maintaining an Active Route for Multicast Group Members ...	4-11
Border Routers and Multiple PIM-SM Domains	4-12
PIM-SM Router Types	4-12
Designated Router (DR)	4-12
Bootstrap Router (BSR)	4-13

Rendezvous Point (RP)	4-14
Static Rendezvous Point (Static-RP)	4-17
Operating Rules and Recommendations	4-19
Configuration Steps for PIM-SM	4-20
Planning Considerations	4-20
Per-Router Global Configuration Context	4-20
Per-VLAN PIM-SM Configuration	4-21
Router PIM Configuration	4-23
Configuring PIM-SM on the Router	4-25
Global Configuration Context for Supporting PIM-SM	4-26
Global Configuration Context Commands	4-26
Example of Configuring for PIM Support at the Global Level ..	4-27
VLAN Context Commands for Configuring PIM-SM	4-28
Enabling or Disabling IGMP in a VLAN	4-28
Enabling or Disabling PIM-SM Per-VLAN	4-29
Changing the Interval for PIM-SM Neighbor Notification	4-30
Changing the Randomized Delay Setting for PIM-SM Neighbor Notification	4-31
Changing the PIM-SM Neighbor Timeout Interval	4-31
Enabling or Disabling LAN Prune Delay	4-32
Changing the LAN-Prune-Delay Interval	4-33
Changing the DR (Designated Router) Priority	4-33
Example of Configuring PIM-SM Support in a VLAN Context ..	4-34
Router PIM Context Commands for Configuring PIM-SM Operation	4-35
Configuring a BSR Candidate	4-35
Configuring Candidate-RPs on PIM-SM Routers	4-37
Enabling, Disabling, or Changing Router PIM Notification Traps	4-41
Changing the Global Join-Prune Interval on the Router	4-42
Changing the Shortest-Path Tree (SPT) Operation	4-42
Statically Configuring an RP To Accept Multicast Traffic	4-42
Example of Configuring PIM-SM Support in the Router PIM Context	4-43
Displaying PIM-SM Data and Configuration Settings	4-46
Displaying Multicast Route Data	4-47

Listing Basic Route Data for Active Multicast Groups	4-47
Listing Data for an Active Multicast Group	4-48
Listing All VLANs Having Currently Active PIM Flows	4-50
Displaying PIM-Specific Data	4-51
Displaying the Current PIM status and Global Configuration . .	4-51
Displaying Current PIM Entries Existing In the Multicast Routing Table	4-52
Displaying a Specific PIM Entry Stored in the Multicast Routing Table	4-53
Listing Currently Configured PIM Interfaces	4-55
Displaying IP PIM VLAN Configurations	4-55
Displaying PIM Neighbor Data	4-57
Displaying BSR Data	4-59
Displaying BSR Status and Configuration	4-59
Listing Non-Default BSR Configuration Settings	4-60
Displaying the Current RP Set	4-61
Displaying Candidate-RP Data	4-63
Displaying the Router's Candidate-RP Status and Configuration	4-63
Listing Non-Default C-RP Configuration Settings	4-64
Operating Notes	4-65
Event Log Messages	4-66

5 IP Routing Features

Contents	5-1
Overview of IP Routing	5-6
IP Interfaces	5-7
IP Tables and Caches	5-7
ARP Cache Table	5-8
IP Route Table	5-8
IP Forwarding Cache	5-9
IP Route Exchange Protocols	5-10
IP Global Parameters for Routing Switches	5-10
IP Interface Parameters for Routing Switches	5-12

Configuring IP Parameters for Routing Switches	5-13
Configuring IP Addresses	5-13
Changing the Router ID	5-13
Configuring ARP Parameters	5-15
How ARP Works	5-15
Enabling Proxy ARP	5-17
Configuring Forwarding Parameters	5-17
Changing the TTL Threshold	5-18
Enabling Forwarding of Directed Broadcasts	5-18
Configuring ICMP	5-19
Disabling ICMP Messages	5-19
Disabling Replies to Broadcast Ping Requests	5-19
Disabling ICMP Destination Unreachable Messages	5-20
Disabling ICMP Redirects	5-21
Configuring Static IP Routes	5-21
Static Route Types	5-21
Other Sources of Routes in the Routing Table	5-22
Static IP Route Parameters	5-22
Static Route States Follow VLAN States	5-23
Configuring a Static IP Route	5-23
Displaying Static Route Information	5-25
Configuring the Default Route	5-25
Configuring RIP	5-26
Overview of RIP	5-26
RIP Parameters and Defaults	5-27
RIP Global Parameters	5-27
RIP Interface Parameters	5-27
Configuring RIP Parameters	5-28
Enabling RIP	5-28
Enabling IP RIP on a VLAN	5-29
Changing the RIP Type on a VLAN Interface	5-29
Changing the Cost of Routes Learned on a VLAN Interface	5-29
Configuring RIP Redistribution	5-30
Define RIP Redistribution Filters	5-30
Modify Default Metric for Redistribution	5-31

Enable RIP Route Redistribution	5-31
Changing the Route Loop Prevention Method	5-33
Displaying RIP Information	5-33
Displaying General RIP Information	5-34
Displaying RIP Interface Information	5-35
Displaying RIP Peer Information	5-37
Displaying RIP Redistribution Information	5-38
Displaying RIP Redistribution Filter (restrict) Information	5-39
Configuring OSPF	5-40
Terminology	5-41
Overview of OSPF	5-43
OSPF Router Types	5-44
Interior Routers	5-44
Area Border Routers (ABRs)	5-44
Autonomous System Boundary Router (ASBR)	5-45
Designated Routers	5-45
OSPF Area Types	5-47
Backbone Area	5-48
Normal Area	5-48
Not-So-Stubby-Area (NSSA)	5-49
Stub Area	5-50
OSPF RFC Compliance	5-50
Reducing AS External LSAs and Type-3 Summary LSAs	5-50
Algorithm for AS External LSA Reduction	5-51
Replacing Type-3 Summary LSAs and Type-7 Default External LSAs with a Type-3 Default Route LSA	5-52
Dynamic OSPF Activation and Configuration	5-52
General Configuration Steps for OSPF	5-53
Configuration Rules	5-54
OSPF Global and Interface Settings	5-54
Configuring OSPF on the Routing Switch	5-56
1. Enable IP Routing	5-56
2. Enable Global OSPF Routing	5-56
3. Changing the RFC 1583 OSPF Compliance Setting	5-57
4. Assign the Routing Switch to OSPF Areas	5-59

5. Assign VLANs and/or Subnets to Each Area	5-64
6. Optional: Assigning Loopback Addresses to an Area	5-65
7. Optional: Configure for External Route Redistribution in an OSPF Domain	5-67
8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone	5-70
9. Optional: Influence Route Choices by Changing the Administrative Distance Default	5-73
10: Optional: Change OSPF Trap Generation Choices	5-74
11. Optional: Adjust Performance by Changing the VLAN or Subnet Interface Settings	5-75
12. Optional: Configuring OSPF Interface Authentication	5-80
13. Configuring an ABR To Use a Virtual Link to the Backbone	5-82
Configuring a Virtual Link	5-83
Optional: Adjust Virtual Link Performance by Changing the Interface Settings	5-85
Configuring OSPF Authentication on a Virtual Link	5-88
Displaying OSPF Information	5-90
Displaying General OSPF Configuration Information	5-91
Displaying OSPF Area Information	5-92
Displaying OSPF External Link State Information	5-93
Displaying OSPF Interface Information	5-95
Displaying OSPF Interface Information for a Specific VLAN or IP Address	5-97
Displaying OSPF Link State Information	5-98
Displaying OSPF Neighbor Information	5-101
Displaying OSPF Redistribution Information	5-103
Displaying OSPF Redistribution Filter (restrict) Information	5-103
Displaying OSPF Virtual Neighbor Information	5-104
Displaying OSPF Virtual Link Information	5-105
Displaying OSPF Route Information	5-107
Displaying OSPF Traps Enabled	5-109
OSPF Equal-Cost Multipath (ECMP) for Different Subnets Available Through the Same Next-Hop Routes	5-109
Displaying the Current IP Load-Sharing Configuration	5-110

Configuring IRDP	5-112
Enabling IRDP Globally	5-113
Enabling IRDP on an Individual VLAN Interface	5-113
Displaying IRDP Information	5-114
Configuring DHCP Relay	5-115
Overview	5-115
DHCP Packet Forwarding	5-115
Unicast Forwarding	5-115
Broadcast Forwarding	5-116
Prerequisites for DHCP Relay Operation	5-116
Enabling DHCP Relay	5-116
Configuring an IP Helper Address	5-117
Operating Notes	5-117
Hop Count in DHCP Requests	5-117
Disabling the Hop Count in DHCP Requests	5-117
Operating Notes	5-118
Verifying the DHCP Relay Configuration	5-118
Displaying the DHCP Relay Setting	5-118
Displaying DHCP Helper Addresses	5-119
Displaying the Hop Count Setting	5-120
DHCP Option 82	5-120
Option 82 Server Support	5-121
Terminology	5-122
General DHCP Option 82 Requirements and Operation	5-123
Option 82 Field Content	5-124
Forwarding Policies	5-127
Configuration Options for Managing DHCP Client Request Packets	5-127
Multiple Option 82 Relay Agents in a Client Request Path	5-128
Validation of Server Response Packets	5-129
Multinetted VLANs	5-131
Configuring Option 82	5-131
Example of Option 82 Configuration	5-133
Operating Notes	5-135

UDP Broadcast Forwarding	5-137
Overview	5-137
Subnet Masking for UDP Forwarding Addresses	5-138
Configuring and Enabling UDP Broadcast Forwarding	5-139
Globally Enabling UDP Broadcast Forwarding	5-139
Configuring UDP Broadcast Forwarding on Individual VLANs	5-139
Displaying the Current IP Forward-Protocol Configuration	5-141
Operating Notes for UDP Broadcast Forwarding	5-142
Messages Related to UDP Broadcast Forwarding	5-142

6 Virtual Router Redundancy Protocol (VRRP)

Contents	6-1
Overview	6-3
Terminology	6-4
General Operation	6-5
Virtual Router (VR)	6-8
Virtual IP Address	6-8
Master Router	6-9
Owner Router	6-9
Backup Router	6-10
Virtual Router MAC Address	6-10
VRRP and ARP	6-11
General Operating Rules	6-11
Steps for Provisioning VRRP Operation	6-13
Basic Configuration Process	6-13
Example Configuration	6-15
Associating More Than One Virtual IP Address With a VR	6-17
Configuring VRRP	6-18
Enabling VRRP in the Global Configuration Context	6-18
Creating a VR and Entering the VR Context	6-19
Configuring a VR Instance on a VLAN Interface	6-20
Changing VR Advertisement Interval and Source IP Address ..	6-22

Preempt Mode on VRRP Backup Routers	6-24
Enabling or Disabling VRRP Operation on a VR	6-24
Displaying VRRP Configuration and Statistics Data	6-25
VRRP Configuration Data	6-25
Displaying the VRRP Global Configuration	6-25
Displaying All VR Configurations on the Router	6-25
Displaying a Specific VR Configuration	6-27
VRRP Statistics Data	6-28
Displaying Global VRRP Statistics Only	6-28
Displaying Statistics for All VRRP Instances on the Router	6-29
Displaying Statistics for All VRRP Instances in a VLAN	6-32
Displaying Statistics for a Specific VRRP Instance	6-33
Standards Compliance	6-33
Operating Notes	6-34
Event Log Messages	6-35

Index

Product Documentation

About Your Switch Manual Set

Note

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

Printed Publications

The two publications listed below are printed and shipped with your switch. The latest version of each is also available in PDF format on the ProCurve Web site, as described in the Note at the top of this page.

- *Read Me First*—Provides software update information, product notes, and other information.
- *Installation and Getting Started Guide*—Explains how to prepare for and perform the physical installation and connect the switch to your network.

Electronic Publications

The latest version of each of the publications listed below is available in PDF format on the ProCurve Web site, as described in the Note at the top of this page.

- *Management and Configuration Guide*—Describes how to configure, manage, and monitor basic switch operation.
- *Advanced Traffic Management Guide*—Explains how to configure traffic management features such as VLANs, MSTP, QoS, and Meshing.
- *Multicast and Routing Guide*—Explains how to configure IGMP, PIM, IP routing, and VRRP features.
- *Access Security Guide*—Explains how to configure access security features and user authentication on the switch.
- *Release Notes*—Describe new features, fixes, and enhancements that become available between revisions of the main product guide.

Software Feature Index

For the software manual set supporting your ProCurve 8212zl switch model, this feature index indicates which manual to consult for information on a given software feature.

Both Intelligent Edge and Premium Edge software features are available on the Procurve 8212zl switch.

Premium Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
OSPF			X	
PIM-DM (Dense Mode)			X	
PIM-SM (Sparse Mode)			X	
VRRP			X	

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
802.1Q VLAN Tagging		X		
802.1X Port-Based Priority	X			
802.1X Multiple Authenticated Clients Per Port				X
ACLs				X
AAA Authentication				X
Authorized IP Managers				X
Authorized Manager List (Web, Telnet, TFTP)				X
Auto MDIX Configuration	X			
BOOTP	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Config File	X			
Console Access	X			
Copy Command	X			
CoS (Class of Service)		X		
Debug	X			
DHCP Configuration	X			
DHCP Option 82			X	
DHCP Snooping				X
DHCP/Bootp Operation	X			
Diagnostic Tools	X			
Downloading Software	X			
Dynamic ARP Protection				X
Eavesdrop Protection				X
Event Log	X			
Factory Default Settings	X			
Flow Control (802.3x)	X			
File Management	X			
File Transfers	X			
Friendly Port Names	X			
Guaranteed Minimum Bandwidth (GMB)	X			
GVRP		X		
Identity-Driven Management (IDM)		X		
IGMP			X	
Interface Access (Telnet, Console/Serial, Web)	X			
IP Addressing	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
IP Routing			X	
Jumbo Packets	X			
LACP	X			
Link	X			
LLDP	X			
LLDP-MED	X			
MAC Address Management	X			
MAC Lockdown				X
MAC Lockout				X
MAC-based Authentication				X
Management VLAN		X		
Meshing		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
ProCurve Manager (PCM)	X			
Ping	X			
Port Configuration	X			
Port Monitoring		X		
Port Security				X
Port Status	X			
Port Trunking (LACP)	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Port-Based Access Control (802.1X)				X
Power over Ethernet (PoE)	X			
Protocol Filters				X
Protocol VLANs		X		
Quality of Service (QoS)		X		
RADIUS Authentication and Accounting				X
RADIUS-Based Configuration				X
Rate-Limiting	X			
Redundant Management	X			
RIP			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, TFTP, Xmodem)	X			
Source-Port Filters				X
Spanning Tree (STP, RSTP, MSTP)		X		
SSHv2 (Secure Shell) Encryption				X
SSL (Secure Socket Layer)				X
Syslog	X			
System Information	X			
TACACS+ Authentication				X

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Telnet Access	X			
TFTP	X			
Time Protocols (TimeP, SNTP)	X			
Traffic Mirroring	X			
Traffic/Security Filters				X
Troubleshooting	X			
Uni-Directional Link Detection (UDLD)	X			
UDP Forwarder			X	
USB Device Support	X			
Virus Throttling (Connection-Rate Filtering)				X
VLANs		X		
VLAN Mirroring (1 static VLAN)		X		
Voice VLAN		X		
Web Authentication RADIUS Support				X
Web-based Authentication				X
Web UI	X			
Xmodem	X			

Getting Started

Contents

Introduction	1-2
Conventions	1-2
Command Syntax Statements	1-2
Command Prompts	1-3
Screen Simulations	1-3
Configuration and Operation Examples	1-3
Keys	1-3
Sources for More Information	1-4
Getting Documentation From the Web	1-6
Online Help	1-6
Menu Interface	1-6
Command Line Interface	1-7
Web Browser Interface	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-8
Physical Installation	1-8

Introduction

This guide is intended for use with the ProCurve Switch 8212zl. It 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 xiii.

You can download documentation from the ProCurve Networking web site, www.procurve.com.

Conventions

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

Command Syntax Statements

Syntax: ip < default-gateway < *ip-addr* >> | routing >

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 8212z1#
```

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, 2007 13:43:13
               K.12.01
               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
```

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 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 “Software Feature Index” on page xiv.

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 1-6.

- **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 1-6.
- **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 1-6.)

- *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, traffic control, and PoE operation
 - Redundant management
 - 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: 802.1D (STP), 802.1w (RSTP), and 802.1s (MSTP)
 - meshing
 - Quality-of-Service (QoS)
 - Access Control Lists (ACLs)
- *Multicast and Routing Guide*—Use this guide for information topics such as:
 - IGMP
 - PIM (SM and DM)
 - IP routing
 - VRRP
- *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.

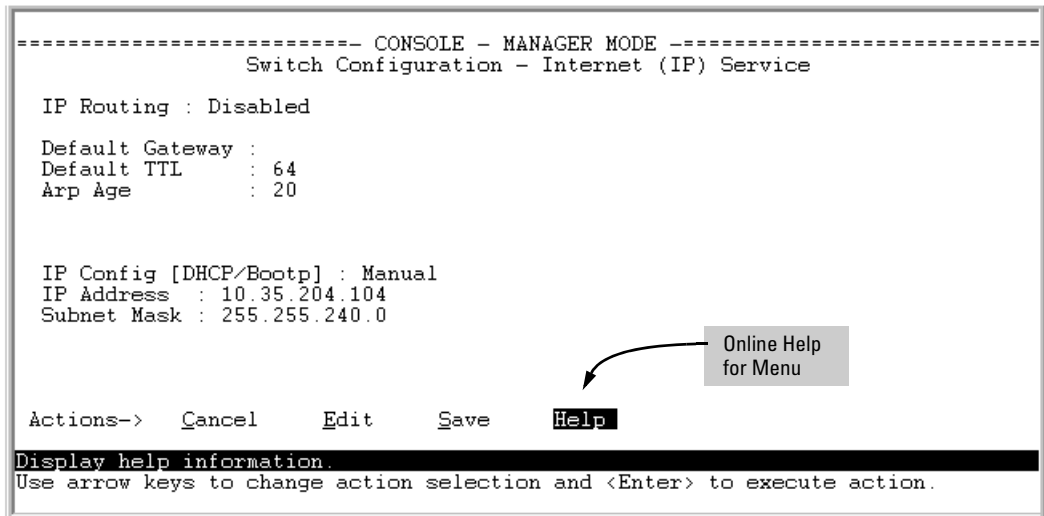
If you need further information on ProCurve switch technology, visit the ProCurve Networking web site at:

www.procurve.com

Online Help

Menu Interface

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



```
----- CONSOLE - MANAGER MODE -----  
Switch Configuration - Internet (IP) Service  
  
IP Routing : Disabled  
  
Default Gateway :  
Default TTL      : 64  
Arp Age         : 20  
  
IP Config [DHCP/Bootp] : Manual  
IP Address      : 10.35.204.104  
Subnet Mask    : 255.255.240.0  
  
Actions->  Cancel  Edit  Save  Help  
  
Display help information.  
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 1-2. Online Help for Menu Interface

Command Line Interface

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

```
ProCurve# write help
Usage: write <memory|terminal>

Description: View or save the running configuration of the switch.

      write terminal - displays the running configuration of the
                    switch on the terminal
      write memory  - saves the running configuration of the
                    switch to flash. The saved configuration
                    becomes the boot-up configuration of the switch
                    the next time it is booted.
```

Figure 1-3. CLI Help

Web Browser Interface

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. You can access the Help by clicking on the question mark button in the upper right corner of any of the web browser interface screens.




Figure 1-4. Button for Web Browser Interface Online Help

Note

To access the online Help for the ProCurve web browser interface, you need either ProCurve Manager (version 1.5 or greater) installed on your network or an active connection to the World Wide Web. Otherwise, Online help for the web browser interface will not be available.

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 1-6.

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
Number of Multicast Filters 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	2-19
CLI Commands for IGMP Proxy Configuration	2-21
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 2,047 (if data driven) or 2,048 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.

Number of Multicast Filters Allowed

The number of multicast filters allowed depends on the number of configured VLANs:

- 16 multicast filters if VLANs \leq 1024
- 8 multicast filters if VLANs $>$ 1024

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.

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 vlan < 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.)

Multimedia Traffic Control with IP Multicast (IGMP)

CLI: Configuring and Displaying IGMP

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 8212zl Switch 6400cl Switch 6200yl Switch 5400zl Switch 5300xl Switch 4200vl Switch 3500yl Switch 3400cl 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”.

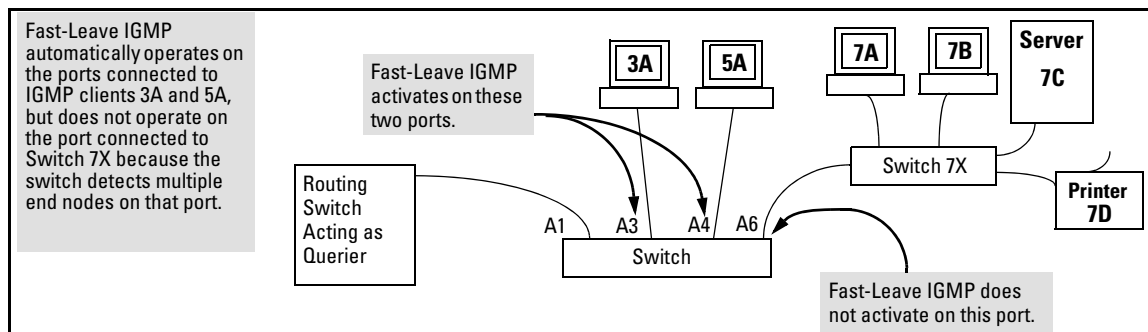


Figure 2-3. Example of Automatic Fast-Leave IGMP Criteria

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

Note

For more information about PIM-DM and PIM-SM, see the chapters “*PIM-DM (Dense Mode)*” and “*PIM-SM (Sparse Mode)*” in this guide.

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.

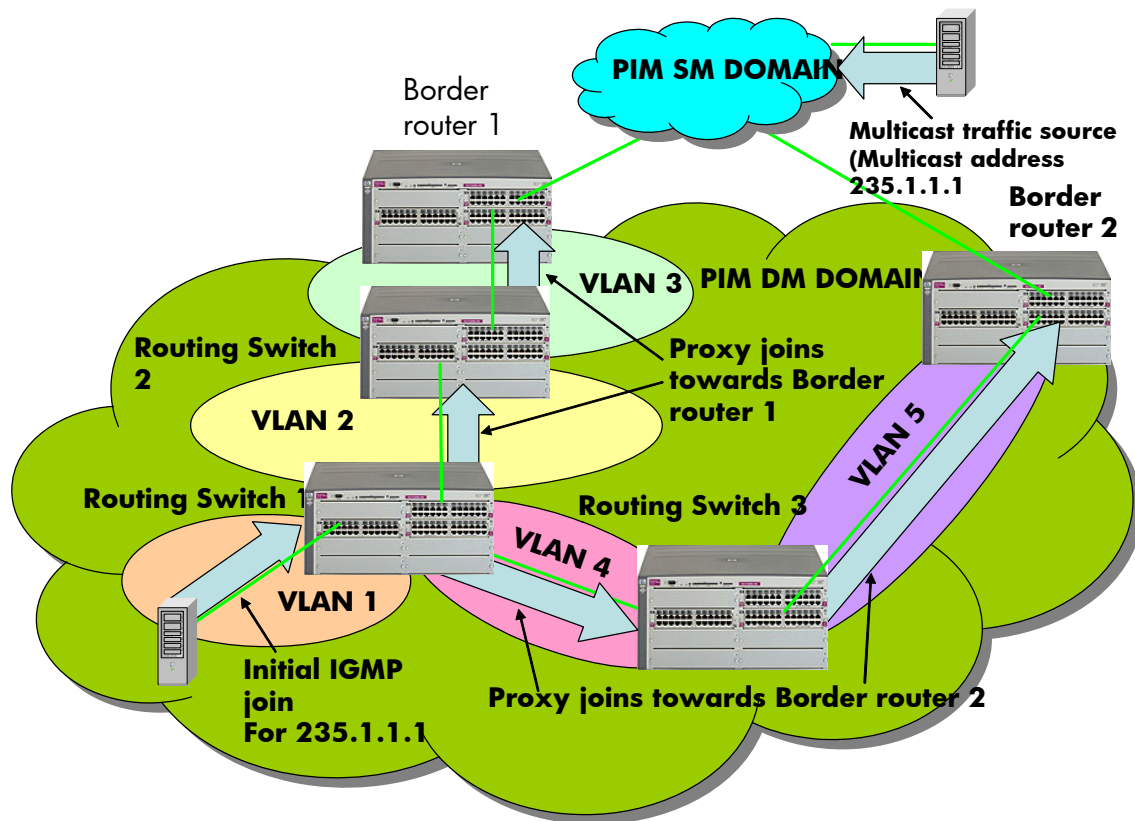


Figure 2-4. IGMP Proxy Example

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 224.0.1.0 - 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.
- 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 Series 5400zl switches, the Series 3500yl switches, the switch 6200yl, the switch 8212zl, 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.

PIM-DM (Dense Mode)

Contents

Overview	3-2
Introduction	3-3
Feature Overview	3-4
PIM-DM Operation	3-5
Multicast Flow Management	3-7
General Configuration Elements	3-9
Terminology	3-9
PIM-DM Operating Rules	3-10
Configuring PIM-DM	3-11
Global and PIM Configuration Contexts	3-12
PIM VLAN (Interface) Configuration Context	3-15
Displaying PIM Data and Configuration Settings	3-22
Displaying PIM Route Data	3-23
Displaying PIM Status	3-28
Operating Notes	3-35
Troubleshooting	3-37
Messages Related to PIM Operation	3-38
Applicable RFCs	3-41
Exceptions to Support for RFC 2932 - Multicast Routing MIB ..	3-42

Overview

This chapter describes protocol-independent multicast routing operation on the switches covered in this guide and how to configure it with the switch's built-in interfaces, and assumes an understanding of multimedia traffic control with IP multicast (IGMP), which is described in chapter 2, "Multimedia Traffic Control with IP Multicast (IGMP)".

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 Web Browser Interface"
- Chapter 6, "Switch Memory and Configuration"

Introduction

Feature	Default	Menu	CLI	Web
Configure PIM Global	n/a	—	3-12	—
Configure PIM VLAN Interface	n/a	—	3-15	—
Display PIM Route Data	Disabled	—	3-23	—
Display PIM Status	0 (Forward All)	—	3-28	—

In a network where IP multicast traffic is transmitted for multimedia applications, such traffic is blocked at routed interface (VLAN) boundaries unless a multicast routing protocol is running. Protocol Independent Multicast (PIM) is a family of routing protocols that form multicast trees to forward traffic from multicast sources to subnets that have used a protocol such as IGMP to request the traffic. PIM relies on the unicast routing tables created by any of several unicast routing protocols to identify the path back to a multicast source (*reverse path forwarding*, or RPF). With this information, PIM sets up the distribution tree for the multicast traffic. The PIM-DM and PIM-SM protocols on the switches covered by this manual enable and control multicast traffic routing.

IGMP provides the multicast traffic link between a host and a multicast router running PIM-DM or PIM-SM. IGMP and either PIM-DM or PIM-SM must be enabled on VLANs whose member ports have directly connected hosts with a valid need to join multicast groups.

PIM-DM is used in networks where, at any given time, multicast group members exist in relatively large numbers and are present in most subnets. PIM-SM (described in Chapter 4 of this guide) is used in networks where multicast sources and group members are sparsely distributed over a wide area and can result in unnecessary multicast traffic on routers outside the distribution paths needed for traffic between a given multicast source and the hosts belonging to the multicast group. In such networks, PIM-SM can be used to reduce the effect of multicast traffic flows in network areas where they are not needed. And because PIM-SM does not automatically flood traffic, it is a logical choice in lower bandwidth situations such as WAN environments.

Feature Overview

PIM-DM on the switches covered in this guide includes:

- **Routing Protocol Support:** PIM uses whichever unicast routing protocol is running on the routing switch. These can include:
 - RIP
 - OSPF
 - Static routes
 - Directly connected interfaces
- **VLAN Interface Support:** Up to 128 outbound VLANs are supported in the multicast routing table (MRT) at any given time. This means the sum of all outbound VLANs across all current flows on a router may not exceed 128. (A single flow may span one inbound VLAN and up to 128 outbound VLANs, depending on the VLAN memberships of the hosts actively belonging to the flow.)
- **Flow Capacity:** Up to 2048 flows are supported in hardware across a maximum of 128 outbound VLANs. (A flow is composed of an IP source address and an IP multicast group address, regardless of the number of active hosts belonging to the multicast group at any given time.)
- **IGMP Compatibility:** PIM-DM is compatible with IGMP versions 1 - 3, and is fully interoperable with IGMP for determining multicast flows.
- **VRRP:** PIM-DM is fully interoperable with VRRP to quickly transition multicast routes in the event of a failover.
- **MIB Support:** With some exceptions, PIM-DM supports the parts of the Multicast Routing MIB applicable to PIM-DM operation. (Refer to “Exceptions to Support for RFC 2932 - Multicast Routing MIB” on page 3-42.)
- **PIM Draft Specifications:** Compatible with PIM-DM draft specification, versions 1 and 2.

PIM-DM Operation

PIM-DM operates at the router level to direct traffic for a particular multicast group along the most efficient path to the VLANs having hosts that have joined that group. A unicast source address and a multicast group address comprise a given source/group (S/G) pair. Multicast traffic moving from a source to a multicast group address creates a *flow* to the area(s) of the network requiring the traffic. That is, the flow destination is the multicast group address, and not a specific host or VLAN. Thus, a single multicast flow has one source and one multicast group address (destination), but may reach many hosts in different subnets, depending on which hosts have issued joins for the same multicast group.

PIM routes the multicast traffic for a particular S/G pair on paths between the source unicast address and the VLANs where it is requested (by joins from hosts connected to those VLANs). Physical destinations for a particular multicast group can be hosts in different VLANs or networks. Individual hosts use IGMP configured per-VLAN to send joins requesting membership in a particular multicast group. All hosts that have joined a given multicast group (defined by a multicast address) remain in that group as long as they continue to issue periodic joins.

On the switches covered in this guide, PIM-DM interoperates with IGMP and the switch's routing protocols. (Note that PIM-DM operates independently of the routing protocol you choose to run on your switches, meaning you can use PIM-DM with RIP, OSPF, or static routes configured.) PIM-DM utilizes a unicast routing table to find the path to the originator of the multicast traffic and sets up multicast "trees" for distributing multicast traffic. (This method is termed *reverse path forwarding*, or *RPF*).

For the flow of a given multicast group, PIM-DM creates a "tree" structure between the source and the VLANs where hosts have joined the group. The tree structure consists of:

- Extended branches to VLANs with hosts that currently belong to the group
- Pruned branches to VLANs with no hosts that belong to the group

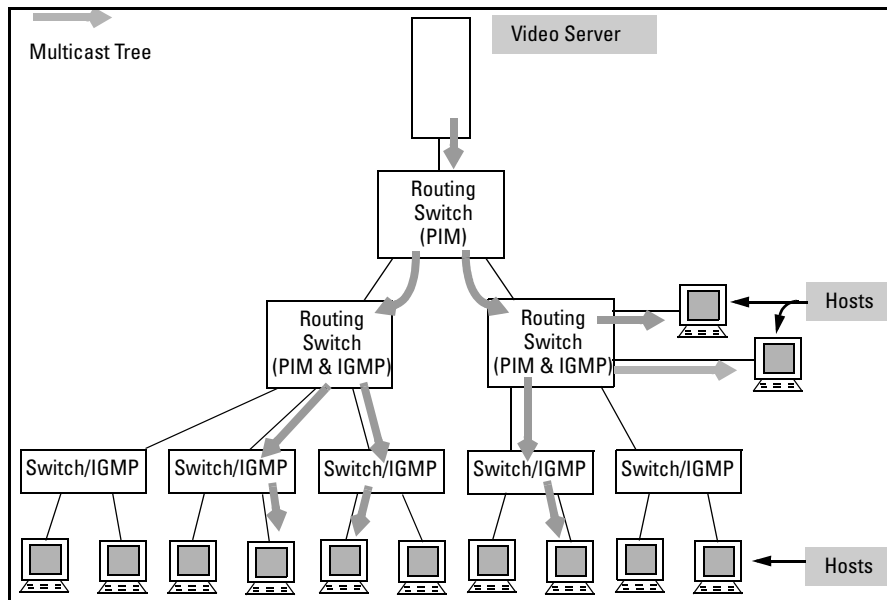


Figure 3-1. Example of Multicast “Tree” for a Given Flow

When the routing switch detects a new multicast flow, it initially floods the traffic throughout the PIM-DM domain, and then prunes the traffic on the branches (network paths) where joins have not been received from individual hosts. This creates the ‘tree’ structure shown above. The routing switch maintains individual branches in the multicast tree as long as there is at least one host maintaining a membership in the multicast group. When all of the hosts in a particular VLAN drop out of the group, PIM-DM prunes that VLAN from the multicast tree. Similarly, if the routing switch detects a join from a host in a pruned VLAN, it adds that branch back into the tree.

Note

Where the multicast routers in a network use one or more multinatted VLANs, there must be at least one subnet common to all routers on the VLAN. This is necessary to provide a continuous forwarding path for the multicast traffic on the VLAN. Refer to the `[ip pim-dense [ip-addr < any | source-ip-address >]` command under “PIM VLAN (Interface) Configuration Context” on page 3-15.

Multicast Flow Management

This section provides details on how the routing switch manages forwarding and pruned flows. This information is useful when planning topologies to include multicast support and when viewing and interpreting the Show command output for PIM-DM features.

Initial Flood and Prune. As mentioned earlier, when a router running PIM-DM receives a new multicast flow, it initially floods the traffic to all downstream multicast routers. PIM-DM then prunes the traffic on paths to VLANs that have no host joins for that multicast address. (Note that PIM-DM does not re-forward traffic back to its source VLAN.)

Maintaining the Prune State. For a multicast group “X” on a given VLAN, when the last host belonging to group “X” leaves the group, PIM places that VLAN in a prune state, meaning the group “X” multicast traffic is blocked to that VLAN. The prune state remains until a host on the same VLAN issues a join for group “X”, in which case the router cancels the prune state and changes the flow to the forwarding state.

State Refresh Packets and Bandwidth Conservation. A multicast switch, if directly connected to a multicast source such as a video conferencing application, periodically transmits *state refresh* packets to downstream multicast routers. On routers that have pruned the multicast flow, the state refresh packets keep the pruned state alive. On routers that have been added to the network after the initial flooding and pruning of a multicast group, the state refresh packets inform the newly added router of the current state of that branch. This means that if all multicast routers in a network support the state refresh packet, then the multicast router directly connected to the multicast source performs only one flood-prune cycle to the edge of the network when a new flow (multicast group) is introduced, and preserves bandwidth for other uses. Note, however, that some vendors’ multicast routers do not offer the state refresh feature. In this case, PIM-DM must periodically advertise an active multicast group to these devices by repeating the flood/prune cycle on the paths to such routers. For better traffic management in multicast-intensive networks where some multicast routers do not offer the state refresh feature, you may want to group such routers where the increased bandwidth usage will have the least effect on overall network performance.

PIM-DM (Dense Mode) PIM-DM Operation

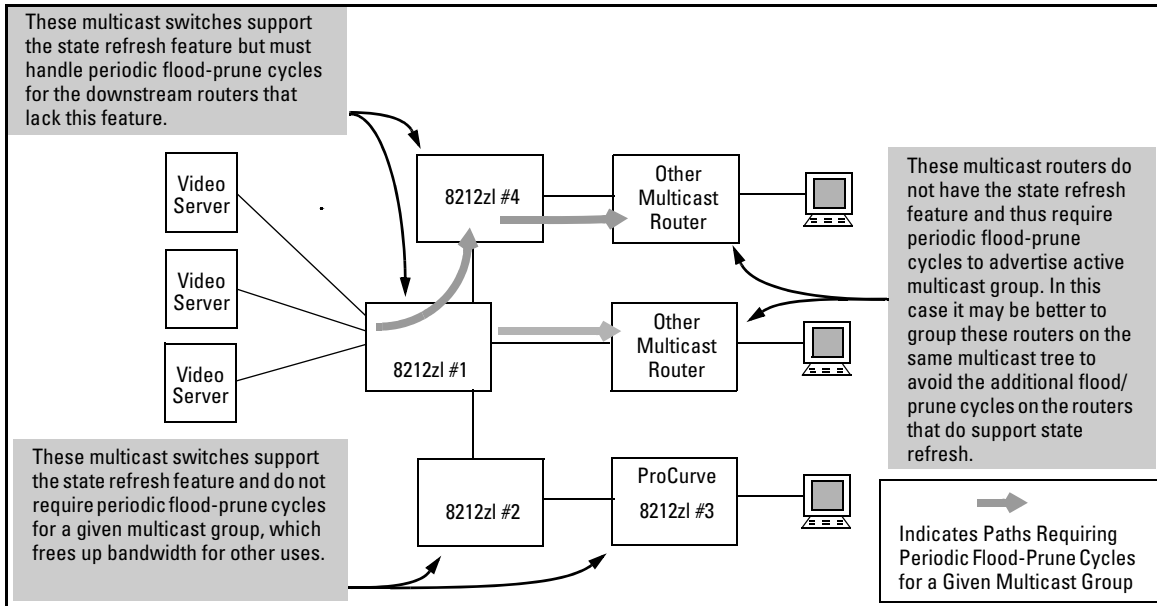


Figure 3-2. Example of Bandwidth Conservation in Switches with PIM-DM State Refresh

General Configuration Elements

The configured elements PIM-DM requires are:

1. IP routing enabled on all routing switches you want to carry routed multicast traffic.
2. Configure the routing method(s) needed to reach the interfaces (VLANs) on which you want multicast traffic available for hosts in your network:
 - Enable RIP or OSPF at both the global and VLAN levels on the routers where there are connected hosts that may issue multicast joins.
 - Configure static routes to and from the destination subnets.
3. Enable IP multicast routing.
4. For each VLAN on which there are hosts that you want to join multicast groups, enable IGMP on that VLAN. Repeat this action on every switch and router belonging to the VLAN.
5. Enable PIM-DM at the global level on the routing switch and on the VLANs where you want to allow routed multicast traffic.

Note

When you initially enable PIM-DM, ProCurve recommends that you leave the PIM-DM configuration parameters at their default settings. You can then assess performance and make configuration changes where a need appears.

Terminology

Flow: Multicast traffic moving between a unicast source and a multicast group. One S/G pair is counted as a single flow, regardless of the number of hosts belonging to the related multicast group.

Host: A client device that requests multicast traffic by transmitting IGMP “joins” for a specific multicast group, such as a video conferencing application.

MRT (Multicast Routing Table). The routing switch creates this table internally to maintain data on each multicast group it supports. The Show commands described later in this chapter display MRT data managed in this table.

Multicast Address: In IP multicast traffic on the switch, this is a single IP address that can be used by a group of related or unrelated clients wanting the same data. A single S/G pair consists of unicast source address and a multicast group address. Sometimes termed a “multicast group address”. See also “Source” and “S/G Pair”.

Multicast Routing: A method for transmitting multicast datagrams from a source in one IP network to a multicast address in one or more other IP networks.

PIM Neighbor: On a routing switch configured for PIM operation, a PIM neighbor is another PIM-configured routing switch or router that is either directly connected to the first routing switch or connected through networked switches and/or hubs.

Prune: To eliminate branches of a multicast tree that have no hosts sending joins to request or maintain membership in that particular multicast group.

S/G Pair: The unicast address of the server transmitting the multicast traffic and the multicast address to which the server is transmitting the traffic.

Source (S): In IP multicast traffic on the switch, the source (S) is the unicast address of the server transmitting the multicast traffic. A single S/G pair consists of unicast source address and a multicast group address. See also “S/G Pair”.

PIM-DM Operating Rules

- The routing switch supports 2048 multicast flows in hardware. (For more on this topic, refer to “Flow Capacity” on page 3-36.)
- The multicast routing table (MRT) that PIM-DM creates allows up to 128 outbound VLANs, meaning that at any given time, PIM-DM supports multicast routing across 128 VLANs.
- The routing switch allows one instance of PIM per VLAN. Thus, in networks using multinetted VLANs, all routers on a given VLAN intended to route multicast packets must have a least one common subnet on that VLAN. Thus, in the case of multinetting, you must select one subnet on the multinetted VLAN to use for multicast routing. To facilitate this, the routing switch provides a command for specifying which IP address PIM will use on each VLAN.

Configuring PIM-DM

Command	Page
PIM Global Context Commands	
[no] ip multicast-routing	3-12
[no] router pim	3-12
state-refresh	3-13
trap	3-13
PIM Interface Context Commands	
[no] ip pim-dense	3-15
[ip-addr < any <i>source-ip-address</i> >]	3-15
[hello-interval]	3-15
[hello-delay]	3-16
[graft-retry-interval]	3-16
[max-graft-retries]	3-17
[lan-prune-delay]	3-17
[propagation-delay]	3-18
[override-delay]	3-18
[ttl-threshold]	3-19

PIM-DM requires configuration on both the global level and on the VLAN (interface) level. The recommended configuration order is:

1. Enable IGMP on all VLANs where hosts may join a multicast group.
2. Enable the following at the global level.
 - IP routing
 - IP multicast routing
 - Router PIM and any non-default, global PIM settings you want to apply
 - Router RIP, Router OSPF, and/or a static route
3. If you selected RIP or OSPF in step 2, then on each VLAN where you want multicast routing to operate, enable the same option.
4. Enable the following in each VLAN context where you want multicast routing to operate:
 - IP RIP or IP OSPF
 - IP PIM
 - Any non-default, VLAN-level IP PIM settings you want to apply

Global and PIM Configuration Contexts

Note

PIM-DM operation requires a routing protocol enabled on the routing switch. You can use RIP, OSPF, and/or static routing. The examples in this section use RIP. For more on these topics, refer to Chapter 5, “IP Routing Features” in this guide.

Syntax: [no] ip multicast-routing

Enables or disables IP multicast routing on the routing switch. IP routing must be enabled. (Default: Disabled.)

Syntax: [no] router pim

Enables or disables PIM at the global level and places the CLI in the PIM context. IP routing must be enabled first. (Default: Disabled.)

Syntax: router pim state-refresh < 10 - 300 >

Executed in the PIM context, this command sets the interval in seconds between successive State Refresh messages originated by the routing switch. Note that only the routing switch connected directly to the unicast source initiates state-refresh packets. All other PIM routers in the network only propagate these state-refresh packets. (Range: 10 - 300 seconds; Default: 60 seconds)

Syntax: [no] router pim trap < all | neighbor-loss | hardware-mrt-full | software-mrt-full >

Executed in the PIM context, this command enables and disables these PIM SNMP traps:

all — Enable/Disable all PIM notification traps.

neighbor-loss — Enable/Disable the notification trap sent when the timer for a multicast router neighbor expires and the switch has no other multicast router neighbors on the same VLAN with a lower IP address. (Default: Disabled.)

hardware-mrt-full — Enable/Disable notification trap when the hardware multicast routing table (MRT) is full (1023 active flows). In this state, any additional flows are handled by the software MRT, which increases processing time for the affected flows. (Default: Disabled.)

software-mrt-full — Enable/Disable notification trap when the routing switch's software multicast routing table is full (that is, when routing resources for active flows are exhausted). (Default: Disabled.) Note that in this state, the routing switch does not accept any additional flows.

Example of Configuring PIM in the Global and PIM Contexts. In figure 3-2 on page 3-8, the “#1” routing switch is directly connected to the multicast sources for the network. In this case, suppose that you want to do the following:

- Reduce the state-refresh time from the default 60 seconds to 30 seconds. Note that the routing switch transmits state-refresh packets only if it is directly connected to the multicast source.
- Configure an SNMP trap to notify your network management station if the routing switch's hardware multicast routing table becomes filled to the maximum of 1023 active flows.

PIM-DM (Dense Mode)

Configuring PIM-DM

To configure global-level PIM operation for the “8212zl #1” routing switch, you would use the commands shown in figure 3-3, below.

```
ProCurve(config)# ip routing
ProCurve(config)# ip multicast-routing
ProCurve(config)# router rip
ProCurve(rip)# exit
ProCurve(config)# router pim
ProCurve(pim)# state-refresh 45
ProCurve(pim)# trap hardware-mrt-full
ProCurve(pim)# write mem
ProCurve(pim)# exit
```

← Enables IP routing.
← Enables multicast routing.
← Enables RIP.
← Exits from the RIP context.
← Enables PIM and enters the PIM context.
← Configures a non-default State Refresh timer.
← Sets an SNMP trap to notify an SNMP management station if the hardware multicast routing table fills with active flows.

Using **show run** displays the configuration changes resulting from the above commands.

```
ProCurve(config)# show run
Running configuration:
; J9091A Configuration Editor; Created on release #K.12.XX
hostname "ProCurve"
module 1 type J8702A
module 2 type J8702A
ip routing
snmp-server community "public" Unrestricted
vlan 1
.
.
.
vlan 29
.
.
.
vlan 25
  name "VLAN25"
  untagged A20-A24
  ip address 10.38.10.1 255.255.255.0
  exit
ip multicast-routing
router rip
  exit
router pim
| state-refresh 45
| trap hardware-mrt-full
| exit
```

Figure 3-3. Example of Configuring PIM-DM on a Routing Switch at the Global Level

After configuring the global-level PIM operation on a routing switch, go to the device's VLAN context level for each VLAN you want to include in your multicast routing domain. (Refer to “PIM VLAN (Interface) Configuration Context”, below.

PIM VLAN (Interface) Configuration Context

Syntax: [no] ip pim-dense
[no] vlan < vid > ip pim

*Enables multicast routing on the VLAN interface to which the CLI is currently set. The **no** form disables PIM on the VLAN. Default: Disabled.*

Syntax: [no] ip pim-dense [ip-addr < any | source-ip-address >]
[no] vlan < vid > ip pim-dense [ip-addr < any | source-ip-address >]

*In networks using multinetted VLANs, all routers on a given VLAN intended to route multicast packets must have a least one common subnet on that VLAN. Use this command when the VLAN is configured with multiple IP addresses (multinetting) to specify the IP address to use as the source address for PIM protocol packets outbound on the VLAN. Use < ip-address > to designate a single subnet in cases where multicast routers on the same multinetted VLAN are not configured with identical sets of subnet IP addresses . Use < all > if the multinetted VLAN is configured with the same set of subnet addresses. (**Default:** The Primary VLAN.)*

Syntax: ip pim-dense [hello-interval < 5 - 30 >]
vlan < vid > ip pim-dense [hello-interval < 5 - 30 >]

*Changes the frequency at which the routing switch transmits PIM “Hello” messages on the current VLAN. The routing switch uses “Hello” packets to inform neighboring routers of its presence. The routing switch also uses this setting to compute the **Hello Hold Time**, which is included in Hello packets sent to neighbor routers. **Hello Hold Time** tells neighbor routers how long to wait for the next Hello packet from the routing switch. If another packet does not arrive within that time, the router removes the neighbor adjacency on that VLAN from the routing table, which removes any flows running on that interface. Shortening the Hello interval reduces the Hello Hold Time. This has the effect of changing how quickly other routers will stop sending traffic to the routing switch if they do not receive a new Hello packet when expected.*

*For example, if multiple routers are connected to the same VLAN and the routing switch requests multicast traffic, all routers on the VLAN receive that traffic. (Those which have pruned the traffic will drop it when they receive it.) If the upstream router loses contact with the routing switch receiving the multicast traffic (that is, fails to receive a Hello packet when expected), then the shorter Hello Interval causes it to stop transmitting multicast traffic onto the VLAN sooner, resulting in less unnecessary bandwidth usage. Not used with the **no** form of the **ip pim-dense** command.*

Syntax: ip pim-dense [hello-delay < 0 - 5 >]
vlan < vid > ip pim-dense [hello-delay < 0 - 5 >]

*Changes the maximum time in seconds before the routing switch actually transmits the initial PIM Hello message on the current VLAN. In cases where a new VLAN activates with connections to multiple routers, if all of the connected routers sent Hello packets at the same time, then the receiving router could become momentarily overloaded. This value randomizes the transmission delay to a time between **0** and the **hello delay** setting. Using “**0**” means no delay. After the routing switch sends the initial Hello Packet to a newly detected VLAN interface, it sends subsequent Hello packets according to the current **Hello Interval** setting. Not used with the **no** form of the **ip pim-dense** command. Default: 5 seconds.*

Syntax: ip pim-dense [graft-retry-interval < 1-10 >]
vlan < vid > ip pim-dense [graft-retry-interval < 1-10 >]

*Graft packets result when a downstream router transmits a request to join a flow. The upstream router responds with a graft acknowledgment packet. If the Graft Ack is not received within the time period of the **graft-retry-interval**, it resends the graft packet. This command changes the interval (in seconds) the routing switch waits for the Graft Ack (acknowledgement) from another router before resending the Graft request. Not used with the **no** form of the **ip pim-dense** command. (Default: 3 seconds.)*

Syntax: ip pim-dense [max-graft-retries < 1 - 10 >
vlan < vid > ip pim-dense [max-graft-retries < 1 - 10 >

*Changes the number of times the routing switch will retry sending the same graft packet to join a flow. If a Graft Ack response is not received after the specified number of retries, the routing switch ceases trying to join the flow. In this case the flow is removed until either a state refresh from upstream re-initiates the flow or an upstream router floods the flow. Increasing this value helps to improve multicast reliability. Not used with the **no** form of the **ip pim-dense** command. (Default: 3 attempts.)*

Syntax: ip pim-dense [lan-prune-delay]
vlan < vid > ip pim-dense [lan-prune-delay]

*Enables the LAN Prune Delay option on the current VLAN. With **lan-prune-delay** enabled, the routing switch informs downstream neighbors how long it will wait before pruning a flow after receiving a prune request. Other, downstream routers on the same VLAN must send a Join to override the prune before the **lan-prune-delay** time if they want the flow to continue. This prompts any downstream neighbors with hosts continuing to belong to the flow to reply with a Join. If no joins are received after the **lan-prune-delay** period, the routing switch prunes the flow. The **propagation-delay** and **override-interval** settings (below) determine the **lan-prune-delay** setting.*

*Uses the **no** form of the **ip pim-dense** command to disable the LAN Prune Delay option. (Default: Enabled.)*

Syntax: ip pim-dense [propagation-delay < 250-2000 >]
vlan < vid > ip pim-dense [propagation-delay < 250-2000 >]

ip pim-dense [override-interval < 500 - 6000 >]
vlan < vid > ip pim-dense [override-interval < 500 - 6000 >]

*A routing switch sharing a VLAN with other multicast routers uses these two values to compute the **lan-prune-delay** setting (above) for how long to wait for a PIM-DM join after receiving a prune packet from downstream for a particular multicast group. For example, a network may have multiple routing switches sharing VLAN "X". When an upstream routing switch initially floods traffic from multicast group "X" to VLAN "Y", if one of the routing switches on VLAN "Y" does not want this traffic it issues a prune response to the upstream neighbor. The upstream neighbor then goes into a "prune pending" state for group "X" on VLAN "Y". (During this period, the upstream neighbor continues to forward the traffic.) During the "pending" period, another routing switch on VLAN "Y" can send a group "X" Join to the upstream neighbor. If this happens, the upstream neighbor drops the "prune pending" state and continues forwarding the traffic. But if no routers on the VLAN send a Join, then the upstream router prunes group "X" from VLAN "Y" when the **lan-prune-delay** timer expires. (Defaults: **propagation-delay** = 500 milliseconds; **override-interval** = 2500 milliseconds.)*

Syntax: ip pim-dense [ttl-threshold < 0 - 255 >]
vlan < vid > ip pim-dense [ttl-threshold < 0 - 255 >]

Sets the multicast datagram time-to-live (router hop-count) threshold for the VLAN. Any IP multicast datagrams or state refresh packets with a TTL less than this threshold will not be forwarded out the interface. The default value of 0 means all multicast packets are forwarded out the interface.

*This parameter provides a method for containing multicast traffic within a network, or even within specific areas of a network. Initially, the multicast traffic source sets a TTL value in the packets it transmits. Each time one of these packets passes through a multicast routing device, the TTL setting decrements by 1. If the packet arrives with a TTL lower than the **mroute ttl-threshold**, the routing switch does not forward the packet. Changing this parameter on a routing switch requires knowledge of the TTL setting of incoming multicast packets. A value that is too high can allow multicast traffic to go beyond your internal network. A value that is too low may prevent some intended hosts from receiving the desired multicast traffic. (Default: 0 — forwards multicast traffic regardless of packet TTL setting.)*

Example of Configuring PIM-DM Operation at the VLAN Level. The network in figure 3-4 uses VLAN 25 for multicast traffic. However, this VLAN is multinetted and there is only one subnet (10.38.10.x) in VLAN 25 that is common to all three routing switches. Thus, when configuring VLAN 25 on these routing switches to perform multicast routing, it is necessary to use **ip pim-dense < source-ip-address >** to designate the common subnet as the source address for outbound multicast traffic on VLAN 25. (If only identical subnets were present in the multinetted VLAN 25 configuration on all three devices, then the **ip pim-dense ip-addr any** command would be used instead.) Note that the other VLANs in the network are not multinetted and therefore do not require the **ip pim-dense ip-addr < any | source-ip-address >** option.

For this example, assume that the VLANs and IP addressing are already configured on the routing switch.

PIM-DM (Dense Mode)
Configuring PIM-DM

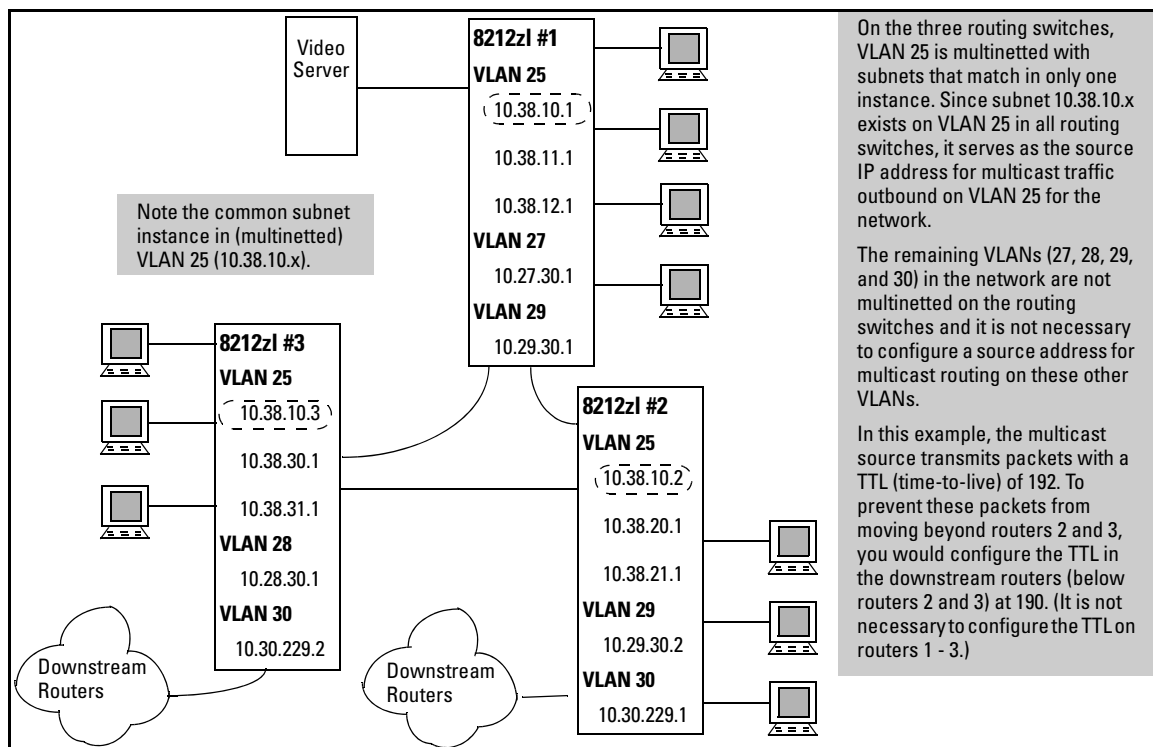


Figure 3-4. Example of a Multicast Network with a Multinetted VLAN

Figure 3-5 illustrates the steps for configuring multicast routing at the VLAN level for the 8212zl switch #1 shown in figure 3-4.

```
ProCurve(config)# vlan 25
ProCurve(vlan-25)# ip igmp
ProCurve(vlan-25)# ip rip
ProCurve(vlan-25)# ip pim-dense ip-addr 10.38.10.1
ProCurve(vlan-25-pim-dense)# vlan 27
ProCurve(vlan-27)# ip igmp
ProCurve(vlan-27)# ip rip
ProCurve(vlan-27)# ip pim-dense
ProCurve(vlan-27-pim-dense)# vlan 29
ProCurve(vlan-29)# ip igmp
ProCurve(vlan-29)# ip rip
ProCurve(vlan-29)# ip pim-dense
ProCurve(vlan-29-pim-dense)# write mem
ProCurve(vlan-29-pim-dense)# exit
ProCurve(vlan-29)# exit
```

Figure 3-5. VLAN-Level Configuration Steps for PIM-DM on Routing Switch #1

```

ProCurve(config)# show run
...
ip routing
...
vlan 29
  name "VLAN29"
  untagged A11-A15,A17
  ip address 10.29.30.1 255.255.255.0
  ip igmp
  exit
vlan 25
  name "VLAN25"
  [ untagged A20-A24 ]
  | ip address 10.38.10.1 255.255.255.0 |
  | ip address 10.38.11.1 255.255.255.0 |
  | ip address 10.38.12.1 255.255.255.0 |
  [ ip igmp ]
  exit
vlan 27
  name "VLAN27"
  untagged A6-A10,A18
  ip address 10.27.30.1 255.255.255.0
  ip igmp
  exit
[ ip multicast-routing ]
[ router rip ]
[   exit ]
[ router pim ]
[   state-refresh 45 ]
[   trap hardware-mrt-full ]
[   exit ]
vlan 25
  [ ip rip 10.38.10.1 ]
  [ ip rip 10.38.11.1 ]
  | ip pim-dense |
  | ip-addr 10.38.10.1 |
  |   exit ]
vlan 27
  [ ip rip 10.27.30.1 ]
  [ ip pim-dense ]
  [   ip-addr any ]
  [   exit ]
vlan 29
  [ ip rip 10.29.30.1 ]
  [ ip pim-dense ]
  [   ip-addr any ]

```

← Enables IP routing; required for multicast routing.

↙ Multinetting and IGMP enabled in VLAN 25.

↖ Multicast Routing Configuration for Global Level..

↖ Indicates the source-IP-address for multicast packets forwarded on this VLAN.

↖ Multicast Routing Configuration for VLAN 25.

↖ Multicast Routing Configurations for VLANs 27 and 29.

Note: Dashed lines indicate configuration settings affecting multicast routing.

Figure 3-6. The Multicast Routing Configuration on Switch #1 in Figure 3-4 (Page 3-20)

Displaying PIM Data and Configuration Settings

Command	Page
show ip mroute	3-23
[interface < vid >]	3-24
[< multicast-ip-addr > < source-ip-addr >]	3-25
show ip pim	3-28
[interface	3-29
[< vid >]]	3-30
[mroute	3-31
[< multicast-group-address > < multicast-source-address >]]	3-32
neighbor	3-34
[< ip-address >]	3-35

Displaying PIM Route Data

Syntax: show ip mroute

Without parameters, lists all VLANs actively forwarding routed, multicast traffic.

Group Address: *The multicast address of the specific multicast group (flow).*

Source Address: *The unicast address of the multicast group source.*

Neighbor: *The IP address of the upstream multicast router interface (VLAN) from which the multicast traffic is coming. A blank field for a given multicast group indicates that the multicast server is directly connected to the routing switch.*

VLAN: *The interface on which the multicast traffic is moving.*

For example, the next figure displays the show ip route output on the “8212zl #2” routing switch in figure 3-4 on page 3-20. This case illustrates two multicast groups from the same multicast server source.

```
ProCurve(config)# show ip mroute
```

Group Address	Source Address	Neighbor	VLAN
239.255.255.1	10.27.30.2	10.29.30.1	29
239.255.255.5	10.27.30.2	10.29.30.1	29

IP Multicast Route Entries
Total number of entries : 2

Indicates the upstream multicast router interface (VLAN) from which the multicast traffic is coming.

Figure 3-7. Example Showing the Route Entry Data on the “#2” Routing Switch in Figure 3-4 on Page 3-20

Syntax: show ip mroute [interface < vid >]

Lists these settings:

VLAN: *The VID specified in the command.*

Protocol Identity: *PIM-DM only.*

TTL: *The time-to-live threshold for packets forwarded through this VLAN. When configured, the routing switch drops multicast packets having a TTL lower than this value. (When a packet arrives, the routing switch decrements it's TTL by 1, then compares the decremented packet TTL to the value set by this command.) A TTL Threshold setting of 0 (the default) means all multicast packets are forwarded regardless of the TTL value they carry. A multicast packet must have a TTL greater than 1 when it arrives at the routing switch. Otherwise the routing switch drops the packet instead of forwarding it on the VLAN.*

```
ProCurve(config)# show ip mroute interface 29

IP Multicast Interface

VLAN      : 29
Protocol  : PIM-DM

TTL Threshold : 0
```

Figure 3-8. Example of the Above Command on Routing Switch “#2” in Figure 3-4 on Page 3-20

Syntax: show ip mroute [< multicast-ip-addr > < source-ip-addr >]

Lists the following data for the specified flow (multicast group):

Group Address: *The multicast group IP address for the current group.*

Source Address: *The multicast source address < source-ip-addr > for the current group.*

Source Mask: *The subnet mask applied to the multicast source address < source-ip-addr >.*

Neighbor: *Lists the IP address of the upstream next-hop router running PIM-DM; that is, the router from which the routing switch is receiving datagrams for the current multicast group. This value is 0.0.0.0 if the routing switch has not detected the upstream next-hop router's IP address. This field is empty if the multicast server is directly connected to the routing switch.*

VLAN: *Lists the VLAN ID (VID) on which the routing switch received the specified multicast flow.*

Up Time (Sec): *The elapsed time in seconds since the routing switch learned the information for the current instance of the indicated multicast flow.*

Expiry Time (Sec): *Indicates the remaining time in seconds before the routing switch ages-out the current flow (group membership). This value decrements until:*

- *Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.)*
- *Reset by a new flow for the current multicast group on the VLAN.*
- *The timer expires (reaches 0). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.*

Multicast Routing Protocol: *Identifies the multicast routing protocol through which the current flow was learned.*

Unicast Routing Protocol: *Identifies the routing protocol through which the routing switch learned the upstream interface for the current multicast flow. The listed protocol will be either **RIP**, **OSPF**, or **Static Route**.*

Downstream Interfaces:

VLAN: *Lists the VID of the VLAN that the routing switch is using to send the outbound packets of the current multicast flow to the next-hop router.*

State: *Indicates whether the outbound VLAN and next-hop router for the current multicast flow are receiving datagrams.*

- **Pruned:** *The routing switch has not detected any joins from the current multicast flow and is not currently forwarding datagrams in the current VLAN.*
- **Forwarding:** *The routing switch has received a join for the current multicast flow and is forwarding datagrams in the current VLAN.*

Up Time (Sec): *Indicates the elapsed time in seconds since the routing switch learned the displayed information about the current multicast flow.*

Expiry Time: *Shows the remaining time in seconds until the Next-Hop routing switch ages-out the current flow (group membership) on the indicated VLAN. Includes the date calculated for the age-out event. This value decrements until:*

- *Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.*
- *Reset by a new flow for the current multicast group on the VLAN.*
- *The timer expires (reaches 0). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.*

Note that the “Next-Hop routing switch” is the next multicast routing switch in the path from the current multicast routing switch to the source for the displayed multicast flow.

```
ProCurve(config)# show ip mroute 239.255.255.5
10.27.30.2

IP Multicast Route Entry

Group Address : 239.255.255.5
Source Address : 10.27.30.2
Source Mask : 255.255.255.0
Neighbor : 10.30.229.1
VLAN : 27
Up time (sec) : 408
Expire Time (sec) : 150

Multicast Routing Protocol : PIM-DM
Unicast Routing Protocol : rip

Downstream Interfaces

VLAN State Up time (sec) Expire Time (sec)
-----
28 pruned 408 98
```

A blank **Neighbor** field indicates that the multicast server is directly connected to the routing switch.

Figure 3-9. Example Output for Routing Switch “#1” in Figure 3-4 on Page 3-20

Displaying PIM Status

Syntax: show ip pim

Displays PIM status and global parameters.

PIM Status: Shows either **enabled** or **disabled**.

State Refresh Interval (sec): A PIM routing switch originates state refresh messages to inform its neighbors of the active flows it is currently routing. This updates the current flow data on PIM routers that join or rejoin a multicast network after the initial flood and prune. This enables hosts on such routers to join a multicast group without having to wait for a “flood and prune” cycle. PIM routers having the state refresh capability can eliminate all but an initial flood and prune cycle. PIM routers without this capability periodically trigger a flood and prune cycle on the path between the PIM router and the multicast source. (Range: 10 - 300 seconds; Default: 60 seconds.)

Join/Prune Interval (sec): Indicates the frequency with which the router transmits join and prune messages for the multicast groups the router is forwarding.

SPT Threshold: This is the “Shortest Path Tree Threshold” used with PIM-SM. For more information, refer to “Displaying the Current PIM status and Global Configuration” on page 4-51.

Traps: Enables the following SNMP traps:

- **neighbor-loss:** Sends a trap if a neighbor router is lost.
- **hardware-mrt-full:** Sends a trap if the hardware multicast router (MRT) table is full (2,048 active flows).
- **software-mrt-full:** Sends a trap if the software multicast router (MRT) table is full (511 active flows). This can occur only if the hardware MRT is also full.
- **all:** Enables all of the above traps.

```
ProCurve(config)# show ip pim

PIM Global Parameters

PIM Status                : enabled
State Refresh Interval (sec) : 60
Join/Prune Interval (sec)  : 60
SPT Threshold              : Enabled
Traps                      : hardware-mrt-full
```

Figure 3-10. Example Output for Routing Switch “#1” in Figure 3-4 on Page 3-20

Syntax: show ip pim [interface]

Lists the PIM interfaces (VLANs) currently configured in the routing switch.

VLAN: *Lists the VID of each VLAN configured on the switch to support PIM-DM.*

IP Address: *Lists the IP addresses of the PIM interfaces (VLANs).*

Mode: *Shows dense only.*

```
ProCurve(config)# show ip pim interface

PIM Interfaces

VLAN IP Address      Mode
----  -
25   10.38.10.1         dense
27   10.27.30.1         dense
29   10.29.30.1         dense
```

Figure 3-11. Example Output for Routing Switch “#1” in Figure 3-4 on Page 3-20

PIM-DM (Dense Mode)

Displaying PIM Data and Configuration Settings

Syntax: show ip pim [interface [< vid >]]

Displays the current configuration for the specified VLAN (PIM interface). Refer to table 3-1, below.

```
ProCurve(config)# show ip pim interface 29

PIM Interface

VLAN          : 29
IP Address    : 10.29.30.1
Mode          : dense

Hello Interval (sec) : 30
Hello Delay (sec)   : 5

Graft Retry Interval(sec) : 3
Max Graft Retries      : 2
Override Interval (msec) : 2500      Lan Prune Delay           : Yes
Propagation Delay (msec) : 500        Lan Delay Enabled        : No
SR TTL Threshold      : 2          State Refresh Capable    : No
```

Figure 3-12. Example Output for Routing Switch “#1” in Figure 3-4 on Page 3-20

Table 3-1. PIM Interface Configuration Settings

Field	Default	Control Command
VLAN	n/a	vlan < vid > ip pim-dense
IP	n/a	vlan < vid > ip pim-dense < any ip-addr >
Mode	dense	PIM-Dense or PIM-Sparse
Hello Interval (sec)	30	ip pim-dense hello interval < 5 - 30 >
Hello Hold Time	105	The routing switch computes this value from the current “Hello Interval” and includes it in the “Hello” packets the routing switch sends to neighbor routers. Neighbor routers use this value to determine how long to wait for another Hello packet from the routing switch. Refer to the description of the Hello Interval on page 3-15.
Hello Delay	5	vlan < vid > ip pim-dense hello delay < 0 - 5 >
Graft Retry Interval (sec)	3	vlan < vid > ip pim-dense graft-retry-interval < 1 - 10 >
Max Graft Retries	2	vlan < vid > ip pim-dense graft-retries < 1 - 10 >

Field	Default	Control Command
Override Interval (msec)	2500	vlan < vid > ip pim-dense override-interval < 500 - 6000 >
Propagation Delay (msec)	500	vlan < vid > ip pim-dense propagation-delay < 250-2000 >
SR TTL Threshold (router hops)	0	vlan < vid > ip pim-dense ttl-threshold < 0 - 255 >
LAN Prune Delay	Yes	vlan < vid > ip pim-dense lan-prune-delay
LAN Delay Enabled	No	Shows Yes if all multicast routers on the current VLAN interface enabled LAN-prune-delay. Otherwise shows No .
State Refresh Capable	n/a	Indicates whether the VLAN responds to state refresh packets. The VLAN connected to the multicast source does not receive state refresh packets and thus is not state-refresh capable. Downstream VLANs in the switches covered in this guide are state-refresh capable.

Syntax: show ip pim [mroute]

Shows PIM-specific information from the IP multicast routing table (IP MRT). When invoked without parameters, lists all PIM entries currently in the routing switch's IP MRT.

Group Address: Lists the multicast group addresses currently active on the routing switch.

Source Address: Lists the multicast source address for each Group Address.

Metric: Indicates the path cost upstream to the multicast source. Used when multiple multicast routers contend to determine the best path to the multicast source. The lower the value, the better the path. This value is set to **0** (zero) for directly connected routes.

Metric Pref: Used when multiple multicast routers contend to determine the path to the multicast source. When this value differs between routers, PIM selects the router with the lowest value. If **Metric Pref** is the same between contending multicast routers, then PIM selects the router with the lowest **Metric** value to provide the path for the specified multicast traffic. This value is set to **0** (zero) for directly connected routes.

(Metric Pref is based on the routing protocol in use: RIP, OSPF, or static routing. Also, different vendors may assign different values for this setting.)

This output shows the routing switch is receiving two multicast groups from an upstream device at 10.27.30.2. The "0" metric shows that the routing switch is directly connected to the multicast source.

```
ProCurve(config)# show ip pim mroute
```

PIM Route Entries

Group Address	Source Address	Metric	Metric Pref
239.255.255.1	10.27.30.2	0	0
239.255.255.5	10.27.30.2	0	0

Figure 3-13. Example Showing a Routing Switch Detecting two Multicast Groups from a Directly Connected Multicast Server

Syntax: show ip pim [mroute [< multicast-group-address >
< multicast-source-address >]]

Displays the PIM route entry information for the specified multicast group (flow):

Group Address: Lists the specified multicast group address.

Source Address: Lists the specified multicast source address.

Source Mask: Lists the network mask for the multicast source address.

Metric: Lists the number of multicast router hops to the source address.

Metric: Indicates the path cost upstream to the multicast source. Used when multiple multicast routers contend to determine the best path to the multicast source. The lower the value, the better the path.

Metric Pref: Used when multiple multicast routers contend to determine the path to the multicast source. When this value differs between routers, PIM selects the router with the lowest value. If Metric Pref is the same between contending multicast routers, then PIM selects the router with the lowest **Metric** value to provide the path for the specified multicast traffic. (Different vendors assign differing values for this setting.)

Assert Timer: The time remaining until the routing switch ceases to wait for a response from another multicast router to negotiate the best path back to the multicast source. If this timer expires without a response from any contending multicast routers, then the routing switch assumes it is the best path, and the specified multicast group traffic will flow through the routing switch.

DownStream Interfaces:

- **VLAN:** Lists the VID of the destination VLAN on the next-hop multicast router.
- **Prune Reason:** *Identifies the reason for pruning the flow to the indicated VLAN:*
 - **Prune:** *A neighbor multicast router has sent a prune request.*
 - **Assert:** *Another multicast router connected to the same VLAN has been elected to provide the path for the specified multicast group traffic.*
 - **Other:** *Used where the VLAN is in the pruned state for any reason other than the above two reasons (such as no neighbors exist and no directly connected hosts have done joins).*

```
ProCurve(config)# show ip pim mroute 239.255.255.1 10.27.30.2

PIM Route Entry

Group Address   : 239.255.255.1
Source Address  : 10.27.30.2
Source Mask     : 255.255.255.0

Metric         : 3
Metric Pref    : 120
Assert Timer   : 0

DownStream Interfaces

VLAN Prune Reason
---- -
28   prune
```

This example displays the MRT data on the first of the two multicast groups shown in figure 3-13 on page 3-32.

Figure 3-14. Example From Routing Switch “#1” in Figure 3-4 on Page 3-20 Showing a Multicast Group from a Directly Connected Source

Syntax: show ip pim [neighbor]

Lists PIM neighbor information for all PIM neighbors connected to the routing switch:

IP Address: *Lists the IP address of a neighbor multicast router.*

VLAN: *Lists the VLAN through which the routing switch connects to the indicated neighbor.*

Up Time: *Shows the elapsed time during which the neighbor has maintained a PIM route to the routing switch.*

Expire Time: *Indicates how long before the routing switch ages-out the current flow (group membership). This value decrements until:*

- *Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.*
- *Reset by a new flow for the current multicast group on the VLAN.*

The timer expires (reaches 0). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.

If the IP-ADDR is specified then detailed information for the specified neighbor is shown.

```
This example simulates output from Routing Switch "#1" in Figure 3-4 on Page 3-20. The data identifies the first downstream neighbor ("Routing Switch #2").
ProCurve(config)# show ip pim neighbor

PIM Neighbors

IP Address          VLAN Up Time (sec)      Expire Time (sec)
-----
10.29.30.2          29   196                    89
```

Figure 3-15. Example of PIM Neighbor Output

Syntax: show ip pim [neighbor [< ip-address >]]

Lists the same information as show ip pim neighbor (page 3-34) for the specified PIM neighbor:

```
This example simulates output from Routing Switch "#1" in Figure 3-4 on Page 3-20. The data is from the first downstream neighbor (Routing Switch "#2").  
ProCurve(config)# show ip pim neighbor 10.29.30.2  
  
PIM Neighbor  
  
IP Address   : 10.29.30.2  
VLAN        : 29  
  
Up Time (sec)      : 26  
Expire Time (sec) : 79
```

Figure 3-16. Example From Routing Switch "#1" in Figure 3-4 on Page 3-20 Showing a Specific Neighbor (Routing Switch "#2")

Operating Notes

PIM Routers without State Refresh Messaging Capability. A PIM router without a state refresh messaging capability learns of currently active flows in a multicast network through periodic flood and prune cycles on the path back to the source. The switches covered in this guide sense downstream multicast routers that do not have the state refresh capability and will periodically flood active multicast groups to these devices. This periodic flooding is not necessary if all of the downstream multicast routers are switches covered in this guide. (The ProCurve Routing Switch Series 9300 and the routers offered by some other vendors do not offer the state refresh capability.)

Flow Capacity. The routing switch provides an ample multicast environment, supporting 1022 multicast flows in hardware across a maximum of 64 VLANs. (A flow comprises a unicast source address and a multicast group address, regardless of the number of active hosts belonging to the multicast group at any given time.) While the typical multicast environment should not normally exceed 1022 flows, the routing switch can support up to 978 additional flows in software, depending on available system resources. (Because the switch processes flows in hardware much faster than in software, you may notice slower processing times for flows occurring in software.) Also, while the routing switch can support up to 2,000 flows, the total demand on system resources from the combined use of more than 1,022 simultaneous flows, a high number of VLANs supporting multicast routing, and/or other, resource-intensive features can oversubscribe memory resources, which reduces the number of flows the routing switch can support in software. That is, the switch does not route flows in software that oversubscribe current memory resources. If the routing switch regularly exceeds the hardware limit of 1022 flows and begins routing flows in software, you may want to move some hosts that create multicast demand to another routing switch, or reduce the number of VLANs on the routing switch by moving some VLANs to another routing switch. Note that the routing switch generates a log message if it either routes a flow in software or drops a flow intended for software routing because memory is oversubscribed. (Refer to “Messages Related to PIM Operation” on page 3-38.)

IGMP Traffic High-Priority Disabled. Enabling IP multicast routing to support PIM-DM operation has the effect of disabling IGMP traffic high-priority, if configured. (Refer to “Configuring IGMP Traffic Priority” on page 2-11.)

ACLs and PIM. The switch allows ACL filtering on unicast addresses, but not on multicast addresses. Also, an ACL does not take effect on a flow if the flow began before the ACL was configured.

When To Enable IGMP on a VLAN. When PIM is enabled on a VLAN, it is not necessary to also enable IGMP unless there may be Joins occurring on that VLAN. But if IGMP is enabled on a VLAN, you must also enable PIM if you want that VLAN to participate in multicast routing.

IP Address Removed. If you remove the IP address for a VLAN, the switch automatically removes the PIM configuration for that VLAN.

Troubleshooting

Symptom: Noticeable slowdown in some multicast traffic. If the switch is supporting more than 1022 active flows. This generates the message `Unable to learn HW IP multicast groups, table FULL` in the Event Log because there is no room in the hardware Multicast Routing Table to add another Multicast Group. Software will route any multicast packets sent to multicast groups that are not in the hardware Multicast Routing Table, but it will be slower and packets may be dropped if the data rate is greater than 3000 packets per second. Refer to “Flow Capacity” on page 3-36.

Note that the PIM protocol uses one MRT entry for every IP multicast source/group pair that it is routing. An entry is not used if the multicast flow is bridged and not routed. Entries in this table are automatically aged out if they are unused for a period of time.

Heavy Memory Usage. Heavy use of PIM (many S/G flows over many VLANs) combined with other memory-intensive features, can oversubscribe memory resources and impact overall performance. If available memory is exceeded, the switch drops any new multicast flows and generates appropriate Event Log messages. Corrective actions can include reducing the number of VLANs on the switches covered in this guide by moving some VLANs to another device, free up system resources by disabling another, non-PIM feature, and/or moving some hosts to another device. For more information, refer to “Operating Notes” on page 3-35 and “Messages Related to PIM Operation” on page 3-38.

IPv4 Table Operation. The IPv4 table, which contains the active IP multicast addresses the switch is currently supporting, has 128k entries. However, the IPv4 table also contains IP host entries for every IP source or destination that the switch has learned, as well as ACL flow entries. Entries in this table are generally aged out if they are unused for 5 minutes or more.

Messages Related to PIM Operation

These messages appear in the Event Log and, if Syslog Debug is configured, in the designated Debug destinations.

Note

The *<counter>* value displayed at the end of each PIM Event Log message (and SNMP trap messages, if trap receivers are configured) indicates the number of times the switch has detected a recurring event since the last reboot. For more information, refer to “Using the Event Log To Identify Problem Sources” in the “Troubleshooting” appendix of the latest version of the *Management and Configuration Guide* for your switch. (The latest version of all ProCurve switch documentation is available on the ProCurve website at www.procurve.com)

Message	Meaning
<code><alpha-string> pkt, src IP<ip-addr> vid <vlan-id> (not a nbr) (<counter>)</code>	A PIM packet arrived from another router for which no neighbor was found. May indicate a misconfiguration between the sending and receiving router. May also occur if a connected router is disconnected, then reconnected.
Bad TTL in State Refresh pkt from IP <code><source-ip-addr> (<counter>)</code>	The switch detected a TTL of 0 (zero) in the PIM portion of a state refresh packet. (Note that this is not the IP TTL.)
Failed alloc of HW <code><alpha-str></code> for flow <code><multicast-address></code> , <code><source-address></code> (<code><dup-msg-cnt></code>)	There are more than 1022 active flows. The switch routes the excess through software, which processes traffic at a slower rate. If this will be an ongoing or chronic condition, transfer some of the flows to another router.
Failed to alloc a PIM <code><data-type></code> pkt (<code><counter></code>)	The router was unable to allocate memory for a PIM control packet. Router memory is oversubscribed. Reduce the number of VLANs or increase the hello delay and/or the override interval to reduce the number of simultaneous packet transmissions. Note that if the number of flows exceeds 1022, the excess flows are routed in software, which reduces the number of packet transmissions. In this case, reducing the number of flows by moving some clients to other routers can help.
Failed to initialize <code><text-str></code> as a call back routine (<code><counter></code>)	Indicates an internal error. Report the incident to your ProCurve customer care center and re-install the router software.
I/F configured with IP <code><ip-address></code> on vid <code><vlan-id></code> (<code><counter></code>)	Indicates that the interface (VLAN) has been configured with the indicated IP address. At boot-up or when an IP address is changed, the switch generates this message for each PIM-configured VLAN.

Message	Meaning
I/F removal with IP < <i>ip-addr</i> > on vid < <i>vlan-id</i> > (< <i>counter</i> >)	Indicates that a PIM interface (VLAN) has been removed from the router as a result of an IP address change or removal.
MCAST flow < <i>multicast-address</i> > < <i>source-address</i> > not rteing (rsc low) (< <i>counter</i> >)	The indicated multicast flow is not routing. The routing switch is low on memory resources as a result of too many flows for the number of configured VLANs. Remedies include one or more of the following: <ul style="list-style-type: none"> • Reduce the number of configured VLANs by moving some VLANs to another router. • Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters. • Move some hosts that create multicast demand to another router.
MCAST MAC add for < <i>mac-address</i> > failed (< <i>counter</i> >)	Indicates a hardware problem. Check the cabling and router ports.
Multicast Hardware Failed to Initialize (< <i>counter</i> >)	Indicates a hardware failure that halts hardware processing of PIM traffic. The software will continue to process PIM traffic at a slower rate. Contact your ProCurve customer care center.
No IP address configured on VID < <i>vlan-id</i> > (< <i>dup-msg-cnt</i> >)	PIM has detected a VLAN without an IP address. Configure an IP address on the indicated VLAN.
Pkt dropped from < <i>ip-address</i> >, (< <i>cause</i> >) vid < <i>vlan-id</i> > (< <i>counter</i> >)	A PIM packet from < <i>ip-address</i> > was dropped due to one of the following causes: <ul style="list-style-type: none"> • No PIM interface on the VLAN • Bad packet length • Bad IP header length • Bad IP total length
Pkt rcvd with a cksum error from < <i>ip-addr</i> > (< <i>counter</i> >)	A packet having a checksum error was received from < <i>ip-address</i> >. Check the cabling and ports on the local and the remote routers.
Rcvd incorrect hello from < <i>ip-addr</i> > (< <i>counter</i> >)	Indicates receipt of a malformed hello packet. (That is, the packet does not match the current specification.) Ensure that compatible versions of PIM-DM are being used.
Rcvd < <i>text-str</i> > pkt with bad len from < <i>ip-addr</i> > (< <i>counter</i> >)	A peer router may be sending incorrectly formatted PIM packets.
Rcvd hello from < <i>ip-address</i> > on vid < <i>vlan-id</i> > (< <i>counter</i> >)	Indicates a misconfiguration where two routers are directly connected with different subnets on the same connected interface.

PIM-DM (Dense Mode)

Messages Related to PIM Operation

Message	Meaning
Rcvd pkt from rtr <ip-address>, unknwn pkt type <value> (<counter>)	A packet received from the router at <ip-address> is an unknown PIM packet type. (The <value> variable is the numeric value received in the packet.)
Rcvd pkt ver# <ver-num>, from <ip-address>, expected <ver-num> (<counter>)	The versions of PIM-DM on the sending and receiving routers do not match. Differing versions will typically be compatible, but features not supported in both versions will not be available.
Rcvd unknwn addr fmly <addr-type> in <text-str> pkt from <ip-addr> (<counter>)	The router received a PIM packet with an unrecognized encoding. As of February, 2004, the router recognizes IPv4 encoding.
Rcvd unknwn opt <opt-nbr> in <text-string> pkt from <ip-addr> (<counter>)	The router received a PIM packet carrying an unknown PIM option. The packet may have been generated by a newer version of PIM-DM, or is corrupt. In most cases, normal PIM-DM operation will continue.
Send error(<failure-type>) on <packet-type> pkt on VID <vid> (<counter>)	Indicates a send error on a packet. This can occur if a VLAN went down right after the packet was sent. The message indicates the failure type, the packet type, and the VLAN ID on which the packet was sent.
Unable to alloc<text-str> table (<counter>)	The router was not able to create some tables PIM-DM uses. Indicates that the router is low on memory resources. Remedies include one or more of the following: <ul style="list-style-type: none">• Reduce the number of configured VLANs by moving some VLANs to another router.• Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters.• Move some hosts that create multicast demand to another router.
Unable to alloc a buf of size <bytes> for <data-flow> (<counter>)	Multicast routing is unable to acquire memory for a flow. Router memory is oversubscribed. Reduce the number of VLANs or the number of features in use. Remedies include one or more of the following: <ul style="list-style-type: none">• Reduce the number of configured VLANs by moving some VLANs to another router.• Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters.• Move some hosts that create multicast demand to another router.

Message	Meaning
Unable to alloc a msg buffer for < <i>text-message</i> > (< <i>counter</i> >)	Multicast routing is unable to acquire memory for a flow. Router memory is oversubscribed. Reduce the number of VLANs or the number of features in use. Remedies include one or more of the following: <ul style="list-style-type: none">• Reduce the number of configured VLANs by moving some VLANs to another router.• Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters.• Move some hosts that create multicast demand to another router.

Applicable RFCs

PIM on the switches covered in this guide is compatible with these RFCs:

- RFC 3376 - Internet Group Management Protocol, Version 3
- RFC 2365 - Administratively Scoped IP Multicast
- RFC 2932 - Multicast Routing MIB, *with exceptions (Refer to "Exceptions to Support for RFC 2932 - Multicast Routing MIB".)*
- RFC 2933 - IGMP MIB
- RFC 2934 - Protocol Independent Multicast MIB for IPv4
- draft-ietf-ssm-arch-01.txt - Source-Specific Multicast for IP (draft specification, expires May 2003)

Exceptions to Support for RFC 2932 - Multicast Routing MIB

These MIB objects are not supported in the switches covered in this guide.

ipMRouteInterfaceRateLimit

ipMRouteInterfaceInMcastOctets

ipMRouteInterfaceOutMcastOctets

ipMRouteInterfaceHCInMcastOctets

ipMRouteInterfaceHCOutMcastOctets

ipMRouteBoundaryTable

ipMRouteBoundaryEntry

ipMRouteBoundaryIfIndex

ipMRouteBoundaryAddress

ipMRouteBoundaryAddressMask

ipMRouteBoundaryStatus OBJECT-TYPE

ipMRouteScopeNameTable

ipMRouteScopeNameEntry

ipMRouteScopeNameAddress

ipMRouteScopeNameAddressMask

ipMRouteScopeNameLanguage

ipMRouteScopeNameString

ipMRouteScopeNameDefault

ipMRouteScopeNameStatus

PIM-SM (Sparse Mode)

Contents

Introduction	4-4
Feature Overview	4-5
Terminology	4-6
PIM-SM Operation and Router Types	4-9
PIM-SM Operation	4-9
Rendezvous-Point Tree (RPT)	4-9
Shortest-Path Tree (SPT)	4-10
Restricting Multicast Traffic to Rendezvous-Point Trees (RPTs)	4-11
Maintaining an Active Route for Multicast Group Members ...	4-11
Border Routers and Multiple PIM-SM Domains	4-12
PIM-SM Router Types	4-12
Designated Router (DR)	4-12
Bootstrap Router (BSR)	4-13
Rendezvous Point (RP)	4-14
Static Rendezvous Point (Static-RP)	4-17
Operating Rules and Recommendations	4-19
Configuration Steps for PIM-SM	4-20
Planning Considerations	4-20
Per-Router Global Configuration Context	4-20
Per-VLAN PIM-SM Configuration	4-21
Router PIM Configuration	4-23
Configuring PIM-SM on the Router	4-25
Global Configuration Context for Supporting PIM-SM	4-26
Global Configuration Context Commands	4-26
Example of Configuring for PIM Support at the Global Level ..	4-27

VLAN Context Commands for Configuring PIM-SM	4-28
Enabling or Disabling IGMP in a VLAN	4-28
Enabling or Disabling PIM-SM Per-VLAN	4-29
Changing the Interval for PIM-SM Neighbor Notification	4-30
Changing the Randomized Delay Setting for PIM-SM Neighbor Notification	4-31
Changing the PIM-SM Neighbor Timeout Interval	4-31
Enabling or Disabling LAN Prune Delay	4-32
Changing the LAN-Prune-Delay Interval	4-33
Changing the DR (Designated Router) Priority	4-33
Example of Configuring PIM-SM Support in a VLAN Context . .	4-34
Router PIM Context Commands for Configuring PIM-SM Operation	4-35
Configuring a BSR Candidate	4-35
Configuring Candidate-RPs on PIM-SM Routers	4-37
Enabling, Disabling, or Changing Router PIM Notification Traps	4-41
Changing the Global Join-Prune Interval on the Router	4-42
Changing the Shortest-Path Tree (SPT) Operation	4-42
Statically Configuring an RP To Accept Multicast Traffic	4-42
Example of Configuring PIM-SM Support in the Router PIM Context	4-43
Displaying PIM-SM Data and Configuration Settings	4-46
Displaying Multicast Route Data	4-47
Listing Basic Route Data for Active Multicast Groups	4-47
Listing Data for an Active Multicast Group	4-48
Listing All VLANs Having Currently Active PIM Flows	4-50
Displaying PIM-Specific Data	4-51
Displaying the Current PIM status and Global Configuration . .	4-51
Displaying Current PIM Entries Existing In the Multicast Routing Table	4-52
Displaying a Specific PIM Entry Stored in the Multicast Routing Table	4-53
Listing Currently Configured PIM Interfaces	4-55
Displaying IP PIM VLAN Configurations	4-55
Displaying PIM Neighbor Data	4-57
Displaying BSR Data	4-59

Displaying BSR Status and Configuration	4-59
Listing Non-Default BSR Configuration Settings	4-60
Displaying the Current RP Set	4-61
Displaying Candidate-RP Data	4-63
Displaying the Router's Candidate-RP Status and Configuration	4-63
Listing Non-Default C-RP Configuration Settings	4-64
Operating Notes	4-65
Event Log Messages	4-66

Introduction

Feature	Default	CLI
Enable PIM-SM Support	Disabled	4-26
Configure PIM-SM on VLAN Interfaces	Disabled	4-28
Configure Router PIM Context	Disabled	
Bootstrap Router Candidate		4-35
Rendezvous-Point Candidate		4-37
Notification Traps		4-41
Shortest-Path Tree		4-42
Display Multicast Route Data	n/a	4-47
Display PIM-Specific Data	n/a	4-51
Display PIM Neighbor Data	n/a	4-57
Display BSR and C-RP Data	n/a	4-59
Display Current RP-Set	n/a	4-61
Display Candidate-RP Data	n/a	4-63

In a network where IP multicast traffic is transmitted for multimedia applications, such traffic is blocked at routed interface (VLAN) boundaries unless a multicast routing protocol is running. Protocol Independent Multicast (PIM) is a family of routing protocols that form multicast trees to forward traffic from multicast sources to subnets that have used a protocol such as IGMP to request the traffic. PIM relies on the unicast routing tables created by any of several unicast routing protocols to identify the path back to a multicast source (*reverse path forwarding*, or RPF). With this information, PIM sets up the distribution tree for the multicast traffic. The PIM-DM and PIM-SM protocols on the switches covered by this manual enable and control multicast traffic routing.

IGMP provides the multicast traffic link between a host and a multicast router running PIM-SM. Both PIM-SM and IGMP must be enabled on VLANs whose member ports have directly connected hosts with a valid need to join multicast groups.

PIM-DM (described in chapter 3) is used in networks where, at any given time, multicast group members exist in relatively large numbers and are present in most subnets. However, using PIM-DM in networks where multicast sources

and group members are sparsely distributed over a wide area can result in unnecessary multicast traffic on routers outside the distribution paths needed for traffic between a given multicast source and the hosts belonging to the multicast group. In such networks, PIM-SM can be used to reduce the effect of multicast traffic flows in network areas where they are not needed. And because PIM-SM does not automatically flood traffic, it is a logical choice in lower bandwidth situations.

Feature Overview

PIM-SM on the routers covered by this manual includes:

- **Routing Protocol Support:** PIM uses whichever IP unicast routing protocol is running on the router. These can include:
 - RIP
 - OSPF
 - static routes
 - directly connected interfaces
- **VLAN Interface Support:** Up to 511 outbound VLANs (and 1 inbound VLAN) are supported in the multicast routing table (MRT) at any given time. This means the sum of all outbound VLANs across all current flows on a router may not exceed 511. (A single flow may span one inbound VLAN and up to 511 outbound VLANs, depending on the VLAN memberships of the hosts actively belonging to the flow.)
- **Flow Capacity:** Up to 2048 flows are supported in hardware across a maximum of 512 VLANs. (A flow is composed of an IP source address and an IP multicast group address, regardless of the number of active hosts belonging to the multicast group at any given time.)
- **Multicast Group to Rendezvous Point (RP) Mapping:** PIM-SM uses the Bootstrap Router (BSR) protocol to automatically resolve multicast group addresses to Candidate-RP routers. In the current software release, a router administers BSR operation on a PIM-SM domain basis. (BSR zones and PIM border router operation are not currently supported by the software covered in this guide.) Note that BSR operation does not extend to statically configured RPs. (For more on this topic, refer to “Static Rendezvous Point (Static-RP)” on page 4-17.)

- **IGMP Compatibility:** PIM-SM is compatible with IGMP version 2, and is fully interoperable with IGMP for determining multicast flows.
 - **VRRP:** PIM-SM is fully interoperable with VRRP to quickly transition multicast routes in the event of a failover.
 - **MIB Support on the Routers Covered by this Guide:**
 - PIM-SM supports the Protocol Independent Multicast MIB for IPv4 (RFC 2934).
 - With some exceptions, PIM-SM supports the parts of the Multicast Routing MIB (RFC 2932) applicable to PIM-SM operation. (Refer to “Exceptions to Support for RFC 2932 - Multicast Routing MIB” on page 3-42.)
 - **PIM Draft Specifications:** Compatible with PIM-SM draft specification (RFC 2362, version 10).
-

Terminology

Bootstrap Router (BSR). In a given PIM-SM domain, the BSR is the router elected to distribute the RP-set to the candidate rendezvous points (C-RPs) in a PIM-SM domain. The BSR does not interact with static rendezvous points (static-RPs) For more information on BSRs, refer to “Bootstrap Router (BSR)” on page 4-13. See also “RP-Set”, below.

Bootstrap Message (BSM): A message sent from the current BSR to the other PIM-SM routers in the domain to distribute the current RP-set and the status of the sending BSR as the current bootstrap router.

Candidate Rendezvous Point (C-RP): A PIM-SM router configured as the distribution point for all traffic from a multicast traffic source to a particular multicast group (destination). Multiple C-RPs can be configured to support the same multicast group, but only one C-RP will be elected to actually distribute the traffic for that group. (See also **Rendezvous Point**, page 4-7.)

Dynamic RP: A PIM-SM router configured as a Candidate Rendezvous Point (C-RP).

C-RP: See **Candidate Rendezvous Point**, above.

Designated Router (DR): Within a given VLAN or network, the router elected to forward a multicast flow from its IP source (in the VLAN or network) to the appropriate rendezvous point (either an RP or static-RP) in the PIM-SM domain.

Edge Router: Any router directly connected to a host or other endpoint in the network.

Flow: Multicast traffic having one source and one multicast group address (destination). This traffic may reach many hosts in different subnets, depending on which hosts have issued joins for the same multicast group.

Multicast Source: A single device originating multicast traffic for other devices (receivers).

Prune: To eliminate branches of a multicast tree that have no hosts sending joins to request or maintain membership in that particular multicast group.

Rendezvous Point (RP): A router that is either elected from a pool of eligible C-RPs (dynamic RPs) or statically configured (static RP) to support the distribution of traffic for one or more multicast groups and/or ranges of multicast groups. The RP for a given multicast group receives that group's traffic from a DR on the VLAN receiving the traffic from a multicast traffic source. The RP then forwards the traffic to downstream edge or intermediate PIM-SM routers in the path(s) to the requesting hosts (end points). (See also **Candidate Rendezvous Point**, page 4-6).

Rendezvous Point Tree (RPT): The path extending from the DR through any intermediate PIM-SM routers leading to the PIM-SM edge router(s) for the multicast receiver(s) requesting the traffic for a particular multicast group. (Refer to "Rendezvous-Point Tree (RPT)" on page 4-9.)

Reverse Path Forwarding (RPF): This is a methodology that uses the unicast routing table created by IP protocols such as RIP and OSPF to determine the source address of a packet. PIM uses RPF to set up distribution trees for multicast traffic.

Router: In the context of this chapter, a router is any ProCurve switch model covered by this guide and configured with IP routing enabled.

Routing Switch: See **Router**, above.

RP: See **Rendezvous Point**, above.

RPT: See **Rendezvous Point Tree**.

RP-Set: A complete list of multicast-group-to-RP mappings the BSR has learned and distributed to the C-RPs in a given PIM-SM domain. The learned RP-set applies only to C-RPs, and not to static-RPs. (Note, however, that the **show ip pim rp-set** command lists both the learned RP-set from the BSR and any static-RPs configured on the router.)

Shortest Path Tree (SPT): The shortest path from the DR through any intermediate PIM-SM routers leading to the PIM-SM edge router(s) for the multicast receiver(s) requesting the traffic for a particular multicast group. Unless the RPT is in this path, it is excluded from the SPT. (Refer to “Shortest-Path Tree (SPT)” on page 4-10.)

SPT: See **Shortest Path Tree**.

Static Rendezvous Point (Static-RP). A PIM-SM router manually configured as the distribution point for a multicast group or range of contiguous groups. (Refer to “Static Rendezvous Point (Static-RP)” on page 4-17.)

PIM-SM Operation and Router Types

Unlike PIM-DM, PIM-SM assumes that most hosts do not want to receive multicast traffic, and uses a non-flooding multicast model to direct traffic for a particular multicast group from the source to the VLAN(s) where there are multicast receivers that have joined the group. As a result, this model sends traffic only to the routers that specifically request it.

PIM-SM Operation

In a given PIM-SM domain, routers identified as Designated Routers (DRs), Rendezvous Points (RPs), and a Bootstrap Router (BSR) participate in delivering multicast traffic to the IP multicast receivers that request it. This approach avoids the flooding method of distributing multicast traffic (employed by PIM-DM) and is best suited for lower bandwidth situations.

The software supports the following operation to enable multicast traffic delivery within a PIM-SM domain:

- From a pool of eligible DR candidates in each VLAN, one Designated Router (DR) is elected for each VLAN interface having at least one PIM-SM router. In a multinetted domain, this DR supports multicast traffic from a source on any subnet in the VLAN.
- From a pool of eligible Bootstrap Router (BSR) candidates in the domain, one BSR is elected for the entire domain.
- From a pool of eligible Candidate Rendezvous Points (C-RPs), one is elected to support each multicast group or range of groups allowed in the domain, excluding any group supported only by static-RPs. The multicast groups allowed in the domain are determined by the aggregation of the groups allowed by the individually configured RPs and any static-RPs. (Note that RP-Cs and static RP's can be configured with overlapping support for a given set of multicast groups.)

Rendezvous-Point Tree (RPT)

When a DR in a VLAN receives traffic for a particular multicast group from a source on that VLAN, the DR encapsulates the traffic and forwards it to the RP elected to support that multicast group. The RP decapsulates the traffic and forwards it on toward the multicast receiver(s) requesting that group. This forms a *Rendezvous Point Tree* (RPT) extending from the DR through any

intermediate PIM-SM routers leading to the PIM-SM edge router(s) for the multicast receiver(s) requesting the traffic. (If the RP has no current join requests for the group, then the traffic is dropped at the RP.)

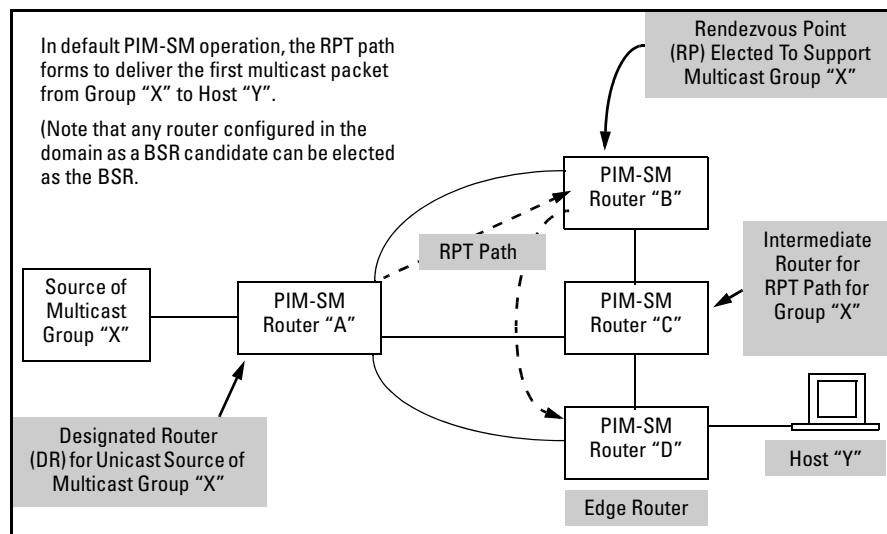


Figure 4-1. Example PIM-SM Domain with RPT Active To Support a Host Joining a Multicast Group

Shortest-Path Tree (SPT)

SPTs are especially useful in high data rate applications where reducing unnecessary traffic concentrations and throughput delays are significant. In the default PIM-SM configuration, SPT operation is automatically enabled. (The software includes an option to disable SPT operation. Refer to "Changing the Shortest-Path Tree (SPT) Operation" on page 4-42.)

Shortest-Path Tree Operation. In the default PIM-SM configuration, after an edge router receives the first packet of traffic for a multicast group requested by a multicast receiver on that router, it uses Reverse Path Forwarding (RPF) to learn the shortest path to the group source. The edge router then stops using the RPT and begins using the *shortest path tree* (SPT) connecting the multicast source and the multicast receiver. In this case, when the edge router begins receiving group traffic from the multicast source through the SPT, it sends a prune message to the RP tree to terminate sending the requested group traffic on that route. (This results in entries for both the RP path and the STP in the routing table. Refer to "Routing Table Entries" on

page 4-65.) When completed, the switchover from the RPT to a shorter SPT can reduce unnecessary traffic concentrations in the network and reduce multicast traffic throughput delays.

Note that the switchover from RPT to SPT is not instantaneous. For a short period, packets for a given multicast group may be received from both the RPT and the SPT. Also, in some topologies, the RPT and the SPT to the same edge router may be identical.

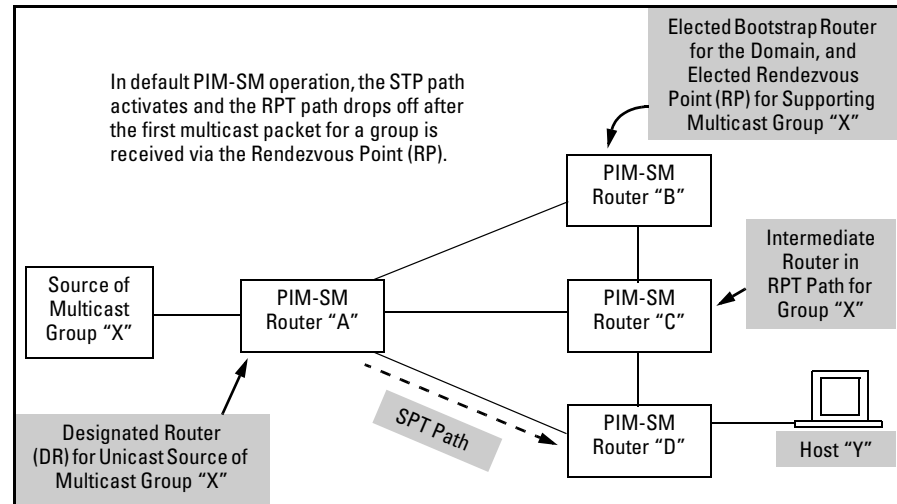


Figure 4-2. Example PIM-SM Domain with SPT Active To Support a Host that Has Joined a Multicast Group

Restricting Multicast Traffic to Rendezvous-Point Trees (RPTs)

An alternate method to allowing the domain to use SPTs is to configure all of the routers in the domain to use only RPTs. However, doing so can increase the traffic load in the network and cause delays in packet delivery.

Maintaining an Active Route for Multicast Group Members

The edge router itself and any intervening routers on the active tree between the members (receivers) of a multicast group and the DR for that group, send periodic joins. This keeps the active route available for as long as there is a multicast receiver requesting the group. When a route times out or is pruned, the DR ceases to send the requested group traffic on that route.

Border Routers and Multiple PIM-SM Domains

Creating multiple domains enables a balancing of PIM-SM traffic within a network. Defining PIM-SM domain boundaries requires the use of PIM border routers (PMBRs), and multiple PMBRs can be used between any two domains.

Note

As of March 2006, the software covered by this guide does not support PMBR operation for PIM-SM networks.

PIM-SM Router Types

Within a PIM-SM domain, PIM-SM routers can be configured to fill one or more of the roles described in this section.

- **Designated Router (DR):** A router performing this function forwards multicast traffic from a unicast source to the appropriate distribution (rendezvous) point. Refer to “Designated Router (DR)”, below.
- **Bootstrap Router (BSR):** A router elected to this function keeps all routers in a PIM-SM domain informed of the currently assigned RP for each multicast group currently known in the domain. Refer to “Bootstrap Router (BSR)” on page 4-13.
- **Rendezvous Point (RP):** A router elected as a rendezvous point for a multicast group receives requested multicast traffic from a DR and forwards it toward the multicast receiver(s) requesting the traffic. Refer to “Rendezvous Point (RP)” on page 4-14.
- **Static Rendezvous Point (Static-RP):** This option forwards traffic in the same way as an RP, but requires manual configuration on all routers in the domain to be effective.

All of the above functions can be enabled on each of several routers in a PIM-SM domain. For more information, refer to the following sections.

Designated Router (DR)

In a VLAN populated by one or more routers running PIM-SM, one such router is elected the *Designated Router* (DR) for that VLAN. When the DR receives a Join request from a multicast receiver on that VLAN, it forwards the Join towards the router operating as the RP for the requested multicast group.

Where multiple PIM-SM routers exist in a VLAN, the following criteria is used to elect a DR:

1. The router configured with the highest DR priority in the VLAN is elected.

2. If multiple routers in the VLAN are configured with the highest DR priority, then the router having the highest IP address is elected.

In a given domain, each VLAN capable of receiving multicast traffic from a unicast source should have at least one DR. (Enabling PIM-SM on a VLAN automatically enables the router as a DR for that VLAN.) Because there is an election process for DR on each VLAN, it is generally recommended that all routers on a VLAN be enabled for DR. Where it is important to ensure that a particular router is elected as the DR for a given VLAN, you can increase the DR priority on that VLAN configuration for that router.

If it is necessary to prevent a router from operating as a DR on a given VLAN, disable DR operation by configuring the DR priority as 0 (zero).

Bootstrap Router (BSR)

Before a DR can forward encapsulated packets for a specific multicast group to an RP, it must know which router in the domain is the elected RP for that multicast group. The bootstrap router (BSR) function enables this operation by doing the following:

1. Learns the group-to-RP mappings on the Candidate Rendezvous Points (C-RPs) in the domain by reading the periodic advertisements each one sends to the BSR.
2. Distributes the aggregate C-RP information as an *RP-set* to the PIM-SM routers in the domain. This is followed by an election to assign a specific multicast group or range of groups to the C-RPs in the domain. (The software supports assignment of up to four multicast addresses and/or ranges of multicast addresses to a C-RP.)

The BSR periodically sends bootstrap messages to the other PIM-SM routers in the domain to maintain and update the RP-set data throughout the domain, and to maintain its status as the elected BSR.

Note

Where static RPs are configured in the domain to support the same multicast group(s) as one or more (dynamic) C-RPs, then the RP-set data has the precedence for assigning RPs for these groups unless the static-RPs have been configured with the **override** option *and* if the multicast group mask for the static-RP equals or exceeds the same mask for the applicable C-RP(s). Refer to the **Note** on page 4-17.

BSR Configuration and Election. There should be multiple BSR candidates configured in a PIM-SM domain so that if the elected BSR becomes unavailable, another router will take its place. In the BSR election process,

the BSR candidate configured with the highest priority number is selected. Where the highest priority setting is shared by multiple candidates, the candidate having the highest IP address is selected. In the event that the selected BSR subsequently fails, another election takes place among the remaining BSR candidates. To facilitate a predictable BSR election, configure a higher priority on the router you want elected as the BSR for the domain. (Refer to “Changing the Priority Setting for a BSR-Candidate Router” on page 4-36.)

Note

A router serving as the BSR for a domain should be central to the network topology. This will help to ensure optimal performance and also reduce the possibility of a network problem isolating the BSR.

BSR Role in Fault Recovery. If the hold-time maintained in the BSR for a given C-RP’s latest advertisement expires before being refreshed by a new advertisement from the C-RP, then the non-reporting C-RP is removed from the domain. In this case, the removed C-RP’s multicast groups are re-assigned to other C-RPs. (If no other C-RPs or static-RPs in the domain are configured to support a multicast group from the non-reporting C-RP, then that group becomes unavailable in the domain.)

Rendezvous Point (RP)

Instead of flooding multicast traffic as is done with PIM-DM, PIM-SM uses a set of multiple routers to operate as *rendezvous points* (RPs). Each RP controls multicast traffic forwarding for one or more multicast groups as follows:

- receives traffic from multicast sources (S) via a DR
- receives multicast joins from routers requesting multicast traffic
- forwards the requested multicast traffic to the requesting routers

Note that the routers requesting multicast traffic are either edge routers directly connected to specific multicast receivers using IGMP to request the traffic, or are intermediate routers on the path between the edge routers and the RP. This operation forms an *RP Tree* (RPT) where only the destination multicast address appears in the RP routing table. This is represented as follows:

(*, G), where:

* = a variable (wildcard) representing the IP address of any multicast source

G = a particular multicast group address.

The software supports up to 100 RPs in a given PIM-SM domain.

Defining Supported Multicast Groups. An RP in the default candidate configuration supports the entire range of possible multicast groups. This range is expressed as a multicast address and mask, where the mask defines whether the address is for a single address or a range of contiguous addresses:

Multicast Address	Mask	Address Range
224.0.0.0	240.0.0.0	224.0.0.0 - 239.255.255.255

An alternate way to express the above (default) address and mask is:

224.0.0.0/4

In non-default candidate configurations, an RP allows up to four ranges of contiguous multicast groups, and/or individual multicast groups. For example:

RP Candidate Configuration	Supported Range of Multicast Groups
235.0.240.0/12	235.0.240.1 - 235.0.255.255
235.0.0.1/28	235.0.0.1 - 235.0.0.15
235.0.0.128/32	235.0.0.128 only
235.0.0.77/32	235.0.0.77 only

Note

If a given multicast group is excluded from all RPs in a given domain, then that group will not be available to the multicast receivers connected in the domain.

For more on this topic, refer to “Configuring Candidate-RPs on PIM-SM Routers” on page 4-37.

Candidate-RP Election. Within a PIM-SM domain, different RPs support different multicast addresses or ranges of multicast addresses. (That is, a given PIM-SM multicast group or range of groups is supported by only one active RP, although other candidate RPs can also be configured with overlapping or identical support.)

A candidate RP’s group-prefix configuration identifies the multicast groups the RP is enabled to support.

If multiple candidate RPs have group prefixes configured so that any of these RPs can support a given multicast group, then the following criteria are used to select the RP to support the group:

1. The C-RP configured with the longest group-prefix mask applicable to the multicast group is selected to support the group. If multiple RP candidates meet this criterion, then step 2 applies.
2. The C-RP configured with the highest priority is selected. If multiple RP candidates meet this criterion, then step 3 applies.
3. A hash function (using the configured **bsr-candidate hash-mask-length** value) generates a series of mask length values that are individually assigned to the set of eligible C-RPs. If the hash function matches a single RP candidate to a longer mask length than the other candidates, that candidate is selected to support the group. If the hash function matches the longest mask length to multiple RP candidates, then step 4 applies.
4. The C-RP having the highest IP address is selected to support the group.

Notes

In a PIM-SM domain where there are overlapping ranges of multicast groups configured on the C-RPs, discrete ranges of these groups are assigned to the domain's C-RPs in blocks of sequential group numbers. The number of multicast groups in the blocks assigned within a given domain is determined by the **bsr-candidate hash-mask-length** value (range = 1 - 32; page 4-36) configured on the elected BSR for the domain. A higher value means fewer sequential group numbers in each block of sequential group numbers, which results in a wider dispersal of multicast groups across the C-RPs in the domain.

As indicated above, multiple C-RPs can be configured to support the same multicast group(s). This is the generally recommended practice, and results in redundancy that helps to prevent loss of support for desired multicast groups in the event that a router in the domain becomes unavailable.

Configuring a C-RP to support a given multicast group does not ensure election of the C-RP to support that group unless the group is excluded from all other RPs in the domain. Refer to “” on page 4-16.

Also, within a PIM-SM domain, a router can be configured as a C-RP available for a given multicast group or range of groups and as the static RP for a given multicast group or range of groups. The recommended practice is to use C-RPs for all multicast groups unless there is a need to ensure that a specific group or range of groups is always supported by the same routing switch. For more on this topic, refer to “Static Rendezvous Point (Static-RP)” on page 4-17.

Redundant Group Coverage Provides Fault-Tolerance. If a C-RP elected to support a particular multicast group or range of groups becomes unavailable, the router will be excluded from the RP-set. If the multicast group configuration of one or more other C-RPs overlaps the configuration in the failed RP, then another C-RP will be elected to support the multicast group(s) formerly relying on the failed RP.

Static Rendezvous Point (Static-RP)

General Application. Like C-RPs, static-RPs control multicast forwarding of specific multicast groups or ranges of contiguous groups. However, static-RPs are not dynamically learned, and increase the configuration and monitoring effort needed to maintain them. As a result static-RPs are not generally recommended for use except where one of the following conditions applies:

- It is desirable to designate a specific router interface as a backup RP for specific group(s).
- Specific multicast groups are expected, and a static-RP would help to avoid overloading a given RP with a high volume of multicast traffic.
- A C-RP for the same group(s) is less reliable than another RP that would not normally be elected to support the group(s).
- tighter traffic control or a higher priority is desired for specific multicast groups

Notes

While use of C-RPs and a BSR enable a dynamic selection of RPs for the multicast group traffic in a network, using static-RPs involves manually configuring all routers in the domain to be aware of each static RP. This can increase the possibility of multicast traffic failure due to misconfigurations within the PIM-SM domain. Also, because a BSR does not administer static-RPs, troubleshooting PIM-SM traffic problems can become more complex. For these reasons, use of static-RPs should be limited to applications where no viable alternatives exist, or where the network is stable and requires configuring and maintaining only a few routers.

If a static-RP operating as the primary RP for a multicast group fails, and the PIM-SM configuration in the domain does not include a (secondary) dynamic RP (C-RP) backup to the static-RP, then new multicast groups assigned to the static-RP will not be available to multicast receivers in the domain. Also, if a static-RP fails, support for existing groups routed through SPTs that *exclude* the failed router will continue, but any existing flows routed through the RPT will fail.

Supporting a Static-RP as Primary . A static-RP can be configured to operate as either a secondary or primary RP. With the primary option, a dynamic (C-RP) backup is recommended. The precedence of a static-RP over a dynamic RP is determined by the following static-RP configuration options:

- **override** enabled on the static-RP
- a group mask on the static-RP that equals or exceeds the group mask on the C-RP for the same multicast group(s)

For **override** configuration information, refer to “Statically Configuring an RP To Accept Multicast Traffic” on page 4-42.

Operating Rules for Static RPs.

- Static-RPs can be configured on the same routers as C-RPs.
- Where a C-RP and a static-RP are configured to support the same multicast group(s), the C-RP takes precedence over the static-RP unless the static-RP is configured to override the C-RP. (Refer to “Supporting a Static-RP as Primary”, above.)
- Any static-RP in a domain must be configured identically on all routers in the domain. Otherwise, some DRs will not know of the static-RP and will not forward the appropriate multicast traffic, and some routers will not know where to send Joins for the groups supported by static-RP.
- Up to four static-RP entries can be configured on a router. Each entry can be for either a single multicast group or a range of contiguous groups.
- Only one interface can be configured as the static RP for a given multicast group or range of groups. For example, a properly configured PIM-SM domain does not support configuring 10.10.10.1 and 10.20.10.1 to both support a multicast group identified as 239.255.255.10.
- Static-RPs are not included in the RP-set messages generated by the BSR, and do not generate advertisements.
- If a static-RP becomes unavailable, it is necessary to remove and/or replace the configuration for this RP in all routers in the domain.

Configuration. Refer to “Statically Configuring an RP To Accept Multicast Traffic” on page 4-42.

Operating Rules and Recommendations

Guideline for Configuring Candidate RPs and BSRs. Routers in a PIM-SM domain should usually be configured as both candidate RPs and candidate BSRs. Doing so can reduce some overhead traffic.

The Shortest-Path-Tree (SPT) Policy Should Be the Same for All RPs in a Domain. Allowing some RPs to remain configured to implement SPTs while configuring other RPs in the same domain to force RPT use can result in unstable traffic flows. (Use the `[no] ip pim-sparse spt-threshold` command to change between SPT and RPT operation on each router.)

Application of RPs to Multicast Groups. In a PIM-SM domain, a given multicast group or range of groups can be supported by only one RP. (Typically, multiple candidate RPs in a domain are configured with overlapping coverage of multicast groups, but only one such candidate will be elected to support a given group.)

Ensuring that the Candidate RPs in a PIM-SM Domain Cover All Desired Multicast Groups. All of the multicast groups you want to allow in a given PIM-SM domain must be included in the aggregate of the multicast groups configured in the domain's candidate RPs. In most cases, all C-RPs in a domain should be configured to support all RP groups (the default configuration for a router enabled as a C-RP). This provides redundancy in case an RP becomes unavailable. (If the C-RP supporting a particular multicast group becomes unavailable, another C-RP is elected to support the group as long as there is redundancy in the C-RP configuration for multiple routers. Note that in cases where routers are statically configured to support a specific group or range of groups, the C-RP prioritization mechanism allows for redundant support.

PIM-SM and PIM-DM. These two features cannot both be enabled on the same router at the same time.

Supporting PIM-SM Across a PIM Domain. To properly move multicast traffic across a PIM-SM domain, all routers in the domain must be configured to support PIM-SM. That is, a router without PIM-SM capability blocks routed multicast traffic in a PIM-SM domain.

Configuration Steps for PIM-SM

This process assumes that the necessary VLANs and IP addressing have already been configured on the routing switch.

Note

The switches covered by this guide do not support PMBR operation in the current software release.

Planning Considerations

- Where multiple routers are available to operate as the DR for a given source, set the DR priority on each router according to how you want the router used.
- Determine whether there are any bandwidth considerations that would call for disabling SPT operation. (If any routers in the domain have SPT operation disabled, then it should be disabled on all RPs in the domain. Refer to “Operating Rules and Recommendations” on page 4-19.)
- Determine the routers to configure as C-BSRs. In many applications, the best choice may be to configure all routers in the domain as candidates for this function.
- Determine the multicast group support you want on each C-RP and any static-RPs in the domain. The easiest option is to enable C-RP to support all possible multicast groups on all routers in the domain. However, if there are traffic control considerations you want to apply, you can limit specific multicast groups to specific routers and/or set priorities so that default traffic routes support optimum bandwidth usage.

Per-Router Global Configuration Context

Use these steps to enable routing and PIM operation in the global configuration context of each PIM-SM router (**ProCurve(config)#**_)

1. Enable routing. (Use **ip routing**.)
2. Enable multicast routing. (Use **ip multicast-routing**.)
3. Enable PIM. (Use **router pim**.)
4. Configure the routing method(s) needed to reach the interfaces (VLANs) on which you want multicast traffic available for multicast receivers in your network:

- Enable RIP or OSPF (Use **router < rip | ospf >**.)
- If desired, configure static routes to the destination subnets.
(Use **ip route < dest-ip-address >/< mask-bits > < next-hop-ip-addr >**.)

Per-VLAN PIM-SM Configuration

These steps configure PIM-SM in the VLAN interface context for each VLAN configured on the router (**ProCurve(vlan-< vid >#_)**).

1. Enable IGMP. (Use **ip igmp**.) Repeat this action on every router (and switch) having membership in the VLAN.
2. Enable the same routing method you enabled in step 4 under “Per-Router Global Configuration Context” on page 4-20. at both the global and VLAN levels on the routers where there are connected multicast receivers that may issue joins or send multicast traffic.
3. Enable PIM-SM on the VLAN interfaces where you want to allow routed multicast traffic. (Default: disabled)
 - a. If these VLANs do not already have static IP addresses, then statically configure one or more IP addresses on each VLAN you want to support PIM-SM operation. (PIM-SM cannot be enabled on a VLAN that does not have a statically configured IP address. That is, PIM-SM cannot use an IP address acquired by DHCP/Bootp.)
 - b. Use **ip pim-sparse** to enter the VLAN’s **pim-sparse** context and do one of the following:
 - Enable PIM-SM on the VLAN and allow the default option (**any**) to dynamically determine the source IP address for the PIM-SM packets sent from this VLAN interface.
 - Enable PIM-SM on the VLAN and specify an IP address for the PIM-SM packets sent from this VLAN interface. (The specified IP address must already be statically configured on the VLAN.)

(This step requires enabling **router pim** on the global configuration context. Refer to step 3 on page 4-20.)
 - c. In the VLAN’s **pim-sparse** context, you also have the option to change the current DR priority (default = 1) to the value wanted for the current router in the current VLAN. (Use **dr-priority < 0 - 4294967295 >**.)

Note

When you initially enable PIM-SM, ProCurve recommends that you leave the PIM-SM traffic control settings (listed in the next step) at their default settings. You can then assess performance and make configuration changes where a need appears.

4. This is an optional step in the initial PIM-SM configuration. (Refer to the preceding Note.) In the **pim-sparse** context of a given VLAN on which PIM-SM is enabled, change one or more of the traffic control settings listed in the following table. (Note that some VLAN context control settings apply to both PIM-SM and PIM-DM.)

Features Accessed in VLAN-< vid>-pim-sparse Context	Operation
ip-addr (page 4-29)	Sets or resets the source IP address for PIM-SM packets sent out on the interface. Also enables PIM-SM on the interface. (Default: any)
hello-interval* (page 4-30)	Resets the interval between transmitted PIM Hello packets on the interface. (Default: 30 seconds)
hello-delay* (page 4-31)	Resets the maximum delay for transmitting a triggered PIM Hello packet on the interface. (Default: 5 seconds)
nbr-timeout (page 4-31)	Resets the neighbor loss time interval for the interface. (Default: 180 seconds)
lan-prune-delay* (page 4-32)	Enables or disables the LAN prune delay feature on the interface. (Default: on)
override-interval* (page 4-33)	Resets the override interval of the LAN Prune Delay configured on the interface. (Default: 2500 milliseconds)
propagation-delay* (page 4-33)	Resets the delay interval for triggering LAN Prune Delay packets on the interface. (Default: 500 milliseconds)
dr-priority (page 4-33)	Resets the priority of the interface in the Designated Router election process. (Default: 1) If you want one router on a given VLAN to have a higher priority for DR than other routers on the same VLAN, use the dr-priority command to reconfigure the DR priority setting as needed. Otherwise, the highest DR priority among multiple routers on the same VLAN interface is assigned to the router having the highest source IP address for PIM-SM packets on that interface.

*Applies to both PIM-SM and PIM-DM operation.

Router PIM Configuration

These steps configure PIM-SM in the Router PIM context (**ProCurve(pim)#_**).

1. Specify the VLAN interface to advertise as the BSR candidate and enable the router to advertise itself as a candidate BSR in a PIM-SM domain. (Use **bsr-candidate source-ip-vlan < vid >**.)
2. Optional: To make BSR candidate selection occur quickly and predictably, set a different priority on each BSR candidate in the domain. (Use **bsr-candidate priority** — page 4-36.)
3. Do one of the following to configure RP operation:
 - Recommended: Enable C-RP operation and configure the router to advertise itself as a candidate RP to the BSR for the current domain. This step includes the option to allow the C-RP to be a candidate for either all possible multicast groups or for up to four multicast groups and/or ranges of groups. (Use **rp-candidate source-ip-vlan < vid > [group-addr/group-mask]**.)
 - Alternative or Additional Option: Use **rp-address < ip-addr > [group-addr/group-mask]** to statically configure the router as the RP for a specified multicast group or range of multicast groups. (This must be configured on all PIM-SM routers in the domain.)
4. Optional: In the PIM router context, change one or more of the traffic control settings in the following table.

Options Accessed in Router PIM Context	Operation
rp-candidate group-prefix < group-addr/group-mask >	Enter an address and mask to define an additional multicast group or a range of groups.
rp-candidate hold-time < 3 - 255 >	Tells the BSR how long it should expect the sending Candidate-RP router to be operative. (Default: 150; 0 if router is not a candidate)
rp-candidate priority < 0 - 255 >	Changes the priority for the Candidate-RP router. When multiple C-RPs are configured for the same multicast group(s), the priority determines which router becomes the RP for such groups. A smaller value means a higher priority. (Default: 192)
[no] spt-threshold (page 4-42)	Disable or enable the router's ability to switch multicast traffic flows to the shortest path tree. (Default: enabled)

PIM-SM (Sparse Mode)
Configuration Steps for PIM-SM

Options Accessed in Router PIM Context	Operation
join-prune-interval < 5 - 65535 > (page 4-30)	Optional: Globally change the interval for the frequency at which join and prune messages are forwarded on the router's VLAN interfaces. (Default: 60 seconds)
trap < neighbor-loss hardware-mrt-full software-mrt-full all > (page 4-41)	Optional: Enable or disable PIM traps. (Default: disabled.)

Configuring PIM-SM on the Router

Command	Page
Global Context Commands	
[no] ip routing	4-26
[no] ip multicast-routing	4-26
[no] router < rip ospf >	4-26
[no] ip route < <i>src-ip-addr/mask</i> >< <i>dest</i> >	4-26
[no] router pim	4-26
VLAN context	4-28
[no] ip igmp	4-28
ip pim-sparse	4-29
[<i>ip-address</i>]	4-33
hello-interval	4-30
hello-delay	4-31
nbr-timeout	4-32
lan-prune-delay	4-32
override-interval	4-33
propagation-delay	4-33
dr-priority	4-33
router pim Context	4-35
bsr-candidate source-ip-vlan	4-35
bsr-candidate	4-35
priority	4-36
hash-mask	4-36
bsm-interval	4-37
rp-candidate source-ip-vlan	4-38
rp-candidate	4-40
group-prefix	4-40
hold-time	4-40
priority	4-41
trap	4-41
ip pim-sparse spt-threshold	4-42
rp-address	4-43

Global Configuration Context for Supporting PIM-SM

Before configuring specific PIM-SM settings, it is necessary to enable IP routing, IP multicast-routing, an IP routing protocol, and PIM in the global configuration context. Also, if the router operates as an edge router for any end points (receivers) expected to join multicast groups, then it is also necessary to enable IGMP on the VLANs supporting such receivers.

Global Configuration Context Commands

Note

PIM-SM operation requires an IP routing protocol enabled on the router. You can use RIP, OSPF, and/or static routing. The examples in this section use RIP. For more on these topics, refer to 5 in this guide.

Syntax: [no] ip routing

*Enables IP routing on the router. The **no** form of the command disables IP routing. Note that before disabling IP routing, it is necessary to disable all other IP routing protocols on the router. (Default: Disabled)*

Syntax: [no] ip multicast-routing

Enables or disables IP multicast routing on the router. IP routing must be enabled first. Note that router PIM must be disabled before disabling IP multicast routing. (Default: Disabled)

Syntax: [no] router < ospf | rip >
[no] ip route < ip-addr/mask-len > [< ip-addr | vlan | reject | blackhole >]

These commands are the options for the IP routing protocol required to support PIM operation. For more on these options, refer to the chapter titled “IP Routing Features” in this guide.

Syntax: [no] router pim

*Enables PIM at the global level and puts the CLI into the PIM context level. Executing the **no** form of the command at the global level disables PIM. IP routing must be enabled before enabling PIM. (Default: Disabled.)*

Example of Configuring for PIM Support at the Global Level

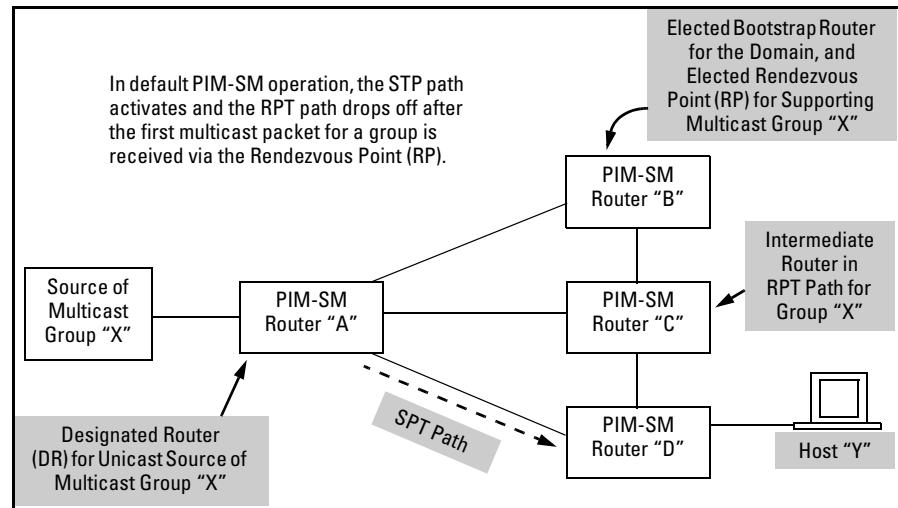


Figure 4-3. Example PIM-SM Domain with SPT Active To Support a Host that Has Joined a Multicast Group

Using the topology shown above, router "B" is directly connected to the DR for multicast group "X". In this case, suppose that you want to globally configure router "B" for PIM operation. On the global level, you would enable the following:

- IP routing
- IP multicast routing
- an IP routing protocol (RIP, OSPF, or static routing; use RIP for this example)

```
ProCurve(config)# ip routing
ProCurve(config)# ip multicast-routing
ProCurve(config)# router rip
ProCurve(rip)# exit
ProCurve(config)# router pim
ProCurve(pim)# exit
ProCurve(config)#
```

Figure 4-4. Global Configuration for Supporting PIM-SM Operation

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

Running configuration:

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

hostname "ProCurve"
module 2 type J8702A
module 1 type J8702A
ip routing
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24, B1-B24
  ip address 10.10.10.1 255.255.255.0
  exit
ip multicast-routing
router rip
  exit
router pim
  exit
-----
```

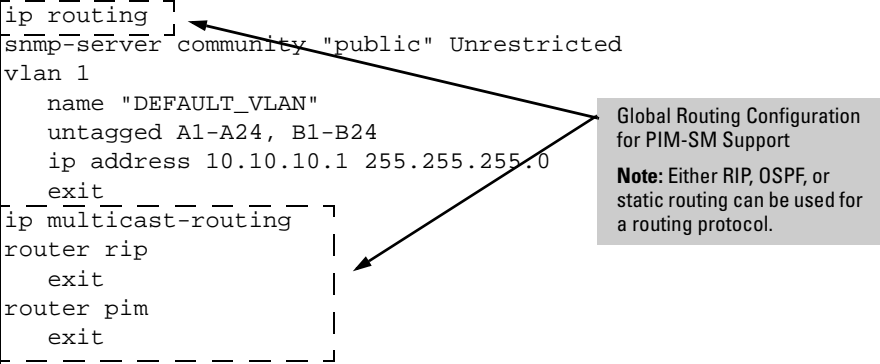


Figure 4-5. Displaying the Running Configuration

VLAN Context Commands for Configuring PIM-SM

PIM-SM must be configured on at least one VLAN in the router before it can be configured as a C-BSR or a C-RP.

Enabling or Disabling IGMP in a VLAN

IGMP must be enabled in VLANs on edge routers where multicast receivers (end points) are connected and will be requesting to join multicast groups.

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

Enables or disables IGMP operation in the current VLAN. Configuring IGMP on the router is required in VLANs supporting edge router operation. For more information, refer to the chapter titled “Multimedia Traffic Control with IP Multicast (IGMP)” in this guide.

Enabling or Disabling PIM-SM Per-VLAN

Syntax: ip pim-sparse [ip-addr < any | < ip-addr >>]
vlan < vid >] ip pim-sparse [ip-addr < any | < ip-addr >>]
no [vlan < vid >] ip pim-sparse

*This command enables or disables PIM-SM in the designated VLAN interface and sets the source (and designated router) IP address for PIM-SM packets sent from the interface. Executing the command without specifying an IP address option causes the router to default to the **any** option. The **no** form of the command disables PIM-SM on the specified VLAN. To change a currently configured (non-default) source IP address setting, use the **ip-addr < ip-addr >** option, below.
(Default: PIM-SM disabled)*

ip-addr any: *Enables the router to dynamically determine from the VLAN's current IP configuration the source IP address to use for PIM-SM packets sent from the VLAN interface.*

Note: *Using this command after a source IP address has already been set does not change that setting.*

— Continued on the next page.—

ip-addr < ip-addr >: *Specifies one of the VLAN's currently existing IP addresses for use as the source IP address for PIM-SM packets sent from the VLAN interface. Note that < ip-addr > must first be statically configured on the VLAN.*

Note: *To change an existing source IP address setting, you must use this command option.*

Changing the Interval for PIM-SM Neighbor Notification

Syntax: ip pim-sparse hello-interval < 5 - 300 >
vlan < vid > ip pim-sparse hello-interval < 5 - 300 >

*Changes the frequency at which the router transmits PIM “Hello” messages on the current VLAN. The router uses “Hello” packets to inform neighboring routers of its presence. The router also uses this setting to compute the **Hello Hold Time**, which is included in Hello packets sent to neighbor routers. **Hello Hold Time** tells neighbor routers how long to wait for the next Hello packet from the router. If another packet does not arrive within that time, the router removes the neighbor adjacency on that VLAN from the routing table, which removes any flows running on that interface. Shortening the Hello interval reduces the Hello Hold Time. This changes how quickly other routers will stop sending traffic to the router if they do not receive a new Hello packet when expected. For example, if multiple routers are connected to the same VLAN and the router requests multicast traffic, all routers on the VLAN receive that traffic. (Those which have pruned the traffic will drop it when they receive it.) If the upstream router loses contact with the router receiving the multicast traffic (that is, fails to receive a Hello packet when expected), then the shorter Hello Interval causes it to stop transmitting multicast traffic onto the VLAN sooner, resulting in less unnecessary bandwidth use. (Default: 30 seconds.)*

Changing the Randomized Delay Setting for PIM-SM Neighbor Notification

Syntax: ip pim-sparse hello-delay < 0 - 5 >
vlan < vid > ip pim-sparse hello-delay < 0 - 5 >

*Changes the maximum time in seconds before the router actually transmits the initial PIM Hello message on the current VLAN. In cases where a new VLAN activates with connections to multiple routers, if all of the connected routers sent Hello packets at the same time, then the receiving router could become momentarily overloaded. This value randomizes the transmission delay to a time between **0** and the **hello delay** setting. Using “0” means no delay. After the router sends the initial Hello Packet to a newly detected VLAN interface, it sends subsequent Hello packets according to the current **Hello Interval** setting. Not used with the **no** form of the **ip pim** command. (Default: 5 seconds.)*

Changing the PIM-SM Neighbor Timeout Interval

Syntax: ip pim-sparse nbr-timeout < 60 - 65535 >
vlan < vid > ip pim-sparse nbr-timeout < 60 - 65535 >

Changes the timeout interval allowed between successive Hello messages from a PIM-SM neighbor (in seconds) after which the neighbor will be considered unreachable. (Default: 180 seconds.)

Enabling or Disabling LAN Prune Delay

Syntax: [no] ip pim-sparse lan-prune-delay
[no] vlan < vid > ip pim-sparse lan-prune-delay

*Enables the LAN Prune Delay option on the current VLAN. With **lan-prune-delay** enabled, the router informs downstream neighbors how long it will wait before pruning a flow after receiving a prune request. Other, downstream routers on the same VLAN must send a Join to override the prune before the **lan-prune-delay** time if they want the flow to continue. This prompts any downstream neighbors with multicast receivers continuing to belong to the flow to reply with a Join. If no Joins are received after the **lan-prune-delay** period, the router prunes the flow. The **propagation-delay** and **override-interval** settings (below) determine the **lan-prune-delay** setting.*

*Uses the **no** form of the command to disable the LAN Prune Delay option.*

(Default: Enabled.)

Changing the LAN-Prune-Delay Interval

Syntax: ip pim-sparse propagation-delay < 250-2000 >
vlan < vid > ip pim-sparse propagation-delay < 250-2000 >

ip pim-sparse override-interval < 500 - 6000 >
vlan < vid > ip pim-sparse override-interval < 500 - 6000 >

*A router sharing a VLAN with other multicast routers uses these two values to compute the **lan-prune-delay** setting (above) for how long to wait for a PIM-SM join after receiving a prune packet from downstream for a particular multicast group. For example, a network may have multiple routers sharing VLAN “X”. When an upstream router is forwarding traffic from multicast group “X” to VLAN “Y”, if one of the routers on VLAN “Y” does not want this traffic it issues a prune response to the upstream neighbor. The upstream neighbor then goes into a “prune pending” state for group “X” on VLAN “Y”. (During this period, the upstream neighbor continues to forward the traffic.) During the “pending” period, another router on VLAN “Y” can send a group “X” Join to the upstream neighbor. If this happens, the upstream neighbor drops the “prune pending” state and continues forwarding the traffic. But if no routers on the VLAN send a Join, then the upstream router prunes group “X” from VLAN “Y” when the **lan-prune-delay** timer expires. (Defaults: **propagation-delay** = 500 milliseconds; **override-interval** = 2500 milliseconds.)*

Changing the DR (Designated Router) Priority

Syntax: ip pim-sparse dr-priority < 0 - 4294967295 >

This command changes the router priority for the DR (Designated Router) election process in the current VLAN. A numerically higher value means a higher priority. If the highest priority is shared by multiple routers in the same VLAN, then the router with the highest IP address is selected as the DR. A 0 (zero) value disables DR operation for the router on the current VLAN.

(Range: 0 - 2147483647; Default: 1)

Example of Configuring PIM-SM Support in a VLAN Context

PIM-SM support must be configured in each VLAN where you want PIM-SM forwarding of multicast traffic. This example illustrates the following per-VLAN configuration steps:

- Enabling PIM-SM on VLAN 120 and allow the default “any” option to select a source IP address for PIM-SM packets forwarded from this VLAN. (Because the VLAN in this example is configured with only one IP address—120.10.10.2—it is this address that will be used for the source.)
- Increasing the Designated Router (DR) priority on this VLAN from the default 1 to 100.
- Leaving the other per-VLAN PIM-SM fields in their default settings.

```
ProCurve(config)# vlan 120
ProCurve(vlan-120)# ip pim-sparse
ProCurve(vlan-120-pim-sparse)# dr-priority 100
ProCurve(vlan-120-pim-sparse)# show ip pim interface 120
```

PIM Interface

VLAN	: 120
IP Address	: 120.10.10.2
Mode	: sparse
Designated Router	: 120.10.10.2
Hello Interval (sec)	: 30
Hello Delay (sec)	: 5
Override Interval (msec)	: 2500
Propagation Delay (msec)	: 500
Neighbour Timeout	: 180
Lan Prune Delay	: Yes
Lan Delay Enabled	: Yes
DR Priority	: 100

The IP Address and Designated Router address correspond to the source IP address most recently configured on the VLAN. Indicates PIM-SM enabled.

Increased DR Priority
All other settings remain at their default-enabled values.

Figure 4-6. Example of Enabling PIM-SM in a VLAN

Router PIM Context Commands for Configuring PIM-SM Operation

This section describes the commands used in the Router PIM context to:

- enable or disable SNMP trap status for PIM events (default: disabled)
- configure candidate Bootstrap Router (BSR) operation
- configure candidate Rendezvous Point (RP) operation or the (optional) static Rendezvous Point (RP) operation

Note

Before configuring BSR, RP, and SNMP trap operation for PIM-SM, it is necessary to enable PIM-SM on at least one VLAN on the router.

Configuring a BSR Candidate

Select the VLAN Interface To Advertise as a BSR Candidate.

Syntax: [no] bsr-candidate source-ip-vlan < vid >
[no] router pim bsr-candidate source-ip-vlan < vid >

*Configures the router to advertise itself as a candidate PIM-SM Bootstrap Router (BSR) on the VLAN interface specified by **source-ip-vlan < vid >**, and enables BSR candidate operation. This makes the router eligible to be elected as the BSR for the PIM-SM domain in which it operates. Note that one BSR candidate VLAN interface is allowed per-router. The **no** form of the command deletes the BSR source IP VLAN configuration and also disables the router from being a BSR candidate if this option has been enabled. (See the **bsr-candidate** command, below.)*

Enable or Disable BSR Candidate Operation on a Router.

Syntax: [no] bsr-candidate
[no] router pim bsr-candidate

*Disables or re-enables the router for advertising itself as a Candidate-BSR on the VLAN interface specified by **source-ip-vlan < vid >**. This command is used to disable and re-enable BSR candidate operation after the **bsr-candidate source-ip-vlan < vid >** command has been used to enable C-BSR operation on the router. (That is, this command operates only after the BSR source-ip-VLAN ID has been configured.) (Default: Disabled.)*

Changing the Priority Setting for a BSR-Candidate Router.

Syntax: bsr-candidate priority < 0 - 255 >
[no] router pim bsr-candidate priority < 0 - 255 >

Specifies the priority to apply to the router when a BSR election process occurs in the PIM-SM domain. The candidate with the highest priority becomes the BSR for the domain. If the highest priority is shared by multiple routers, then the candidate having highest IP address becomes the domain's BSR. Zero (0) is the lowest priority. To make BSR selection easily predictable, use this command to assign a different priority to each candidate BSR in the PIM-SM domain. (Default: 0; Range 0 - 255.)

Note: *Disabling PIM-SM on the elected BSR or disabling the C-BSR functionality on the elected BSR causes the router to send a bootstrap message (BSM) with a priority setting of 0 (zero) to trigger a new BSR election. If all BSRs in the domain are set to the 0 (default) priority, then the election will fail because the result would be to re-elect the BSR that has become unavailable. For this reason, it is recommended that all C-BSRs in the domain be configured with a **bsr-candidate priority** greater than 0.*

Changing the Distribution of Multicast Groups Across a Domain.

Syntax: bsr-candidate hash-mask-length < 1 - 32 >
[no] router pim bsr-candidate hash-mask-length < 1 - 32 >

Controls distribution of multicast groups among the candidate RPs in a domain where there is overlapping coverage of the groups among the RPs. This value specifies the length (number of significant bits) taken into account when allocating this distribution. A longer hash-mask-length results in fewer multicast groups in each block of group addresses assigned to the various RPs. Because multiple blocks of addresses are typically assigned to each candidate RP, this results in a wider dispersal of addresses and enhances load-sharing of the multicast traffic of different groups being used in the domain at the same time. (Default: 30; Range: 1 - 32.)

Changing the Bootstrap Router Message Interval.

Syntax: bsr-candidate bsm-interval < 5 - 300 >
[no] router pim bsr-candidate bsm-interval < 5 - 300 >

Specifies the interval in seconds for sending periodic RP-Set messages on all PIM-SM interfaces on a router operating as the elected BSR in a domain.

*Note: This setting must be smaller than the **rp-candidate hold-time** settings (range of 30 - 255; default 150) configured in the RPs operating in the domain. (Default: 60; Range 5 - 300.)*

Configuring Candidate-RPs on PIM-SM Routers

Note

Before configuring BSR, RP, and SNMP trap operation for PIM-SM, it is necessary to enable PIM-SM on at least one VLAN on the router.

An RP candidate advertises its availability, IP address, and the multicast group or range of groups it supports. The commands in this section are used to configure RP candidate operation. The sequence of steps is as follows:

1. Specify the Source IP VLAN.
2. Enable Candidate-RP operation.
3. Optional: Enable or disable specific multicast address groups.

Specify the Source IP VLAN (and Optionally Configure one or more Multicast Groups or Range of Groups). Specifying the source IP VLAN ID automatically configures the RP candidate to support all multicast groups (unless you include an individual group or range of groups in the command). The recommended approach is to allow all multicast groups unless you have a reason to limit the permitted groups to a specific set.

Syntax: [no] rp-candidate source-ip-vlan < vid > [group-prefix < group-addr/mask >]
[no] router pim rp-candidate source-ip-vlan < vid > [group-prefix < group-addr/mask >]

This command configures C-RP operation as follows:

- *specifies the VLAN interface from which the RP IP address will be selected for advertising the router as an RP candidate. Note that only one VLAN on the router can be configured for this purpose at any time.*
- *enables the router as an RP candidate.*
- *specifies the multicast groups for which the router is a C-RP. (When executed without specifying a multicast group or range of groups, the resulting RP candidate defaults to allowing support for all multicast groups—224.0.0.0/240.0.0.0, or 224.0.0.0/4.*

(Default: Disabled.)

This command is required to initially configure the router as a Candidate-RP.

- *To later add to or change multicast groups, or to delete multicast groups, use **rp-candidate group-prefix < group-addr | group-mask >**, as described under “Adding or Deleting a Multicast Group Address” on page 4-40.*
- *To disable C-RP operation without removing the current C-RP configuration, use **no rp-candidate**. (Refer also to “Enabling or Disabling Candidate-RP Operation” on page 4-40.)*
- *The **no** form of the command:*
 - *deletes the RP source IP VLAN configuration*
 - *deletes the multicast group assignments configured on the router for this RP*
 - *disables the router from being an RP candidate.*

< vid >: *Identifies the VLAN source of the IP address to advertise as the RP candidate address for the router.*

group-prefix < group-addr/mask >: *Specifies the multicast group(s) to advertise as supported by the RP candidate. Use this option when you want to enable the Candidate-RP and simultaneously configure it to support a subset of multicast addresses or ranges of addresses instead of all possible multicast addresses.*

A group prefix can specify all multicast groups (224.0.0.0 - 239.255.255.255), a range (subset) of groups, or a single group. A given address is defined by its nonzero octets and mask. The mask is applied from the high end (leftmost) bits of the address and must extend to the last nonzero bit in the lowest-order, nonzero octet. Any intervening zero or nonzero octet requires eight mask bits. For example:

228.0.0.64/26: Defines a multicast address range of 228.0.0.64 through 228.0.0.127. (The last six bits of the rightmost octet are wildcards.)

228.0.0.64/30: Defines a multicast address range of 228.0.0.64 through 228.0.0.67. (The last two bits of the rightmost octet are wildcards.)

228.0.0.64/32: Defines a single multicast address of 228.0.0.64. (There are no wildcards in this group prefix.)

228.0.0.64/25: Creates an error condition due to the mask failing to include the last (rightmost) nonzero bit in the lowest-order, nonzero octet. (That is, this mask supports an address of 228.0.0.128, but not 228.0.0.64.)

Note that the larger the mask, the smaller the range of multicast addresses supported. A mask of 32 bits always specifies a single multicast address. For example:

230.0.15.240/32: Defines a single multicast address of 230.0.15.240.

Enabling or Disabling Candidate-RP Operation. Use this command when the router is already configured with a source IP VLAN ID and you want to enable or disable C-RP operation on the router.

Syntax: [no] rp-candidate

*Enables Candidate-RP operation on the router. Requires that the source IP VLAN is currently configured, but disabled (page 4-37). The **no** form of the command disables the currently configured Candidate-RP operation, but does not change the configured Candidate-RP settings.*

Adding or Deleting a Multicast Group Address. Use this command if you need to modify the multicast address group configuration for a candidate-RP on the router.

Syntax: [no] rp-candidate group-prefix < group-addr | group-mask >

*Adds a multicast group address to the current Candidate-RP configuration. Requires that the source IP VLAN (page 4-37) is already configured. The **no** form of the command removes a multicast group address from the current Candidate-RP configuration.*

This command does not enable or disable RP candidate operation.

Note: *An RP candidate supports up to four separate multicast address groups. Also, if only one group-prefix address exists in the Router PIM configuration, you cannot delete it unless you first add another group-prefix address.*

Changing the Candidate-RP Hold-Time. Hold-Time is included in the advertisements the Candidate-RP periodically sends to the domain's elected BSR, and updates the BSR on how long to wait after the last advertisement from the reporting RP before assuming that it has become unavailable. For more on this topic, refer to "BSR Role in Fault Recovery" on page 4-14.

Syntax: rp-candidate hold-time < 3 - 255 >

*Changes the hold time a C-RP includes in its advertisements to the BSR. Also, if C-RP is configured, but disabled, this command re-enables it.
(Default: 150 seconds; Range: 3 - 255 seconds.)*

Changing a Candidate-RP's Election Priority. This priority is significant when multiple Candidate-RPs in a given domain are configured to support one or more of the same multicast groups.

Syntax: rp-candidate priority < 0 - 255 >

Changes the current priority setting for a candidate-RP. Where multiple candidate-RPs are configured to support the same multicast group(s), the candidate having the highest priority is elected. Zero (0) is the highest priority; 255 is the lowest priority.

(Default: 192)

Enabling, Disabling, or Changing Router PIM Notification Traps

Syntax: [no] router pim trap < all | neighbor-loss | hardware-mrt-full | software-mrt-full >

Enables and disables these PIM SNMP traps:

all — Enable/Disable all PIM notification traps. (Default: Disabled)

neighbor-loss — Enable/Disable the notification trap sent when the timer for a multicast router neighbor expires and the switch has no other multicast router neighbors on the same VLAN with a lower IP address. (Default: Disabled.)

hardware-mrt-full — Enable/Disable notification trap sent when the hardware multicast routing table (MRT) is full (1023 active flows). In this state, any additional flows are handled by the software MRT, which increases processing time for the affected flows. (Default: Disabled.)

software-mrt-full — Enable/Disable notification trap sent when the router's software multicast routing table is full (that is, when routing resources for active flows are exhausted). Note that in this state, the router does not accept any additional flows. (Default: Disabled.)

Note: Trap operation requires configuring an SNMP trap receiver by using the **snmp-server host < ip-addr >** command at the global configuration level.

Changing the Global Join-Prune Interval on the Router

Syntax: router pim join-prune-interval <5 - 65535>

Sets the interval in seconds at which periodic PIM-SM join/prune messages are to be sent on the router's PIM-SM interfaces. This setting is applied to every PIM-SM interface on the router. (Default: 60 seconds)

Note: *All routers in a PIM-SM domain should have the same join-prune-interval setting.*

Changing the Shortest-Path Tree (SPT) Operation

Generally, using the SPT option eliminates unnecessary levels of PIM-SM traffic in a domain. However, in cases where it is necessary to tightly control the paths used by PIM-SM flows to edge switches, disabling SPT maintains the flows through their original C-RPs regardless of whether shorter paths exist.

Syntax: router pim spt-threshold
[no] router pim spt-threshold

When the router is the edge router for a receiver requesting to join a particular multicast group, this command enables or disables the capability of the router to convert the group's traffic from the RPT (Rendezvous Point Tree) to the SPT (shortest path tree). For more information, refer to "Restricting Multicast Traffic to Rendezvous-Point Trees (RPTs)" on page 4-11. (Default: Enabled.)

Statically Configuring an RP To Accept Multicast Traffic

A given static-RP entry should be manually configured on *all* routers in the PIM-SM domain. *For information on applying static-RPs in a domain, refer to "Static Rendezvous Point (Static-RP)" on page 4-17.*

Syntax: router pim rp-address <rp-ip-addr><group-addr/group-mask> [override]
[no] router pim rp-address <rp-ip-addr> <group-addr/group-mask>
[override]

<rp-ip-addr>: *Statically specifies the IP address of the interface to use as an RP. Up to four static-RP IP addresses can be configured. (Each address can be entered multiple times for different multicast groups or group ranges.)*

<group-addr/group-mask >: *Specifies the multicast group or range of contiguous groups supported by the statically configured RP.*

[override]: *Where a static-RP and a C-RP are configured to support the same multicast group(s) and the multicast group mask for the static RP is equal to or greater than the same mask for the applicable C-RPs, this command assigns the higher precedence to the static-RP, resulting in the C-RP operating only as a backup RP for the configured group. Without override, the C-RP has precedence over a static-RP configured for the same multicast group(s).*

Example of Configuring PIM-SM Support in the Router PIM Context

This example assumes the following:

- IP routing, IP multicast-routing, and at least one routing method (RIP, OSPF, and/or static IP routes) are already configured in the *global configuration context*.
- An IP routing method (RIP or OSPF) and PIM-sparse are already configured in the static VLAN context on which you want to support PIM-SM operation.

Note

Routers configured for C-RP operation can also be configured for C-BSR operation.

Use of static-RP operation must be identically configured on all PIM-SM routers in the domain.

PIM-SM (Sparse Mode)

Configuring PIM-SM on the Router

Figure 4-7 illustrates the following configuration steps for the Router PIM context:

- Enabling BSR operation on the router, including specifying a source IP address.
- Enabling C-RP operation on the router.
- Replacing the default multicast group range (all) with a smaller range (231.128.24.0/18) and a single group address (230.255.1.1/32).
- Enabling static-RP with an **override** on this router for a single group address (231.128.64.255/32) within the range of the C-RP support for the 231.128.24.0 group.
- Leaving the other Router PIM fields in their default settings.

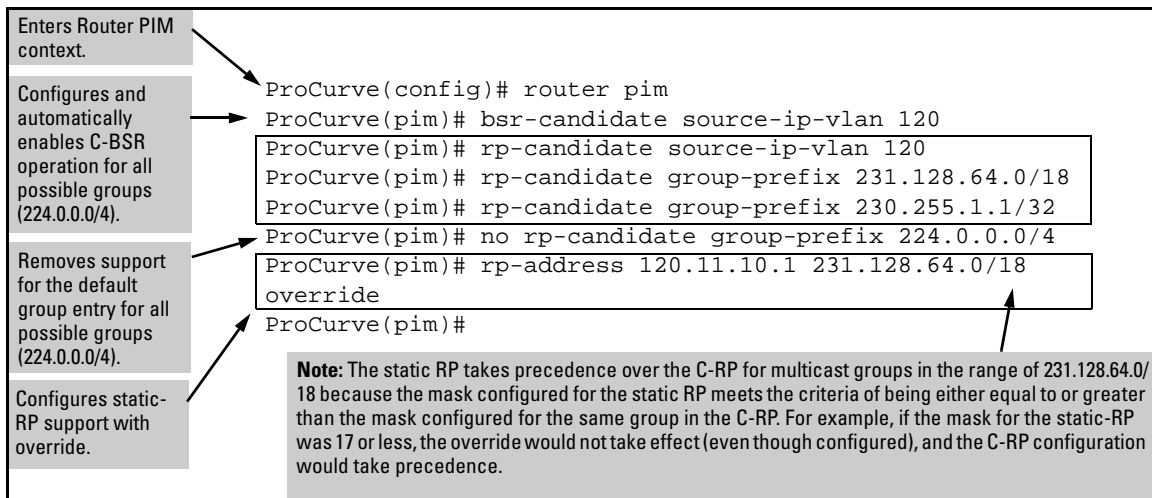


Figure 4-7. Example of Enabling PIM-SM in the Router PIM Context

The next figure illustrates the results of the above commands in the router's running configuration.


```
ProCurve(pim)# show running

Running configuration:
.
.
.
router pim
  bsr-candidate
  bsr-candidate source-ip-vlan 120
  bsr-candidate priority 1
  rp-address 120.10.10.2 231.128.64.255 255.255.255.255
  rp-candidate
  rp-candidate source-ip-vlan 120
  rp-candidate group-prefix 230.255.1.1 255.255.255.255
  rp-candidate group-prefix 231.128.64.0 255.255.192.0
  rp-candidate hold-time 150
exit
```

Figure 4-8. Configuration Results of the Commands in Figure 4-7

Displaying PIM-SM Data and Configuration Settings

Command	Page
show ip mroute	4-47
[< <i>group-addr</i> > < <i>source-ip-addr</i> >]	4-48
[interface [< <i>vid</i> >]]	4-50
show ip pim	4-51
[mroute]	4-52
[< <i>group-address</i> > < <i>source-address</i> >]	4-53
[interface]	4-55
[<i>vid</i>]	4-55
[neighbor]	4-57
[<i>ip-address</i>]	4-58
bsr	4-59
rp-set	4-61
[static learned]	4-61
rp-candidate	4-63
[config]	4-63

Displaying Multicast Route Data

The commands in this section display multicast routing information on packets sent from multicast sources to IP multicast groups detected by the routing switch.

Listing Basic Route Data for Active Multicast Groups

Syntax: show ip mroute

Lists the following data for all VLANs actively forwarding routed, multicast traffic.

Group Address: *The multicast address of the specific multicast group (flow).*

Source Address: *The IP address of the multicast group source.*

Neighbor: *The IP address of the upstream multicast router interface (VLAN) from which the multicast traffic is coming. A blank field for a given multicast group indicates that the multicast server is directly connected to the router.*

VLAN: *The interface on which the multicast traffic is moving.*

For example, the next figure displays the **show ip mroute** output illustrating a case where two multicast groups are from the same multicast server source.

```
ProCurve(config)# show ip mroute
```

Group Address	Source Address	Neighbor	VLAN
234.43.209.12	192.168.1.0	192.168.1.3	1
235.22.22.12	192.168.1.0	192.168.2.4	1

Figure 4-9. Example Showing Route Entry Data

Listing Data for an Active Multicast Group

Syntax: show ip mroute [< group-addr > < source-addr >]

Lists the following data for the specified flow (multicast group):

Group Address: *The multicast group IP address for the current group.*

Source Address: *The source IP address < source-ip-addr > for the current group.*

Source Mask: *The subnet mask applied to the multicast source address < source-ip-addr >.*

Neighbor: *Lists the IP address of the upstream next-hop router running PIM-SM; that is, the router from which the router is receiving datagrams for the current multicast group. This value is 0.0.0.0 if the router has not detected the upstream next-hop router's IP address. This field is empty if the multicast server is directly connected to the router.*

VLAN: *Lists the VLAN ID (VID) on which the router received the specified multicast flow.*

Up Time (Sec): *The elapsed time in seconds since the router learned the information for the current instance of the indicated multicast flow.*

Expire Time (Sec): *Indicates the remaining time in seconds before the router ages-out the current flow (group membership). This value decrements until:*

- *Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.)*
- *Reset by a new flow for the current multicast group on the VLAN.*
- *The timer expires (reaches 0). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.*

Multicast Routing Protocol: *Identifies the IP multicast routing protocol through which the current flow was learned.*

Unicast Routing Protocol: *Identifies the IP routing protocol through which the router learned the upstream interface for the current multicast flow. The listed protocol will be either RIP, OSPF, or Static Route.*

Downstream Interfaces:

VLAN: *Lists the VID of the VLAN the router is using to send the outbound packets of the current multicast flow to the next-hop router.*

State: *Indicates whether the outbound VLAN and next-hop router for the current multicast flow are receiving datagrams.*

- **Pruned:** *The router has not detected any joins from the current multicast flow and is not currently forwarding datagrams in the current VLAN.*
- **Forwarding:** *The router has received a join for the current multicast flow and is forwarding datagrams in the current VLAN.*

Up Time (Sec): *Indicates the elapsed time in seconds since the router learned the displayed information about the current multicast flow.*

```
ProCurve(config)# show ip mroute 234.43.209.12 192.168.1.0

IP Multicast Route Entry

Group Address   : 234.43.209.12
Source Address  : 192.168.1.0
Source Mask     : 255.255.255.0
Neighbor       : 192.168.1.3
VLAN           : 1
Up Time (sec)  : 757
Expire Time (sec) : 173
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol   : RIP

Downstream Interfaces
VLAN  State          Up Time (sec)      Expire Time (sec)
----  -
  2    forwarding    757                12
```

This Neighbor field indicates that the router is receiving multicast traffic from a neighboring PIM router. A blank Neighbor field indicates that the multicast server is directly connected to the router instead of another PIM router.

Figure 4-10. Example Showing Route Entry Data for a Specific Multicast Group

Listing All VLANs Having Currently Active PIM Flows

Syntax: show ip mroute interface [< vid >]

Lists these settings:

VLAN: *The VID specified in the command.*

Protocol: *PIM-SM or PIM-DM.*

TTL: *The time-to-live threshold for packets forwarded through this VLAN. When configured, the router drops multicast packets having a TTL lower than this value. (When a packet arrives, the router decrements its TTL by 1, then compares the decremented packet TTL to the value set by this command.) A **TTL Threshold** setting of **0** (the default) means all multicast packets are forwarded regardless of the TTL value they carry. A multicast packet must have a TTL greater than 1 when it arrives at the router. Otherwise the router drops the packet instead of forwarding it on the VLAN.*

```
ProCurve(config)# show ip mroute interface

IP Multicast Interfaces

VLAN Protocol TTL Threshold
-----
1     PIM-SM  0
80    PIM-SM  15
```

Figure 4-11. Example of Listing the Currently Active Mroute Interfaces

```
ProCurve(config)# show ip mroute interface 29

IP Multicast Interface

VLAN      : 29
Protocol  : PIM-SM
TTL Threshold : 0
```

Figure 4-12. Example of Listing the Mroute Data for a Specific Mroute Interface

Displaying PIM-Specific Data

The commands in this section display PIM-specific multicast routing information for IP multicast groups detected by the router.

Displaying the Current PIM status and Global Configuration

Syntax: show ip pim

Displays PIM status and global parameters.

PIM Status: Shows either **enabled** or **disabled**.

State Refresh Interval (sec): Applies only to PIM-DM operation. Refer to “*Displaying PIM Status*” on page 3-28.

Join/Prune Interval: Indicates the frequency with which the router transmits join and prune messages for the multicast groups the router is forwarding.

SPT Threshold: When **Enabled** indicates that, for a given receiver joining a multicast group, an edge router changes from the RPT to the SPT after receiving the first packet of a multicast flow intended for a receiver connected to the router. When **Disabled**, indicates that the no spt-threshold command has been used to disable SPT operation. (Refer to “*Changing the Shortest-Path Tree (SPT) Operation*” on page 4-42.)

Traps: Enables the following SNMP traps:

- **neighbor-loss:** Sends a trap if a neighbor router is lost.
- **hardware-mrt-full:** Sends a trap if the hardware multicast router (MRT) table is full (2048 active flows).
- **software-mrt-full:** Sends a trap if the software multicast router (MRT) table is full (2048 active flows). This can occur only if the hardware MRT is also full.
- **all:** Enables all of the above traps.

```
ProCurve(config)# show ip pim

PIM Global Parameters

PIM Status                : enabled
State Refresh Interval (sec) : 60
Join/Prune Interval (sec)  : 60
SPT Threshold              : Enabled
Traps                      : all
```

Figure 4-13. Example Output with PIM Enabled

Displaying Current PIM Entries Existing In the Multicast Routing Table

Syntax: show ip pim mroute

Shows PIM-specific information from the IP multicast routing table (IP MRT). When invoked without parameters, lists all PIM entries currently in the router's IP MRT.

Group Address: Lists the multicast group addresses currently active on the router.

Source Address: Lists the multicast source address for each Group Address.

Metric: Indicates the path cost upstream to the multicast source. Used when multiple multicast routers contend to determine the best path to the multicast source. The lower the value, the better the path. This value is set to **0** (zero) for directly connected routes.

Metric Pref: Used when multiple multicast routers contend to determine the path to the multicast source. When this value differs between routers, PIM selects the router with the lowest value. If **Metric Pref** is the same between contending multicast routers, then PIM selects the router with the lowest **Metric** value to provide the path for the specified multicast traffic. This value is set to **0** (zero) for directly connected routes. (**Metric Pref** is based on the IP routing protocol in use: RIP, OSPF, or static routing. Also, different vendors may assign different values for this setting.)

This output shows the routing switch is receiving two multicast groups from an upstream device at 27.27.30.2. The "0" metric shows that the routing switch is directly connected to the multicast source.

```
ProCurve# show ip pim mroute
PIM IP Multicast Route Entries
```

Group Address	Source Address	Metric	Metric Pref
234.43.209.12	100.150.1.0	2	1
235.22.22.12	100.100.25.0	0	1

Figure 4-14. Example Showing a Router Detecting two Multicast Groups from a Directly Connected Multicast Server

Displaying a Specific PIM Entry Stored in the Multicast Routing Table

Syntax: show ip pim mroute [< multicast-group-address >
< multicast-source-address >]

Displays the PIM route entry information for the specified multicast group (flow):

Group Address: Lists the specified multicast group address.

Source Address: Lists the specified multicast source address.

Source Mask: Lists the network mask for the multicast source address.

Metric: Indicates the path cost upstream to the multicast source. Used when multiple multicast routers contend to determine the best path to the multicast source. The lower the value, the better the path.

Metric Pref: Used when multiple multicast routers contend to determine the path to the multicast source. When this value differs between routers, PIM selects the router with the lowest value. If *Metric Pref* is the same between contending multicast routers, then PIM selects the router with the lowest **Metric** value to provide the path for the specified multicast traffic. (Different vendors assign differing values for this setting.)

Assert Timer: The time remaining until the router ceases to wait for a response from another multicast router to negotiate the best path back to the multicast source. If this timer expires without a response from any contending multicast routers, then the router assumes it is the best path, and the specified multicast group traffic will flow through the router.

RPT-bit: A **Yes** setting indicates the route is using the RPT. A **No** setting indicates the route is using the applicable SPT.

PIM-SM (Sparse Mode)

Displaying PIM-SM Data and Configuration Settings

DownStream Interfaces:

- **VLAN:** Lists the VID of the destination VLAN on the next-hop multicast router.
- **Prune Reason:** *Identifies the reason for pruning the flow to the indicated VLAN:*
 - **Prune:** *A neighbor multicast router has sent a prune request.*
 - **Assert:** *Another multicast router connected to the same VLAN has been elected to provide the path for the specified multicast group traffic.*
 - **Other:** *Used where the VLAN is in the pruned state for any reason other than the above two reasons (such as no neighbors exist and no directly connected multicast receivers have issued Joins).*

```
ProCurve# show ip pim mroute 234.43.209.12 192.168.1.0

PIM IP Multicast Route Entry

Group Address   : 234.43.209.12
Source Address  : 192.168.1.0
Source Mask     : 255.255.255.0

Metric          : 20
Metric Pref     : 1
Assert Timer    : 3 min 54 sec
RP-Tree        : Yes
Flags          : rpt, spt

DownStream Interfaces

VLAN Prune Reason
---- -
2    other
3    other
```

Figure 4-15. Example of PIM Mroute Listing for a Specific Multicast Flow

Listing Currently Configured PIM Interfaces

Syntax: show ip pim interface

Lists the PIM interfaces (VLANs) currently configured in the router.

VLAN: Lists the VID of each VLAN configured on the switch to support PIM-DM.

IP Address: Lists the IP addresses of the PIM interfaces (VLANs).

Mode: Shows **dense** or **sparse**, depending on which PIM protocol is configured on the router.

```
ProCurve(config)# show ip pim interface

PIM Interfaces

VLAN IP Address      Mode
----  -
1     10.1.10.1          sparse
2     10.2.10.1          sparse
```

Figure 4-16. Example Showing Two PIM Interfaces Configured

Displaying IP PIM VLAN Configurations

Syntax: show ip pim interface [< vid >]

Displays the current configuration for the specified VLAN (PIM interface). Refer to table 4-1 on page 4-56.

PIM-SM (Sparse Mode)

Displaying PIM-SM Data and Configuration Settings

```
ProCurve(config)# show ip pim interface 1

PIM Interface

VLAN      : 1
IP Address : 10.1.10.1
Mode      : sparse

Designated Router : 10.1.10.1

Hello Interval (sec) : 30
Hello Delay (sec)   : 5

Override Interval (msec) : 2500      Lan Prune Delay      : Yes
Propagation Delay (msec) : 500       Lan Delay Enabled    : No
Neighbour Timeout      : 180        DR Priority           : 1
```

Figure 4-17. Example Showing a PIM-SM Interface Configured on VLAN 1

Table 4-1. PIM Interface Configuration Settings

Field	Default	Control Command
VLAN	n/a	vlan < vid > ip pim
IP	n/a	vlan < vid > ip pim < all ip-addr >
Mode	dense	n/a; PIM Dense only
Hello Interval (sec)	300	ip pim hello interval < 5 - 30 >
Hello Delay	5	The router includes this value in the "Hello" packets the it sends to neighbor routers. Neighbor routers use this value to determine how long to wait for another Hello packet from the router. Refer to "Changing the Interval for PIM-SM Neighbor Notification" on page 4-30.
Override Interval (msec)	2500	vlan < vid > ip pim override-interval < 500 - 6000 >
Propagation Delay (msec)	500	vlan < vid > ip pim propagation-delay < 250-2000 >
Neighbor Timeout	180	ip pim-sparse nbr-timeout < 60 - 65535 >
LAN Prune Delay	Yes	vlan < vid > ip pim lan-prune-delay
LAN Delay Enabled	No	Shows Yes if all multicast routers on the current VLAN interface enabled LAN-prune-delay. Otherwise shows No .
DR Priority	1	ip pim-sparse dr-priority < 0 - 4294967295 >

Displaying PIM Neighbor Data

These commands enable listings of either all PIM neighbors the router detects or the data for a specific PIM neighbor.

Syntax: show ip pim neighbor

Lists PIM neighbor information for all PIM neighbors connected to the router:

IP Address: Lists the IP address of a neighbor multicast router.

VLAN: Lists the VLAN through which the router connects to the indicated neighbor.

Up Time: Shows the elapsed time during which the neighbor has maintained a PIM route to the router.

Expire Time: Indicates how long before the router ages-out the current flow (group membership). This value decrements until:

- Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.
- Reset by a new flow for the current multicast group on the VLAN.

The timer expires (reaches 0). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.

DR Priority: Shows the currently configured priority for Designated Router (DR) operation on the interface.

```
ProCurve(config)# show ip pim neighbor
```

PIM Neighbors

IP Address	VLAN	Up Time (sec)	Expire Time (sec)	DR Priority
10.10.10.2	100	348	90	1
10.20.10.1	200	410	97	1

Figure 4-18. Example of Output Listing all PIM Neighbors Detected

PIM-SM (Sparse Mode)

Displaying PIM-SM Data and Configuration Settings

Syntax: show ip pim neighbor [< ip-address >]

Lists the same information as show ip pim neighbor (page 3-34) for the specified PIM neighbor.

```
ProCurve(config)# show ip pim neighbor 10.10.10.2

PIM Neighbor

IP Address   : 10.10.10.2
VLAN        : 100

Up Time (sec)      : 678
Expire Time (sec) : 93
DR Priority       : 1
```

Figure 4-19. Example Output for a Specific PIM Neighbor

Displaying BSR Data

The router provides BSR information through both IP PIM and the running configuration.

Displaying BSR Status and Configuration

Syntax: show ip pim bsr

Lists the identity, configuration, and time data of the currently elected BSR for the domain, plus the BSR-candidate configuration, the Candidate-RP configuration and the supported multicast groups on the current router.

```

ProCurve(config)# show ip pim bsr

Status and Counters - PIM-SM Bootstrap Router Information

[E-BSR Address       : 10.10.10.2
E-BSR Priority       : 0
E-BSR Hash Mask Length : 30
E-BSR Up Time       : 53 mins
Next Bootstrap Message : 88 secs

[C-BSR Admin Status   : This system is a Candidate-BSR
C-BSR Address        : 10.10.10.1
C-BSR Priority        : 0
C-BSR Hash Mask Length : 30
C-BSR Message Interval : 60
C-BSR Source IP VLAN  : 100

[C-RP Admin Status   : This system is a Candidate-RP
C-RP Address        : 10.10.10.1
C-RP Hold Time      : 150
C-RP Advertise Period : 60
C-RP Priority        : 192
C-RP Source IP VLAN  : 100

Group Address      Group Mask
-----
224.0.0.0         240.0.0.0
229.0.1.0         255.255.255.0
239.100.128.0    255.255.128.0
    
```

Figure 4-20. Example of Listing BSR Data for the Domain and the Immediate Router

Listing Non-Default BSR Configuration Settings

The **show running** command includes the current non-default BSR configuration settings on the router.

```
ProCurve(config)# show running

Running configuration:
.
.
.
ip routing
snmp-server community "public" Unrestricted
vlan 1
.
.
.
vlan 120
.
.
.
ip multicast-routing
router rip
  exit
router pim
  bsr-candidate
  | bsr-candidate source-ip-vlan 120 |
  | bsr-candidate priority 1       |
  rp-candidate
  rp-candidate source-ip-vlan 120
  rp-candidate group-prefix 224.0.0.0 240.0.0.0
  rp-candidate hold-time 150
  exit
vlan 120
  ip rip 120.10.10.2
  ip pim-sparse
  ip-addr any
  exit
  exit
.
.
.
```

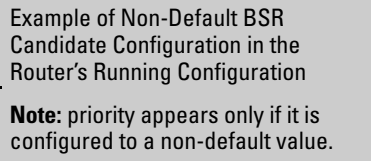


Figure 4-21. Example of Non-Default BSR Configuration Listing

Displaying the Current RP Set

The BSR sends periodic RP updates to all Candidate RPs in the domain. These updates include the set of multicast group data configured on and reported by all Candidate-RPs in the domain. This data does not include any static-RP entries configured on any router in the domain. (To view the static RP-set information for any static-RPs configured on a particular router, you must access the CLI of that specific router.)

Syntax: show ip pim rp-set [learned | static]

Without options, this command displays the multicast group support for both the learned (elected) Candidate-RP assignments and any statically configured RP assignments.

learned: *Displays only the elected Candidate-RP assignments the router has learned from the latest BSR message.*

static: *Displays only the statically configured RP assignment(s) configured on the router.*

```
ProCurve(config)# show ip pim rp-set
```

Status and Counters - PIM-SM Static RP-Set Information			
Group Address	Group Mask	RP Address	Override
231.100.128.255	255.255.255.255	100.10.10.1	Yes

The static RP-set applies only to the current routing switch.

The **Yes** override indicates that the static-RP has precedence over any Candidate-RP routers for supporting the indicated group..

```
Status and Counters - PIM-SM Learned RP-Set Information
```

Group Address	Group Mask	RP Address	Hold Time	Expire Time
231.100.128.0	255.255.240.0	100.10.10.1	150	92
232.240.255.252	255.255.255.252	100.10.10.1	150	92
237.255.248.1	255.255.255.255	100.10.10.1	150	92
239.10.10.240	255.255.255.240	120.10.10.2	150	92
239.10.10.240	255.255.255.252	120.10.10.2	150	92

The Learned RP-set is received from the BSR and includes an aggregation of reports it has received from all accessible candidate-RPs in the domain.

Figure 4-22. Listing Both the Learned and Static RP-Set Data

PIM-SM (Sparse Mode)

Displaying PIM-SM Data and Configuration Settings

```
ProCurve(config)# show ip pim rp-set learned

Status and Counters - PIM-SM Learned RP-Set Information

Group Address      Group Mask          RP Address          Hold Time  Expire Time
-----
231.100.128.0      255.255.240.0      100.10.10.1        150        150
232.240.255.252    255.255.255.252    100.10.10.1        150        150
237.255.248.1      255.255.255.255    100.10.10.1        150        150
239.10.10.240      255.255.255.240    120.10.10.2        150        150
239.10.10.240      255.255.255.252    120.10.10.2        150        150
```

Figure 4-23. Example of Displaying Only the Learned RP-Set Data for the PIM-SM Domain

```
ProCurve(config)# show ip pim rp-set static

Status and Counters - PIM-SM Static RP-Set Information

Group Address      Group Mask          RP Address          Override
-----
231.100.128.255    255.255.255.255    100.10.10.1        Yes
```

Figure 4-24. Example of Displaying only the Static RP-Set Data (Applies to Current Router Only)

Displaying Candidate-RP Data

Displaying the Router's Candidate-RP Status and Configuration

Syntax: show ip pim rp-candidate [config]

rp-candidate: Lists the current Candidate-RP status and, if the status is enabled for C-RP operation, includes the current C-RP configuration on the router.

rp-candidate config: Lists the current Candidate-RP status and the current C-RP configuration on the router, regardless of whether C-RP operation is currently enabled.

```
ProCurve(pim)# show ip pim rp-candidate  
  
This system is not a Candidate-RP
```

Figure 4-25. Example Listing for a Router that is Not Configured as a C-RP

```
ProCurve(pim)# show ip pim rp-candidate config  
  
Status and Counters - PIM-SM Candidate-RP Information  
  
C-RP Admin Status      : This system is not a Candidate-RP  
C-RP Address           : 120.10.10.2  
C-RP Hold Time        : 150  
C-RP Advertise Period : 60  
C-RP Priority          : 192  
C-RP Source IP VLAN   : 120  
  
Group Address  Group Mask  
-----  
239.10.10.240 255.255.255.252
```

Status Line → C-RP Admin Status : This system is not a Candidate-RP

Configuration → C-RP Priority : 192

Indicates that this router is **not** enabled for C-RP operation.

Example of a Candidate-RP configuration for supporting multicast groups in the range of 239.10.10.240 to 239.10.10.243.

Figure 4-26. Example of the Full Candidate-RP Configuration Listing

Listing Non-Default C-RP Configuration Settings

The **show running** command includes the current non-default C-RP configuration settings on the router.

```
ProCurve(config)# show running

Running configuration:
.
.
.
ip routing
snmp-server community "public" Unrestricted
vlan 1
.
.
.
vlan 120
.
.
.
ip multicast-routing
router rip
  exit
router pim
  bsr-candidate
  bsr-candidate source-ip-vlan 120
  bsr-candidate priority 1
  rp-candidate
  rp-candidate source-ip-vlan 120
  rp-candidate group-prefix 224.0.0.0 240.0.0.0
  rp-candidate hold-time 150
  exit
vlan 120
  ip rip 120.10.10.2
  ip pim-sparse
  ip-addr any
.
.
.
```

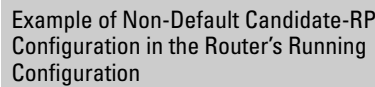


Figure 4-27. Example of Non-Default C-RP Configuration Listing

Operating Notes

Eliminating Redundancy in Support for a Multicast Group. Configuring only one router in a domain as an RP for supporting traffic for a specific multicast group eliminates support redundancy for that group. In this case, if that router becomes unavailable then the group will be excluded from the domain.

Excluding Multicast Groups. If all of the C-RPs and static-RPs (if any) in a domain are configured to exclude some multicast groups or ranges of groups, then multicast traffic for such groups will be dropped when received by a DR, and will not be forwarded to any RP. (Such groups will still be switched locally if IGMP is enabled on the VLAN where the excluded group traffic is received from a multicast traffic source.)

Routing Table Entries. For multicast traffic from a source to the edge router supporting a multicast receiver requesting the traffic, when an SPT forms, the routing table (on the edge router) will contain both of the following for the supported group:

- an (S,G) entry for the source IP address and IP multicast group address supported by the SPT
- an (*,G) entry for the “any” (wildcard) source and (same) multicast group supported by the RP tree

Flow Capacity. The router supports up to 2048 flows. Note that a router acting as a DR or RP will have a significantly higher CPU load than other routers in a PIM-SM domain.

IP Addresses Acquired Through DHCP. PIM-SM operation requires statically configured IP addresses and does not operate with IP addresses acquired from a DHCP server.

Event Log Messages

Message	Meaning
< <i>multicast-addr</i> > / < <i>mask</i> > Inconsistent address and mask.	The mask entered for the specified multicast address does not specify sufficient bits to include the nonzero bits in the mask.
< <i>pkt-type</i> > pkt, src IP < <i>ip-addr</i> > vid < <i>vid-#</i> > (not a nbr)	A PIM packet was received that doesn't have a neighbor.
Bad <parameter-name> in < <i>pkt-type</i> > pkt from IP < <i>ip-addr</i> >	The PIM packet was dropped due to a bad parameter in the packet from the IP address shown.
BSM send to < <i>ip-addr</i> > failed	A BSM (Bootstrap Message) send failed. The IP address shown is the BSM destination address.
Candidate BSR functionality disabled < <i>pkt-type</i> >	Candidate BSR functionality has been disabled.
Candidate RP functionality disabled	Candidate RP functionality has been disabled.
C-RP advertisement send to < <i>ip-addr</i> > failed	A C-RP advertisement send failed. The IP address shown is the destination address of the message.
Enabled as Candidate BSR using address: < <i>ip-addr</i> >	Candidate BSR functionality has been enabled at the indicated IP address.
Enabled as Candidate RP using address: < <i>ip-addr</i> >	Candidate RP functionality has been enabled at the indicated IP address.
Failed alloc of HW < <i>flow</i> > for flow < <i>src-ip-addr</i> >, < <i>multicast-addr</i> >	Hardware resources are consumed and software routing is being done for the flow.
Failed to initialize < <i>pkt-type</i> > as a call back routine	The IP address manager PIM callback routine failed to initialize.
Failed to alloc a < <i>pkt-type</i> > pkt (vid < <i>vid-#</i> >)	Allocation of a packet buffer failed message.
I/F configured with IP < <i>ip-addr</i> > on vid < <i>vid-#</i> >	The IP address on the PIM interface has changed to the indicated address.
I/F removal with IP < <i>ip-addr</i> > on vid < <i>vid-#</i> >	The PIM interface has been removed due to IP address removal or change of the indicated IP address.

Message	Meaning
Illegal operation in BSR state machine	An illegal state/event combination has been detected in the BSR state machine.
Malformed Candidate-RP adv recvd from <ip-addr>	The switch received a malformed C-RP-advertisement.
MCAST MAC add for <mac-addr> failed	The indicated interface could not join the multicast group for PIM packets.
MCAST flow <src-ip-addr>, <multicast-addr> not rteing (rsc low)	A multicast flow has been dropped due to low resources
Multicast Hardware Failed to initialize	The multicast hardware cannot be enabled.
No IP address configured on VID <vid-#>	An IP address is not configured for the indicated interface enabled with PIM.
No route to source/rp <ip-addr>	PIM was unable to find a route to the specified IP address.
No RP for group <ip-addr>	PIM-SM needed an RP for the indicated group address, but none was found.
Inconsistent address and mask	The group prefix needs a route/mask entry, for example, if you want, 224.x.x.x/4, you input 224.0.0.0/4.
Pkt dropped from <ip-addr> <reason>, vid <vid-#>	Received a packet from the indicated IP address and VLAN, and dropped it.
Pkt rcvd with a cksum error from <ip-addr>	A packet arrived from the indicated IP address with a checksum error.
PIM socket error	There was an error regarding the PIM socket, either on a sockopt call or a rcvfrom call.
Rcvd pkt ver# <#>, from <ip-addr>, expected <#>	Received a packet from the indicated IP address with the wrong PIM version number.
Rcvd pkt from rtr <ip-addr>, unkwn pkt type <pkt-type>	Unknown PIM packet type received from the indicated IP address.
Rcvd hello from <ip-addr> on vid <vid-#>	A misconfiguration exists between the routers.
Rcvd incorrect hello from <ip-addr>	An incorrect HELLO packet was received from the indicated IP address.
Rcvd unkwn opt <#> in <pkt-type> pkt from <ip-addr>	A PIM packet with an unknown option number was received from the indicated IP address.

PIM-SM (Sparse Mode)

Event Log Messages

Message	Meaning
Rcvd unkwn addr fmly <i><add-family></i> in <i><pkt-type></i> pkt from <i><ip-addr></i>	A PIM packet with an unknown address family was received.
Rcvd <i><pkt-type></i> pkt with bad len from <i><ip-addr></i>	A PIM packet with an inconsistent length was received from the indicated IP address.
Send error(<i><error-#></i>) on <i><packet-type></i> pkt on VID <i><vid-#></i>	Send packet failed on the indicated VLAN.
Static RP configuration failure: <i><src-ip-addr></i> , <i><multicast-addr></i>	The configuration of a static RP for the indicated multicast group has failed on the indicated interface.
Unable to alloc a buf of size <i><size></i> for <i><memory element></i>	PIM_DM could not allocate memory for the indicated buffer.
Unable to alloc a msg buffer for <i><system-event></i>	Informs the user that a message buffer could not be allocated for the indicated system event.
Unable to allocate <i><table-type></i> table	The PIM interface has been removed due to an IP address removal or change.
Unexpected state/event <i><state>/<event></i> in <i><statemachine></i> statemach	PIM received an event type in a state that was not expected.
VLAN is not configured for IP.	A VLAN must be statically configured with a primary IP address before enabling PIM-SM on that VLAN. If the VLAN has no IP address or is configured to acquire a primary IP address by using DHCP/Bootp, it cannot be configured to support PIM-SM.

IP Routing Features

Contents

Overview of IP Routing	5-6
IP Interfaces	5-7
IP Tables and Caches	5-7
ARP Cache Table	5-8
IP Route Table	5-8
IP Forwarding Cache	5-9
IP Route Exchange Protocols	5-10
IP Global Parameters for Routing Switches	5-10
IP Interface Parameters for Routing Switches	5-12
Configuring IP Parameters for Routing Switches	5-13
Configuring IP Addresses	5-13
Changing the Router ID	5-13
Configuring ARP Parameters	5-15
How ARP Works	5-15
Enabling Proxy ARP	5-17
Configuring Forwarding Parameters	5-17
Changing the TTL Threshold	5-18
Enabling Forwarding of Directed Broadcasts	5-18
Configuring ICMP	5-19
Disabling ICMP Messages	5-19
Disabling Replies to Broadcast Ping Requests	5-19
Disabling ICMP Destination Unreachable Messages	5-20
Disabling ICMP Redirects	5-21
Configuring Static IP Routes	5-21
Static Route Types	5-21
Other Sources of Routes in the Routing Table	5-22
Static IP Route Parameters	5-22

Static Route States Follow VLAN States	5-23
Configuring a Static IP Route	5-23
Displaying Static Route Information	5-25
Configuring the Default Route	5-25
Configuring RIP	5-26
Overview of RIP	5-26
RIP Parameters and Defaults	5-27
RIP Global Parameters	5-27
RIP Interface Parameters	5-27
Configuring RIP Parameters	5-28
Enabling RIP	5-28
Enabling IP RIP on a VLAN	5-29
Changing the RIP Type on a VLAN Interface	5-29
Changing the Cost of Routes Learned on a VLAN Interface	5-29
Configuring RIP Redistribution	5-30
Define RIP Redistribution Filters	5-30
Modify Default Metric for Redistribution	5-31
Enable RIP Route Redistribution	5-31
Displaying RIP Information	5-33
Displaying General RIP Information	5-34
Displaying RIP Interface Information	5-35
Displaying RIP Peer Information	5-37
Displaying RIP Redistribution Information	5-38
Displaying RIP Redistribution Filter (restrict) Information	5-39
Configuring OSPF	5-40
Terminology	5-41
Overview of OSPF	5-43
OSPF Router Types	5-44
Interior Routers	5-44
Area Border Routers (ABRs)	5-44
Autonomous System Boundary Router (ASBR)	5-45
Designated Routers	5-45
OSPF Area Types	5-47
Backbone Area	5-48
Normal Area	5-48

Not-So-Stubby-Area (NSSA)	5-49
Stub Area	5-50
OSPF RFC Compliance	5-50
Reducing AS External LSAs and Type-3 Summary LSAs	5-50
Algorithm for AS External LSA Reduction	5-51
Replacing Type-3 Summary LSAs and Type-7 Default External LSAs with a Type-3 Default Route LSA	5-52
Dynamic OSPF Activation and Configuration	5-52
General Configuration Steps for OSPF	5-53
Configuration Rules	5-54
OSPF Global and Interface Settings	5-54
Configuring OSPF on the Routing Switch	5-56
1. Enable IP Routing	5-56
2. Enable Global OSPF Routing	5-56
3. Changing the RFC 1583 OSPF Compliance Setting	5-57
4. Assign the Routing Switch to OSPF Areas	5-59
5. Assign VLANs and/or Subnets to Each Area	5-64
6. Optional: Assigning Loopback Addresses to an Area	5-65
Changing the Route Loop Prevention Method	5-33
7. Optional: Configure for External Route Redistribution in an OSPF Domain	5-67
8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone	5-70
9. Optional: Influence Route Choices by Changing the Administrative Distance Default	5-73
10. Optional: Change OSPF Trap Generation Choices	5-74
11. Optional: Adjust Performance by Changing the VLAN or Subnet Interface Settings	5-75
12. Optional: Configuring OSPF Interface Authentication	5-80
13. Configuring an ABR To Use a Virtual Link to the Backbone	5-82
Configuring a Virtual Link	5-83
Optional: Adjust Virtual Link Performance by Changing the Interface Settings	5-85
Configuring OSPF Authentication on a Virtual Link	5-88
Displaying OSPF Information	5-90
Displaying General OSPF Configuration Information	5-91
Displaying OSPF Area Information	5-92

Displaying OSPF External Link State Information	5-93
Displaying OSPF Interface Information	5-95
Displaying OSPF Interface Information for a Specific VLAN or IP Address	5-97
Displaying OSPF Link State Information	5-98
Displaying OSPF Neighbor Information	5-101
Displaying OSPF Redistribution Information	5-103
Displaying OSPF Redistribution Filter (restrict) Information .	5-103
Displaying OSPF Virtual Neighbor Information	5-104
Displaying OSPF Virtual Link Information	5-105
Displaying OSPF Route Information	5-107
Displaying OSPF Traps Enabled	5-109
OSPF Equal-Cost Multipath (ECMP) for Different Subnets Available Through the Same Next-Hop Routes	5-109
Displaying the Current IP Load-Sharing Configuration	5-110
Configuring IRDP	5-112
Enabling IRDP Globally	5-113
Enabling IRDP on an Individual VLAN Interface	5-113
Displaying IRDP Information	5-114
Configuring DHCP Relay	5-115
Overview	5-115
DHCP Packet Forwarding	5-115
Unicast Forwarding	5-115
Broadcast Forwarding	5-116
Prerequisites for DHCP Relay Operation	5-116
Enabling DHCP Relay	5-116
Configuring an IP Helper Address	5-117
Operating Notes	5-117
Hop Count in DHCP Requests	5-117
Disabling the Hop Count in DHCP Requests	5-117
Operating Notes	5-118
Verifying the DHCP Relay Configuration	5-118
Displaying the DHCP Relay Setting	5-118
Displaying DHCP Helper Addresses	5-119
Displaying the Hop Count Setting	5-120

DHCP Option 82	5-120
Option 82 Server Support	5-121
Terminology	5-122
General DHCP Option 82 Requirements and Operation	5-123
Option 82 Field Content	5-124
Forwarding Policies	5-127
Configuration Options for Managing DHCP Client Request Packets	5-127
Multiple Option 82 Relay Agents in a Client Request Path	5-128
Validation of Server Response Packets	5-129
Multinetted VLANs	5-131
Configuring Option 82	5-131
Example of Option 82 Configuration	5-133
Operating Notes	5-135
UDP Broadcast Forwarding	5-137
Overview	5-137
Subnet Masking for UDP Forwarding Addresses	5-138
Configuring and Enabling UDP Broadcast Forwarding	5-139
Globally Enabling UDP Broadcast Forwarding	5-139
Configuring UDP Broadcast Forwarding on Individual VLANs	5-139
Displaying the Current IP Forward-Protocol Configuration	5-141
Operating Notes for UDP Broadcast Forwarding	5-142
Messages Related to UDP Broadcast Forwarding	5-142

Overview of IP Routing

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

- **IP Static Routes** – up to 256 static routes
- **RIP** (Router Information Protocol) – supports RIP Version 1, Version 1 compatible with Version 2 (default), and Version 2
- **OSPF** (Open Shortest Path First) – the standard routing protocol for handling larger routed networks
- **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, such as Routing Information Protocol (RIP).

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 5-15.

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
- A route learned through RIP
- A route learned through OSPF

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	Gateway	VLAN	Type	Sub-Type	Metric	D
10.10.10.1/32	10.10.12.1		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, and for OSPF routes, the sub type, and the route’s IP metric (cost). 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 5-23

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 Route Exchange Protocols

The switch supports the following IP route exchange protocols:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

These protocols provide routes to the IP route table. You can use one or more of these protocols, in any combination. The protocols are disabled by default. For configuration information, see the following:

- “Configuring RIP” on page 5-26
- “Configuring OSPF” on page 5-40

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 5-1. IP Global Parameters for Routing Switches

Parameter	Description	Default	See page
Router ID	The value that routers use to identify themselves to other routers when exchanging route information. OSPF uses the router ID to identify routers. RIP does not use the router ID.	The lowest-numbered IP address configured on the lowest-numbered routing interface.	5-13
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	5-15
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. (Can be set using the menu interface to be as long as 240 minutes. Go to Menu > Switch Configuration > IP Config.)	Five minutes.	n/a

Parameter	Description	Default	See page
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	5-17
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 5-2 on page 5-12.	Disabled	5-18
ICMP Router Discovery Protocol (IRDP)	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. 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. <ul style="list-style-type: none"> • Forwarding method (broadcast or multicast) • Hold time • Maximum advertisement interval • Minimum advertisement interval • Router preference level 	Disabled	5-112 5-113
Static route	An IP route you place in the IP route table.	No entries	5-21
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	5-25

IP Interface Parameters for Routing Switches

5-2 lists the interface-level IP parameters for routing switches.

Table 5-2. 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	*
Metric	A numeric cost the router adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1 (one)	5-28
ICMP Router Discovery Protocol (IRDP)	Locally overrides the global IRDP settings. See table 5-1 on page 5-10 for global IRDP information.	Disabled	5-113
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	5-117

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

Changing the Router ID

In most configurations, a routing switch has multiple IP addresses, usually configured on different VLAN interfaces. As a result, a routing switch’s identity to other devices varies depending on the interface to which the other device is attached. Some routing protocols, including Open Shortest Path First (OSPF), identify a routing switch by just one of the IP addresses configured on the routing switch, regardless of the interfaces that connect the routing switches. This IP address is the router ID.

Note

Routing Information Protocol (RIP) does not use the router ID.

If no router ID is configured, then, by default, the router ID on a ProCurve routing switch is the first IP address that becomes physically active at reboot. This is usually the lowest numbered IP interface configured on the device. However, if no router ID is configured and one or more user-configured loopback interfaces are detected at reboot, then the lowest-numbered (user-configured) loopback interface becomes the router ID. If the lowest-numbered loopback interface has multiple IP addresses, then the lowest of these addressees will be selected as the router ID. Once a router ID is selected, it will not automatically change unless a higher-priority interface is configured

IP Routing Features

Configuring IP Parameters for Routing Switches

on the routing switch *and* OSPF is restarted with a reboot. (User-Configured loopback interfaces are always higher priority than other configured interfaces.) However, you prefer, you can explicitly set the router ID to any valid IP address, as long as the IP address is not in use on another device in the network.

Note

To display the router ID, enter the **show ip ospf** CLI command at any Manager EXEC CLI level.

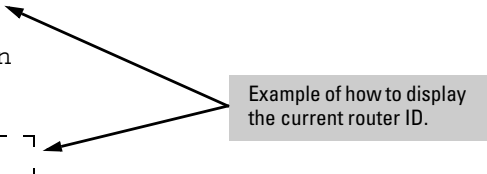
```
ProCurve(ospf)# show ip ospf

OSPF Configuration Information

 OSPF_protocol : enabled
 Router ID     : 10.10.10.1
-----

Currently defined areas:

Area ID          Type      Stub      Stub      Stub
-----          -
backbone         normal  1          send      ospf metric
0.0.0.2          nssa   10         send      external type 2
0.0.0.3          stub   2          send      ospf metric
0.0.0.4          stub   10         send      ospf metric
```



Reconfiguring the Router ID (Optional). If you want to change the router ID setting, do the following:

1. Go to the global config context. When you do so, the CLI prompt will appear similar to the following:
ProCurve(config)#_
2. If OSPF is not enabled, go to step 3. But if OSPF is enabled, then use **no router ospf** to disable OSPF operation.
3. Use **ip router-id <ip-addr>** to specify a new router ID. (This IP address must be unique in the routing switch configuration.)
4. If you disabled OSPF operation (step 2), then use **router ospf** to re-enable OSPF operation.

For more information on the router ID, refer to “IP Global Parameters for Routing Switches” on page 5-10 and “Changing the Router ID” on page 5-13.

To change the router ID, enter a command such as the following:

```
ProCurve(config)# ip router-id 209.157.22.26
```

Syntax: Syntax: `ip router-id <ip-addr>`

The `<ip-addr>` can be any valid, unique IP address.

Note

You can specify an IP address used for an interface on the ProCurve routing switch, but do not specify an IP address in use by another device.

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 5-2 on page 5-25 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.
- **RIP:** If RIP is enabled, the routing switch can learn about routes from the advertisements other RIP routers send to the routing switch. If the RIP route has a lower administrative distance than any other routes from different sources to the same destination, the routing switch places the route in the IP route table. (Refer to “Administrative Distance” on page 5-9.)
- **OSPF:** See RIP, but substitute “OSPF” for “RIP”.
- **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 5-25.

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 5-9):** 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, routing switch "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. Routing switch "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 > [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.

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 5-9. (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 5-1. 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 5-2 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 5-21

Figure 5-2. 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 RIP

This section describes how to configure RIP using the CLI interface.

To display RIP configuration information and statistics, see “Displaying RIP Information” on page 5-33.

Overview of RIP

Routing Information Protocol (RIP) is an IP route exchange protocol that uses a *distance vector* (a number representing distance) to measure the cost of a given route. The cost is a distance vector because the cost often is equivalent to the number of router hops between the ProCurve routing switch and the destination network.

A ProCurve routing switch can receive multiple paths to a destination. The software evaluates the paths, selects the best path, and saves the path in the IP route table as the route to the destination. Typically, the best path is the path with the fewest hops. A hop is another router through which packets must travel to reach the destination. If the ProCurve routing switch receives a RIP update from another router that contains a path with fewer hops than the path stored in the ProCurve routing switch's route table, the routing switch replaces the older route with the newer one. The routing switch then includes the new path in the updates it sends to other RIP routers, including ProCurve routing switches.

RIP routers, including ProCurve routing switches, also can modify a route's cost, generally by adding to it, to bias the selection of a route for a given destination. In this case, the actual number of router hops may be the same, but the route has an administratively higher cost and is thus less likely to be used than other, lower-cost routes. A RIP route can have a maximum cost of 15. Any destination with a higher cost is considered unreachable. Although limiting to larger networks, the low maximum hop count prevents endless loops in the network.

The switches covered in this guide support the following RIP types:

- Version 1
- V1 compatible with V2
- Version 2 (the default)

Note

ICMP Host Unreachable Message for Undeliverable ARPs. If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP timeout 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.

RIP Parameters and Defaults

The following tables list the RIP parameters, their default values, and where to find configuration information.

RIP Global Parameters

5-3 lists the global RIP parameters and their default values.

Table 5-3. RIP Global Parameters

Parameter	Description	Default
RIP state	Routing Information Protocol V2-only.	Disabled
auto-summary	Enable/Disable advertisement of summarized routes.	Enabled
metric	Default metric for imported routes.	1
redistribution	RIP can redistribute static, connected, and OSPF routes. (RIP redistributes connected routes by default, when RIP is enabled.)	Disabled

RIP Interface Parameters

5-4 lists the VLAN interface RIP parameters and their default values.

Table 5-4. RIP Interface Parameters

Parameter	Description	Default
RIP version	The version of the protocol that is supported on the interface. The version can be one of the following: <ul style="list-style-type: none"> • Version 1 only • Version 2 only • Version 1 or version 2 	V2-only

Parameter	Description	Default
metric	A numeric cost the routing switch adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1
IP address	The routes that a routing switch learns or advertises can be controlled.	The routing switch learns and advertises all RIP routes on all RIP interfaces
loop prevention	The method the routing switch uses to prevent routing loops caused by advertising a route on the same interface as the one on which the routing switch learned the route. <ul style="list-style-type: none">• Split horizon - the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.• Poison reverse - the routing switch assigns a cost of 16 ("infinite" or "unreachable") to a route before advertising it on the same interface as the one on which the routing switch learned the route.	Poison reverse
receive	Define the RIP version for incoming packets	V2-only
send	Define the RIP version for outgoing packets	V2-only

Configuring RIP Parameters

Use the following procedures to configure RIP parameters on a system-wide and individual VLAN interface basis.

Enabling RIP

RIP is disabled by default. To enable it, use one of the following methods. When you enable RIP, the default RIP version is **RIPv2-only**. You can change the RIP version on an individual interface basis to **RIPv1** or **RIPv1-or-v2** if needed.

To enable RIP on a routing switch, enter the following commands:

```
ProCurve(config)# ip routing
ProCurve(config)# router rip
ProCurve(rip)# exit
ProCurve(config)# write memory
```

Syntax: [no] router rip

Note

IP routing must be enabled prior to enabling RIP. The first command in the preceding sequence enables IP routing.

Enabling IP RIP on a VLAN

To enable RIP on all IP addresses in a VLAN, use **ip rip** in the VLAN context. When the command is entered without specifying any IP address, it is enabled in all configured IP addresses of the VLAN.

To enable RIP on a specific IP address in a VLAN, use **ip rip [*ip-addr*>|all]** in the VLAN context and enter a specific IP address. If you want RIP enabled on all IP addresses, you can specify **all** in the command instead of a specific IP address.

Changing the RIP Type on a VLAN Interface

When you enable RIP on a VLAN interface, **RIPv2-only** is enabled by default. You can change the RIP type to one of the following on an individual VLAN interface basis:

- Version 1 only
- Version 2 only (the default)
- Version 1 - or - version 2

To change the RIP type supported on a VLAN interface, enter commands such as the following:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# ip rip v1-only
ProCurve(vlan-1)# exit
ProCurve(config)# write memory
```

Syntax: [no] ip rip <v1-only | v1-or-v2 | v2-only >

Changing the Cost of Routes Learned on a VLAN Interface

By default, the switch interface increases the cost of a RIP route that is learned on the interface. The switch increases the cost by adding one to the route's metric before storing the route.

You can change the amount that an individual VLAN interface adds to the metric of RIP routes learned on the interface.

Note

RIP considers a route with a metric of 16 to be unreachable. Use this metric only if you do not want the route to be used. In fact, you can prevent the switch from using a specific interface for routes learned through that interface by setting its metric to 16.

To increase the cost a VLAN interface adds to RIP routes learned on that interface, enter commands such as the following:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# ip rip metric 5
```

These commands configure vlan-1 to add 5 to the cost of each route learned on the interface.

Syntax: ip rip metric < 1-16 >

Configuring RIP Redistribution

You can configure the routing switch to redistribute connected, static, and OSPF routes into RIP. When you redistribute a route into RIP, the routing switch can use RIP to advertise the route to its RIP neighbors.

To configure redistribution, perform the following tasks:

1. Configure redistribution filters to permit or deny redistribution for a route based on the destination network address or interface. (optional)
2. Enable redistribution

Define RIP Redistribution Filters

Route redistribution imports and translates different protocol routes into a specified protocol type. On the switches covered in this guide, redistribution is supported for static routes, directly connected routes, and OSPF routes. Redistribution of any other routing protocol into RIP is not currently supported. When you configure redistribution for RIP, you can specify that static, connected, or OSPF routes are imported into RIP routes. Likewise, OSPF redistribution supports the import of static, connected, or RIP routes into OSPF routes.

To configure for redistribution, define the redistribution tables with “restrict” redistribution filters. In the CLI, use the **restrict** command for RIP at the RIP router level.

Note

Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

Example: To configure the switch to filter out redistribution of static, connected, or OSPF routes on network 10.0.0.0, enter the following commands:

```
ProCurve(config)# router rip
ProCurve(rip)# restrict 10.0.0.0 255.0.0.0
ProCurve(rip)# write memory
```

Note

The default configuration permits redistribution for all default connected routes only.

Syntax: `restrict < ip-addr > < ip-mask > | < ip-addr /< prefix length >`

This command prevents any routes with a destination address that is included in the range specified by the address/mask pair from being redistributed by RIP.

Modify Default Metric for Redistribution

The default metric is a global parameter that specifies the cost applied to all RIP routes by default. The default value is 1. You can assign a cost from 1 – 15.

Example: To assign a default metric of 4 to all routes imported into RIP, enter the following commands:

```
ProCurve(config)# router rip
ProCurve(rip)# default-metric 4
```

Syntax: `default-metric < value >`

The `< value >` can be from 1 – 15. The default is 1.

Enable RIP Route Redistribution

Note

Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

To enable redistribution of connected and static IP routes into RIP, enter the following commands.

```
0(config)# router rip
ProCurve(rip)# redistribute connected
ProCurve(rip)# redistribute static
ProCurve(rip)# redistribute ospf
ProCurve(rip)# write memory
```

Syntax: [no] redistribute < connected | static | ospf >

Changing the Route Loop Prevention Method

RIP can use the following methods to prevent routing loops:

- **Split horizon** - the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.
- **Poison reverse** - the routing switch assigns a cost of 16 (“infinity” or “unreachable”) to a route before advertising it on the same interface as the one on which the routing switch learned the route. This is the default.

These loop prevention methods are configurable on an individual VLAN interface basis.

Note

These methods are in addition to RIP's maximum valid route cost of 15.

Poison reverse is enabled by default. Disabling poison reverse causes the routing switch to revert to **Split horizon**. (Poison reverse is an extension of Split horizon.) To disable Poison reverse on an interface, and thereby enable Split horizon, enter the following:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# no ip rip poison-reverse
```

Syntax: [no] ip rip poison-reverse

Entering the command without the “no” option will re-enable Poison reverse.

Displaying RIP Information

All RIP configuration and status information is shown by the CLI command **show ip rip** and options off that command. The following RIP information can be displayed:

RIP Information Type	Page
General Information	5-34
Interface Information	5-35
Peer Information	5-37
Redistribute Information	5-38
Restrict Information	5-39

Displaying General RIP Information

To display general RIP information, enter **show ip rip** at any context level. The resulting display will appear similar to the following:

```
ProCurve(config)# show ip rip

RIP global parameters

RIP protocol      : enabled
Auto-summary     : enabled
Default Metric   : 4
Distance         : 120
Route changes    : 0
Queries          : 0

RIP interface information

IP Address      Status      Send mode      Recv mode      Metric      Auth
-----
100.1.0.1       enabled    V2-only        V2-only         5           none
100.2.0.1       enabled    V2-only        V2-only         5           none
100.3.0.1       enabled    V2-only        V2-only         5           none
100.4.0.1       enabled    V2-only        V2-only         5           none
100.10.0.1      enabled    V2-only        V2-only         5           none
100.11.0.1      enabled    V2-only        V2-only         5           none
100.12.0.1      enabled    V2-only        V2-only         5           none

RIP peer information

IP Address      Bad routes  Last update timeticks
-----
```

Figure 5-3. Example of General RIP Information Listing

The display is a summary of Global RIP information, information about interfaces with RIP enabled, and information about RIP peers. The following fields are displayed:

- **RIP protocol** – Status of the RIP protocol on the router. RIP must be enabled here and on the VLAN interface for RIP to be active. The default is **disabled**.
- **Auto-summary** – Status of Auto-summary for all interfaces running RIP. If auto-summary is enabled, then subnets will be summarized to a class network when advertising outside of the given network.
- **Default Metric** – Sets the default metric for imported routes. This is the metric that will be advertised with the imported route to other RIP peers. A RIP metric is a measurement used to determine the 'best' path to network; 1 is the best, 15 is the worse, 16 is unreachable.
- **Route changes** – The number of times RIP has modified the routing switch's routing table.

- **Queries** – The number of RIP queries that have been received by the routing switch.
- **RIP Interface Information** – RIP information on the VLAN interfaces on which RIP is enabled.
 - **IP Address** – IP address of the VLAN interface running rip.
 - **Status** – Status of RIP on the VLAN interface.
 - **Send mode** – The format of the RIP updates: RIP 1, RIP 2, or RIP 2 version 1 compatible.
 - **Recv mode** – The switch can process RIP 1, RIP 2, or RIP 2 version 1 compatible update messages.
 - **Metric** – The path “cost”, a measurement used to determine the 'best' RIP route path; 1 is the best, 15 is the worse, 16 is unreachable.
 - **Auth** – RIP messages can be required to include an authentication key if enabled on the interface.
- **RIP Peer Information** – RIP Peers are neighboring routers from which the routing switch has received RIP updates.
 - **IP Address** – IP address of the RIP neighbor.
 - **Bad routes** – The number of route entries which were not processed for any reason.
 - **Last update timeticks** – How many seconds have passed since we received an update from this neighbor.

Syntax: show ip rip

Displaying RIP Interface Information

To display RIP interface information, enter the show ip rip interface command at any context level. The resulting display will appear similar to the following:

```
ProCurve# show ip rip interface
```

RIP interface information					
IP Address	Status	Send mode	Recv mode	Metric	Auth
100.1.0.1	enabled	V2-only	V2-only	1	none
100.2.0.1	enabled	V2-only	V2-only	1	none
100.3.0.1	enabled	V2-only	V2-only	1	none
100.4.0.1	enabled	V2-only	V2-only	1	none

Figure 5-4.Example of Show IP RIP Interface Output

See “RIP Interface Information” on the previous page for definitions of these fields.

You can also display the information for a single RIP VLAN interface, by specifying the VLAN ID for the interface, or specifying the IP address for the interface.

Displaying RIP interface information by VLAN ID: For example, to show the RIP interface information for VLAN 1000, use the **show ip rip interface vlan < vid >** command.

```
ProCurve# show ip rip interface vlan 4
RIP configuration and statistics for VLAN 4
RIP interface information for 100.4.0.1
  IP Address : 100.4.0.1
  Status      : enabled
  Send mode   : V2-only
  Recv mode   : V2-only
  Metric      : 1
  Auth        : none
  Bad packets received : 0
  Bad routes received  : 0
  Sent updates : 0
```

Figure 5-5. Example of RIP Interface Output by VLAN

The information in this display includes the following fields, which are defined under ““RIP Interface Information” on page 5-35: **IP Address**, **Status**, **Send mode**, **Recv mode**, **Metric**, and **Auth**.

The information also includes the following fields:

- **Bad packets received** – The number of packets that were received on this interface and were not processed for any reason.
- **Bad routes received** – The number of route entries that were received on this interface and were not processed for any reason.
- **Sent updates** – The number of RIP routing updates that have been sent on this interface.

Displaying RIP interface information by IP Address: For example, to show the RIP interface information for the interface with IP address 100.2.0.1, enter the **show ip rip interface** command as shown below:

```
ProCurve# show ip rip interface 100.2.0.1

RIP interface information for 100.2.0.1

  IP Address : 100.2.0.1
  Status      : enabled

  Send mode   : V2-only
  Recv mode   : V2-only
  Metric      : 1
  Auth        : none

  Bad packets received : 0
  Bad routes received  : 0
  Sent updates         : 0
```

Figure 5-6. Example of Show IP RIP Interface Output by IP Address

The information shown in this display has the same fields as for the display for a specific VLAN ID. See the previous page for the definitions of these fields.

Syntax: show ip rip interface [*ip-addr* | vlan < *vlan-id* >]

Displaying RIP Peer Information

To display RIP peer information, enter the **show ip rip peer** command at any context level.

The resulting display will appear similar to the following:

```
ProCurve# show ip rip peer

RIP peer information

  IP Address      Bad routes  Last update timeticks
  -----
  100.1.0.100     0           1
  100.2.0.100     0           0
  100.3.0.100     0           2
  100.10.0.100    0           1
```

Figure 5-7. Example of Show IP RIP Peer Output

This display lists all neighboring routers from which the routing switch has received RIP updates. The following fields are displayed:

- **IP Address** – IP address of the RIP peer neighbor.

- **Bad routes** – The number of route entries that were not processed for any reason.
- **Last update timeticks** – How many seconds have passed since the routing switch received an update from this peer neighbor.

Displaying RIP information for a specific peer: For example, to show the RIP peer information for the peer with IP address 100.1.0.100, enter **show ip rip peer 100.1.0.100**.

```
ProCurve# show ip rip peer 100.0.1.100

RIP peer information for 100.0.1.100

  IP Address : 100.1.0.100

  Bad routes  : 0
  Last update timeticks : 2
```

Figure 5-8. Example of Show IP RIP Peer <ip-addr> Output

This display lists the following information for a specific RIP peer:

- **IP Address** – IP address of the RIP peer neighbor.
- **Bad routes** – The number of route entries which were not processed for any reason.
- **Last update timeticks** – How many seconds have passed since the routing switch received an update from this neighbor.

Displaying RIP Redistribution Information

To display RIP redistribution information, enter the **show ip rip redistribute** command at any context level:

```
ProCurve# show ip rip redistribute

RIP redistributing

Route type Status
-----
connected enabled
static      disabled
ospf       disabled
```

Figure 5-9. Example of Show IP RIP Redistribute Output

RIP automatically redistributes connected routes that are configured on interfaces that are running RIP, and all routes that are learned via RIP. The **router rip redistribute** command, described on page 5-30, configures the routing switch to cause RIP to advertise connected routes that are not running RIP, static routes, and OSPF routes. The display shows whether RIP redistribution is enabled or disabled for connected, static, and OSPF routes.

Displaying RIP Redistribution Filter (restrict) Information

To display RIP restrict filter information, enter the **show ip rip restrict** command at any context level:

```
ProCurve# show ip rip restrict
RIP restrict list
  IP Address      Mask
  -----

```

Figure 5-10. Example of Show IP RIP Restrict Output

The display shows if any routes, identified by the IP Address and Mask fields are being restricted from redistribution. The restrict filters are configured by the **router ip restrict** command described on page 5-30.

Configuring OSPF

Feature	Default	Page
Enable IP Routing and Global OSPF Routing	disabled	5-56
Changing the RFC 1583 OSPF Compliance Setting	enabled	5-57
Assign the Routing Switch to OSPF Areas	n/a	5-59
Assign VLANs and/or Subnets to Each Area	n/a	5-64
External Route Redistribution	disabled	5-67
Configure Ranges on an ABR To Reduce Advertising	n/a	5-70
Use Administrative Distance To Influence Route Choices		5-73
Generate OSPF Traps	enabled	5-74
Cost Per Interface	1	5-76
Dead Interval Per Interface	40 sec.	5-76
Hello Interval Per Interface	10 sec.	5-77
Priority Per interface	1	5-77
Retransmit Interval Per Interface	5 sec.	5-78
Transit Delay Per Interface	1 sec.	5-78
Password and MD5 Authentication	disabled	5-80, 5-81
Virtual Link Configuration	n/a	5-82
Dead Interval on a Virtual Link	40 sec.	5-85
Hello Interval on a Virtual Link	10 sec.	5-86
Retransmit Interval on a Virtual Link	5 sec.	5-86
Transit Delay on a Virtual Link	1 sec.	5-87
Password and MD5 Authentication on a Virtual Link	disabled	5-88, 5-89
Displaying OSPF Information	n/a	5-90

This section describes how to configure OSPF using the CLI interface.

Terminology

Area Border Router (ABR): An OSPF-enabled router having interfaces on two or more OSPF areas. (Refer to “Area Border Routers (ABRs)” on page 5-44.)

Autonomous System (AS): A single interior gateway protocol (IGP) domain such as an OSPF or RIP domain.

Autonomous System Boundary Router (ASBR): An OSPF-enabled router having interfaces in multiple IGP domains, such as an ASBR with membership in both a normal area of an OSPF domain and a RIP domain. (Refer to “Autonomous System Boundary Router (ASBR)” on page 5-45.)

Backbone Area: Required in any OSPF domain, this is the transit area for all advertisements and routed traffic between non-backbone areas. (Refer to “Backbone Area” on page 5-48.)

Backup Designated Router (BDR): If the DR for a network becomes inaccessible, the BDR takes over the DR function. (See also “Designated Router”, below, and refer to “Designated Routers” on page 5-45.)

Default Route: A route defined as 0.0.0.0/0. OSPF uses type 3 (summary) defaults and type 7 (external) default routes.

Designated Router (DR): Used in networks having two or more routers, and serves as the distribution point for forwarding updates throughout the network. (See also “Backup Designated Router”, above, and refer to “Designated Routers” on page 5-45.)

External Type-5 Link-State Advertisement: An LSA summarizing known external links for the backbone and normal areas. Refer to Table 5-5 on page 5-43. (See also “Link State Advertisement”.)

External Type-7 Link State Advertisement: An LSA originating with an ASBR in an NSSA and allowed only in the NSSA. Refer to Table 5-5 on page 5-43. (See also “Link State Advertisement”.)

Interior Gateway Protocol (IGP): A method for forwarding traffic between autonomous routing domains. Commonly used between OSPF and RIP domains.

Interior Router: An OSPF-enabled routing switch having interfaces in only one OSPF area. (Refer to “Interior Routers” on page 5-44.)

Link-State Advertisement (LSA): A message sent by a router to its neighbors to advertise the existence of a route to a destination known by the originating router. Refer to Table 5-5 on page 5-43.

Normal Area: Exists within an OSPF domain and connects to the backbone area through one or more ABRs (either physically or through a virtual link). Supports summary link-state advertisements and external link-state advertisements to and from the backbone area, as well as ASBRs.

NSSA (Not-So-Stubby-Area): An OSPF area that limits advertisement of external and summary routes to the backbone area and allows controls on advertisements entering the area from the backbone. (Refer to “Not-So-Stubby-Area (NSSA)” on page 5-49.)

Stub Area: An OSPF area that does not allow an internal ASBR or external type-5 LSAs. (Refer to “Stub Area” on page 5-50.)

Summary Link-State Advertisement: A type-3 LSA summarizing the available links within an OSPF area. This advertisement is sent by the ABR for an area to the backbone area for distribution to the other areas in the OSPF domain. Refer to Table 5-5 on page 5-43. (See also “Link State Advertisement”.)

Type-3 LSA: See “Summary Link-State Advertisement”.

Type-5 LSA: See “External Type-5 Link State Advertisement”.

Type-7 LSA: See “External Type-7 Link State Advertisement”.

Topological Database: See “Link State Database”.

Virtual Link: Used to provide connectivity from a normal area to the backbone when the subject area does not have an ABR physically linked to the backbone area. Refer to “13. Configuring an ABR To Use a Virtual Link to the Backbone” on page 5-82.

Overview of OSPF

OSPF is a link-state routing protocol applied to routers grouped into OSPF areas identified by the routing configuration on each routing switch. The protocol uses link-state advertisements (LSAs) transmitted by each router to update neighboring routers regarding its interfaces and the routes available through those interfaces. Each routing switch in an area also maintains a Link State Database (LSDB) that describes the area topology. (All routers in a given OSPF area have identical LSDBs.) The routing switches used to connect areas to each other flood summary link LSAs and external link LSAs to neighboring OSPF areas to update them regarding available routes. Through this means, each OSPF router determines the shortest path between itself and a desired destination router in the same OSPF domain (Autonomous System). Routed traffic in an OSPF AS is classified as one of the following:

- intra-area traffic
- inter-area traffic
- external traffic

The switches covered in this guide support the following types of LSAs, which are described in RFCs 2328 and 3101:

Table 5-5. OSPF LSA Types

LSA Type	LSA Name	Use
1	Router Link	Describes the state of each interface on a router for a given area. Not propagated to backbone area.
2	Network Link	Describes the OSPF routers in a given network. Not propagated to backbone area.
3	Summary Link	Describes the route to networks in another OSPF area of the same Autonomous System (AS). Propagated through backbone area to other areas.
4	Autonomous System (AS) Summary Link	Describes the route to an ASBR in an OSPF Normal or Backbone area of the same AS. Propagated through backbone area to other areas.
5	AS External Link	Describes the route to a destination in another AS (external route). Originated by ASBR in normal or backbone areas of an AS and propagates through backbone area to other normal areas. For injection into an NSSA, ABR converts type-5 LSAs to a type-7 LSA advertising the default route (0.0.0.0/0).
7	AS External Link in an NSSA Area	Describes the route to a destination in another AS (external route). Originated by ASBR in NSSA. ABR converts type-7 LSAs to type-5 LSAs for injection into the backbone area.

OSPF Router Types

Interior Routers

This type of OSPF router belongs to only one area. Interior routers flood type-1 LSAs to all routers in the same area, and maintain identical link state databases (LSDBs). In figure 5-11, below, routers R1, R3, R4, and R6 are all interior routers because all of their links are to other routers in the same area.

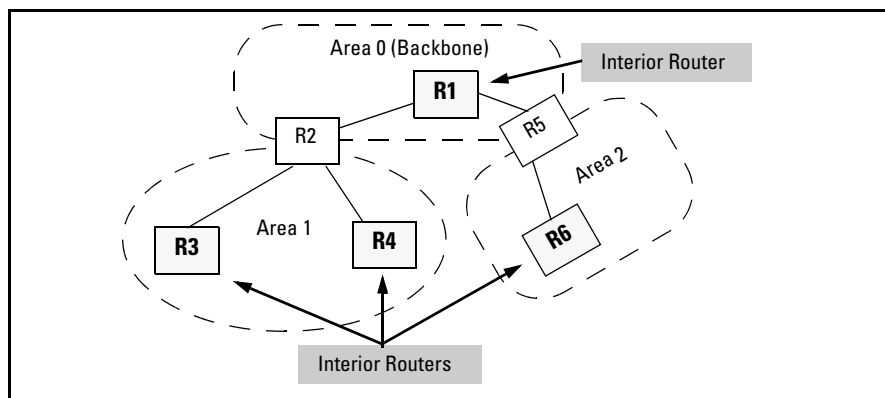


Figure 5-11. Example of Interior Routers

Area Border Routers (ABRs)

This type of OSPF router has membership in multiple areas. ABRs are used to connect the various areas in an AS to the backbone area for that AS. Multiple ABRs can be used to connect a given area to the backbone, and a given ABR can belong to multiple areas other than the backbone. An ABR maintains a separate LSDB for each area to which it belongs. (All routers within the same area have identical LSDBs.) The ABR is responsible for flooding summary LSAs between its border areas. You can reduce summary LSA flooding by configuring area ranges. An area range enables you to assign an aggregate address to a range of IP addresses. This aggregate address is advertised instead of all the individual addresses it represents. You can assign up to eight ranges in an OSPF area. In figure 5-12, below, routers R2 and R5 are Area Border Routers (ABRs) because they both have membership in more than one area.

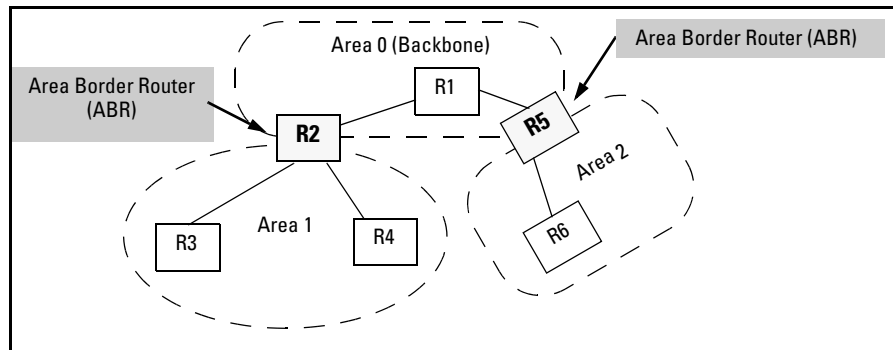


Figure 5-12. Example of Deploying ABRs To Connect Areas to the Backbone

Autonomous System Boundary Router (ASBR)

This type of OSPF router runs multiple Interior Gateway protocols and serves as a gateway to other autonomous systems operating with interior gateway protocols. The ASBR imports and translates different protocol routes into OSPF through *redistribution*. ASBRs can be used in backbone areas, normal areas, and NSSAs, but not in stub areas. For more details on redistribution and configuration examples, see “2. Enable Route Redistribution” on page 5-69.

Designated Routers

In an OSPF network having two or more routers, one router is elected to serve as the designated router (DR) and another router to act as the backup designated router (BDR). All other routers in the area forward their routing information to the DR and BDR, and the DR forwards this information to all of the routers in the network. This minimizes the amount of repetitive information that is forwarded on the network by eliminating the need for each individual router in the area to forward its routing information to all other routers in the network. If the area includes multiple networks, then each network elects its own DR and BDR.

In an OSPF network with no designated router and no backup designated router, the neighboring router with the highest priority is elected as the DR, and the router with the next highest priority is elected as the BDR. If the DR goes off-line, the BDR automatically becomes the DR, and the router with the next highest priority then becomes the new BDR. If multiple ProCurve routing switches on the same OSPF network are declaring themselves as DRs, then both priority and router ID are used to select the designated router and backup designated routers.

Priority is configurable by using the **vlan < vid > ip ospf priority < 0-255 >** command at the interface level. You can use this parameter to help bias one router as the DR. (For more on this command, refer to “Priority Per-Interface” on page 5-77.) If two neighbors share the same priority, the router with the highest router ID is designated as the DR. The router with the next highest router ID is designated as the BDR.

For example, in figure 5-13, the DR and BDR for 10.10.10.0 network in area 5 are determined as follows:

- Router A Priority: 0 Cannot become a DR or BDR.
- Router B Priority: 1 DR for the 10.10.10.0 network.
- Router C Priority: 2 BDR for the 10.10.10.0 network.
- Router D Priority: 0 Cannot become a DR or BDR.
- Router E Priority: 3 Becomes the new BDR if router B becomes unavailable and router C becomes the new DR.

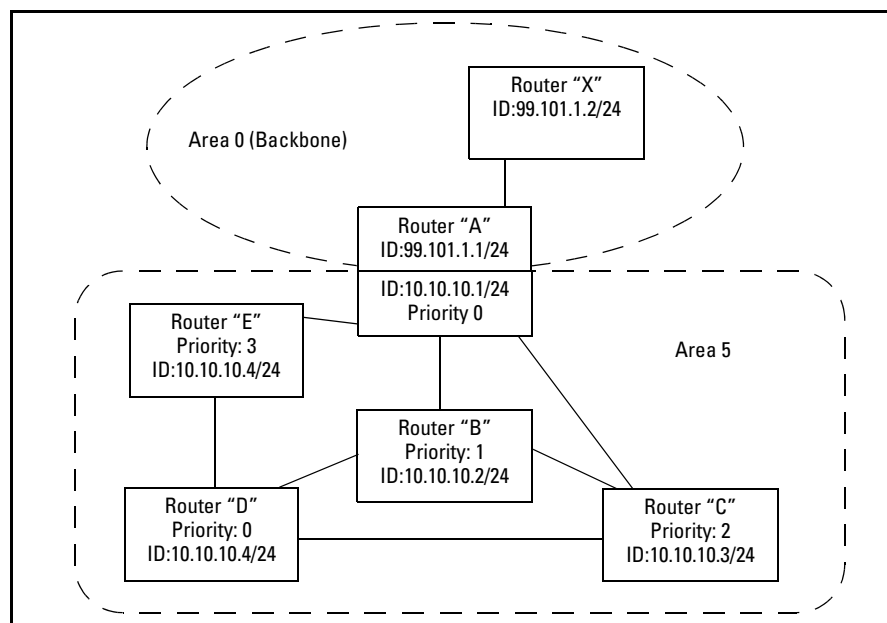


Figure 5-13. Example of Designated Routers in an OSPF Area

To learn the router priority on an interface, use the **show ip ospf interface** command and check the **Pri** setting under **OSPF interface configuration**.

Notes

By default, the router ID is typically the lowest-numbered IP address or the lowest-numbered (user-configured) loopback interface configured on the device. For more information or to change the router ID, see “Changing the Router ID” on page 5-13.

If multiple networks exist in the same OSPF area, the recommended approach is to ensure that each network uses a different router as its DR. Otherwise, if a router is a DR for more than one network, latency in the router could increase due to the increased traffic load resulting from multiple DR assignments.

When only one router on an OSPF network claims the DR role despite neighboring routers with higher priorities or router IDs, this router remains the DR. This is also true for BDRs.

The DR and BDR election process is performed when one of the following events occurs:

- an interface is in a waiting state and the wait time expires
- an interface is in a waiting state and a hello packet is received that addresses the BDR
- a change in the neighbor state occurs, such as:
 - a neighbor state transitions from 2 or higher
 - communication to a neighbor is lost
 - a neighbor declares itself to be the DR or BDR for the first time

OSPF Area Types

OSPF is built upon a hierarchy of network areas. All areas for a given OSPF domain reside in the same *Autonomous System (AS)*. An AS is defined as a number of contiguous networks, all of which share the same interior gateway routing protocol.

An AS can be divided into multiple areas. Each area represents a collection of contiguous networks and hosts, and the topology of a given area is not known by the internal routers in any other area. Areas define the boundaries to which types 1 and 2 LSAs are broadcast, which limits the amount of LSA flooding that occurs within the AS and also helps to control the size of the link-state databases (LSDBs) maintained in OSPF routers. An area is represented in OSPF by either an IP address or a number. Area types include:

- backbone
- not-so-stubby (NSSA)
- normal
- stub

All areas in an AS must connect with the backbone through one or more area border routers (ABRs). If a *normal* area is not directly connected to the backbone area, it must be configured with a *virtual link* to an ABR that is directly connected to the backbone. The remaining area types do not allow virtual link connections to the backbone area.

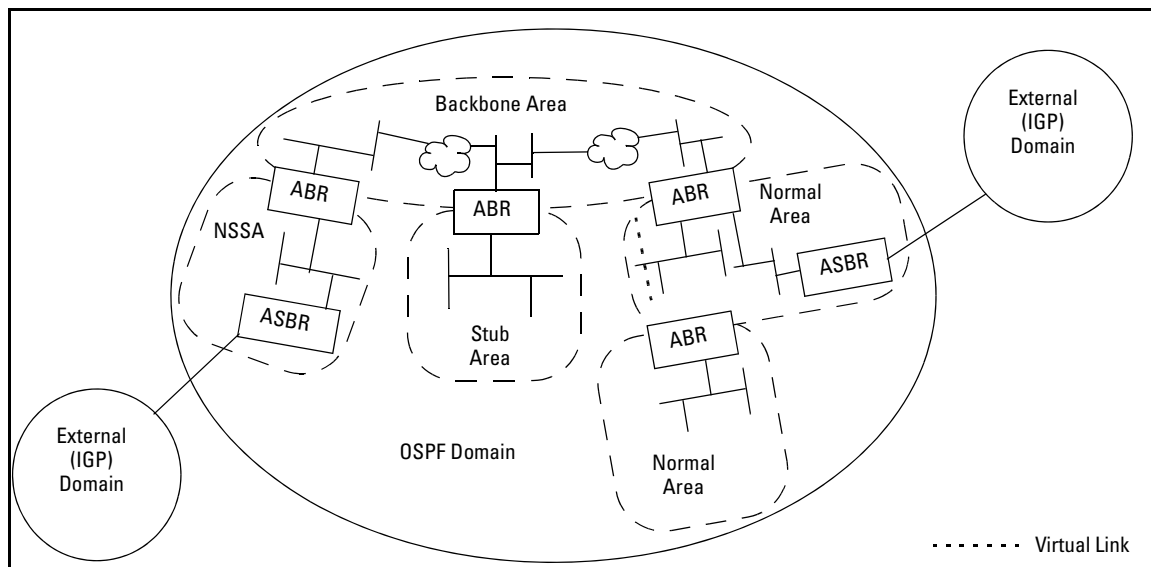


Figure 5-14. Example of an Autonomous System (AS) with Multiple Areas and External Routes

Backbone Area

Every AS must have one (and only one) backbone area (identified as area 0 or 0.0.0.0). The ABRs of all other areas in the same AS connect to the backbone area, either physically through an ABR or through a configured, virtual link. The backbone is a transit area that carries the type-3 summary LSAs, type-5 AS external link LSAs and routed traffic between non-backbone areas, as well as the type-1 and type-2 LSAs and routed traffic internal to the area. ASBRs are allowed in backbone areas.

Normal Area

This area connects to the AS backbone area through one or more ABRs (physically or through a virtual link) and supports type-3 summary LSAs and type-5 external link LSAs to and from the backbone area. ASBRs are allowed in normal areas.

Not-So-Stubby-Area (NSSA)

Beginning with software release K.12.33, this area is available and connects to the backbone area through one or more ABRs. NSSAs are intended for use where an ASBR exists in an area where you want to control the following:

- advertising the ASBR's external route paths to the backbone area
- advertising the NSSA's summary routes to the backbone area
- allowing LSAs from the backbone area to advertise in the NSSA:
 - summary routes (type-3 LSAs) from other areas
 - external routes (type-5 LSAs) from other areas as a default external route (type-7 LSAs)

In the above operation, the ASBR in the NSSA injects external routes as type 7 LSAs. (Type 5 LSAs are not allowed in an NSSA.) The ABR connecting the NSSA to the backbone converts the type 7 LSAs to type 5 LSAs and injects them into the backbone area for propagation to networks in the backbone and to any normal areas configured in the AS. The ABR also injects type-3 summary LSAs:

- from the NSSA into the backbone area
- from the backbone into the NSSA

As mentioned above, if the ABR detects type-5 external LSAs on the backbone, it injects a corresponding type-7 LSA default route (0.0.0.0/0) into the NSSA

You can also configure the NSSA ABR to do the following:

- Suppress advertising some or all of the area's summarized internal or external routes into the backbone area. (Refer to "8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone" on page 5-70.)
- Replace all type-3 summary routes and the type-7 default route with the type-3 default summary route (0.0.0.0/0).

Virtual links are not allowed for NSSAs.

Stub Area

This area connects to the AS backbone through one or more ABRs. It does not allow an internal ASBR, and does not allow external (type 5) LSAs. A stub area supports these actions:

- Advertise the area's summary routes to the backbone area.
- Advertise summary routes from other areas.
- Use the default summary (type-3) route to advertise both of the following:
 - summary routes to other areas in the AS
 - external routes to other autonomous systems

You can configure the stub area ABR to do the following:

- Suppress advertising some or all of the area's summarized internal routes into the backbone area.
- Suppress LSA traffic from other areas in the AS by replacing type-3 summary LSAs and the default external route from the backbone area with the default summary route (0.0.0.0/0).

Virtual links are not allowed for stub areas.

OSPF RFC Compliance

The OSPF features covered in this guide comply with the following:

- RFC 2328 OSPF version 2
- RFC 3101 OSPF NSSA option (s/w release K.12.*xx* and greater)
- RFC 1583 (Enabled in the default OSPF configuration. Refer to the following Note.)

Note

If all of the routers in your OSPF domain support RFC 2178, RFC 2328, or later, you should disable RFC 1583 compatibility on all routers in the domain. Refer to “3. Changing the RFC 1583 OSPF Compliance Setting” on page 5-57.

Reducing AS External LSAs and Type-3 Summary LSAs

An OSPF ASBR uses AS External LSAs to originate advertisements of a route to another routing domain, such as a RIP domain. These advertisements are

- flooded in the area in which the ASBR operates

- injected into the backbone area and then propagated to any other OSPF areas (except stub areas) within the local OSPF Autonomous System (AS). If the AS includes an NSSA, there are two additional options:
 - If the NSSA includes an ASBR, you can suppress advertising some or all of its summarized external routes into the backbone area.
 - Replace all type-3 summary LSAs and the default external route from the backbone area with the default summary route (0.0.0.0/0).

In some cases, multiple ASBRs in an AS can originate equivalent external LSAs. The LSAs are equivalent when they have the same cost, the same next hop, and the same destination. In such cases, the ProCurve switch optimizes OSPF by eliminating duplicate AS External LSAs. That is, the ASBR with the highest router ID floods the AS External LSAs for the external domain into the OSPF AS, while the other ASBRs flush the equivalent AS External LSAs from their databases. As a result, the overall volume of route advertisement traffic within the AS is reduced and the switches that flush the duplicate AS External LSAs have more memory for other OSPF data.

This enhancement implements the portion of RFC 2328 that describes AS External LSA reduction. This enhancement is enabled by default, requires no configuration, and cannot be disabled.

Algorithm for AS External LSA Reduction

The AS External LSA reduction feature behavior changes under the following conditions:

- There is one ASBR advertising (originating) a route to the external destination, but one of the following happens:
 - A second ASBR comes on-line
 - A second ASBR that is already on-line begins advertising an equivalent route to the same destination.

In either case above, the ProCurve switch with the higher router ID floods the AS External LSAs and the other ProCurve switch flushes its equivalent AS External LSAs.

- One of the ASBRs starts advertising a route that is no longer equivalent to the route the other ASBR is advertising. In this case, the ASBRs each flood AS External LSAs. Since the LSAs either no longer have the same cost or no longer have the same next-hop router, the LSAs are no longer equivalent, and the LSA reduction feature no longer applies.
- The ASBR with the higher router ID becomes unavailable or is reconfigured so that it is no longer an ASBR. In this case, the other ASBR floods the AS External LSAs.

Replacing Type-3 Summary LSAs and Type-7 Default External LSAs with a Type-3 Default Route LSA

By default, a routing switch operating as an ABR for a stub area or NSSA injects non-default, summary routes (LSA type 3) into the stub areas and NSSAs. For NSSAs, the routing switch also injects a type-7 default external route. You can further reduce LSA traffic into these areas by using **no-summary**. This command option configures the routing switch to:

- Replace type-3 summary LSA injection into a stub area or NSSA with a type-3 default summary route (0.0.0.0/0).
- Disable injection of the type-7 default external route into an NSSA.

You can enable this behavior when you first configure the stub area or NSSA, or at a later time. (For the full command to use, refer to “Configuring a Stub or NSSA Area” on page 5-61.)

The **no-summary** command does not affect intra-area advertisements, meaning the switch still accepts summary LSAs from OSPF neighbors within its area and floods them to other neighbors. The switch can form adjacencies with other routers regardless of whether summarization is enabled or disabled for areas on each switch.

When you use **no-summary**, the change takes effect immediately. If you apply the option to a previously configured area, the switch flushes all of the summary LSAs it has generated (as an ABR) from the area.

Note

This feature applies only when the switch is configured as an Area Border Router (ABR) for a stub area or NSSA. To completely prevent summary LSAs from injection into the area, use **no-summary** to disable the summary LSAs on each OSPF router that is an ABR for the area.

To implement the above operation for a stub area or NSSA, enter a command such as the following:

```
ProCurve(ospf)# area 40 stub 3 no-summary
```

Dynamic OSPF Activation and Configuration

OSPF automatically activates when enabled with **router ospf**. All configuration commands affecting OSPF (except reconfiguring the router ID) are dynamically implemented, and can be used without restarting OSPF routing. (To reconfigure the router ID, refer to “Changing the Router ID” on page 5-13.)

Note

OSPF is automatically enabled without a system reset.

General Configuration Steps for OSPF

To begin using OSPF on the routing switch, perform the steps outlined below:

1. In the global config context, use **ip routing** to enable routing (page 5-56).
2. Execute **router ospf** to place the routing switch in the **ospf** context and to enable OSPF routing (page 5-56).
3. Change the OSPF RFC 1583 compliance, if needed. (Refer to “3. Changing the RFC 1583 OSPF Compliance Setting” on page 5-57.)
4. Use **area** to assign the areas to which the routing switch will be attached (page 5-59).
5. Assign interfaces to the configured areas per-VLAN or per-subnet by moving to each VLAN context and using one of the following commands:
 - **ip ospf area < ospf-area-id >** assigns all interfaces in the VLAN to the same area. Use this option when there is only one IP address configured on the VLAN or you want all subnets in the VLAN to belong to the same OSPF area.
 - **ip ospf < ip-address > area < ospf-area-id >** assigns an individual subnet to the specified area.

(Refer to page 5-64.)

6. Optional: Assign loopback interfaces to OSPF areas by using the **ip ospf area** command at the loopback interface configuration level.
(Refer to page 5-65.)
7. Optional: On each routing switch used as an ASBR in your OSPF domain, configure redistribution to enable importing the routes you want to make available in the domain.
 - i. On an ASBR in a backbone, normal, or NSSA area where you want to import external routes, configure redistribution filters to define the external routes you *do not* want imported.
 - ii. Enable redistribution.

Refer to “7. Optional: Configure for External Route Redistribution in an OSPF Domain” on page 5-67.

8. Optional: Configure ranges on ABRs to reduce inter-area route advertising.
9. Optional: Use Administrative Distance to influence route choices.
10. Optional: Change OSPF trap generation.
11. Optional: Reconfigure default parameters in the interface context, if needed. Includes **cost**, **dead-interval**, **hello-interval**, **priority**, and others.

12. Optional: Configure OSPF interface authentication.
13. Configure virtual links for any areas not directly connected to the backbone.

Configuration Rules

- If the switch is to operate as an ASBR, you must enable redistribution (step 7 on page 5-53). When you do that, ASBR capability is automatically enabled. For this reason, you should first configure redistribution filters on the ASBR. Otherwise, all possible external routes will be allowed to flood the domain. (Refer to “7. Optional: Configure for External Route Redistribution in an OSPF Domain” on page 5-67.)
- Each VLAN interface on which you want OSPF to run must be assigned to one of the defined areas. When a VLAN interface is assigned to an area, the IP address is automatically included in the assignment. To include additional addresses, you must enable OSPF on them separately, or use the “all” option in the assignment.

OSPF Global and Interface Settings

When first enabling OSPF, you may want to consider configuring ranges and restricting redistribution (if an ASBR is used) to avoid unwanted advertisements of external routes. You may also want to enable the OSPF trap and authentication features to enhance troubleshooting and security. However, it is generally recommended that the remaining parameters with non-null default settings be left as-is until you have the opportunity to assess OSPF operation and determine whether any adjustments to non-default settings is warranted.

The following tables list the global and per-interface commands used with OSPF. For information on when to use these commands, refer to “General Configuration Steps for OSPF” on page 5-53. For detailed information on each command, refer to the page listed for each command.

Table 5-6. OSPF Default Global Settings

Parameter	Default	Page
area < area-#> virtual-link < ip-addr >	None	
default-metric	10	
distance < external inter-area intra-area >	110	
range	All IP Addresses	
redistribute	Disabled	

Parameter	Default	Page
restrict	Disabled	
rfc-1583-compatibility	Enabled	
metric-type	type2	
trap < ospf-trap >	Enabled	

Note

Set global level parameters in the **ospf** context of the CLI. To access this context level, ensure that routing is enabled, then execute **router ospf** at the global CONFIG level. For example:

```
ProCurve (config)# router ospf
ProCurve (ospf)#
```

Table 5-7. OSPF Default Interface Settings

Parameter	Default	Page
all	n/a	
area	None	
authentication-key	None	
cost	1	
dead-interval	40 seconds	
hello-interval	10 seconds	
IP-ADDR	None	
md5-auth-key-chain	None	
priority	1	
retransmit-interval	5 seconds	
transit-delay	1 second	

Note

Use the VLAN interface context to set interface level OSPF parameters for the desired VLAN. To access this context level, use **vlan < vid >** either to move to the VLAN context level or to specify that context from the global config level. For example, both of the following two command sets achieve the same result:

```
ProCurve(config)# vlan 20
ProCurve(vlan-20)# cost 15

ProCurve(config)# vlan 20 cost 15
```

Configuring OSPF on the Routing Switch

1. Enable IP Routing

Syntax: [no] ip routing

Executed at the global configuration level to enable IP routing on the routing switch.

Default: Disabled

*The **no** form of the command disables IP routing. (Global OSPF and RIP routing must be disabled before you disable IP routing.)*

```
ProCurve(config)# ip routing
```

2. Enable Global OSPF Routing

Syntax: [no] router ospf

Executed at the global configuration level to enable OSPF on the routing switch and to enter the OSPF router context. This enables you to proceed with assigning OSPF areas, including ABR and ASBR configuration, and to modify OSPF global parameter settings as needed. Global IP routing must be enabled before executing this command.

Default: Disabled

*The **no** form of the command disables OSPF routing.*

Note: *If you disable OSPF, the switch retains all the configuration information for the disabled protocol in flash memory. If you subsequently restart OSPF, the existing configuration will be applied.*

For example:

```
ProCurve(config)#router ospf  
ProCurve(ospf)#
```


3. Changing the RFC 1583 OSPF Compliance Setting

In OSPF domains supporting multiple external routes from different areas to the same external destination, multiple AS-external-LSAs advertising the same destination are likely to occur. This can cause routing loops and the network problems that loops typically generate. On the routing switches covered by this guide, if RFC 1583 compatibility is disabled, the preference rules affecting external routes are those stated in RFC-2328, which minimize the possibility of routing loops when AS-external-LSAs for the same destination originate from ASBRs in different areas. However, because all routers in an OSPF domain must support the same routing-loop prevention measures, if the domain includes any routers that support only RFC 1583 preference rules, then all routers in the domain must be configured to support RFC 1583.

Note

The routing switch is configured, by default, to be compliant with the RFC 1583 OSPF V2 specification. (Use **show ip ospf general** to view the current RFC 1583 configuration setting.)

All routes in an AS should be configured with the same compliance setting for preference rules affecting external routes. Thus, if any routers in an OSPF domain support only RFC 1583, then all routers must be configured with 1583 compatibility. In the default OSPF configuration, RFC 1583 support is enabled for the routing switches covered by this guide.

If all routers in the domain support RFC 2178 or RFC 2328, then you should disable RFC 1583 compatibility on all of the routers, since conformance to these later RFCs provides more robust protection against routing loops on external routes.

Syntax: [no] rfc1583-compatibility

Executed at the global configuration level to toggle routing switch operation compliance between RFC 1583 and RFC 2328.

rfc1583-compatibility: *Configures the routing switch for external route preference rules compliant with RFC 1583.*

no rfc1583-compatibility: *Configures the routing switch for external route preference rules compliant with RFC 2328.*

Default: Compliance enabled

For example, to disable RFC 1583 compatibility on a routing switch in an OSPF domain where RFC 2178 and RFC 2328 are universally supported:

```
ProCurve(config)# router ospf
ProCurve(ospf)# no rfc1583-compatibility
```

```
ProCurve(config)# router ospf
ProCurve(ospf)# no rfc1583-compatibility
ProCurve_8212(ospf)# show ip ospf general
```

OSPF General Status

OSPF protocol	: enabled
Router ID	: 10.10.51.1
RFC 1583 compatibility	: non-compatible
Intra-area distance	: 110
Inter-area distance	: 110
AS-external distance	: 110
Default import metric	: 10
Default import metric type	: external type 2
Area Border	: no
AS Border	: yes
External LSA Count	: 9
External LSA Checksum Sum	: 408218
Originate New LSA Count	: 24814
Receive New LSA Count	: 14889

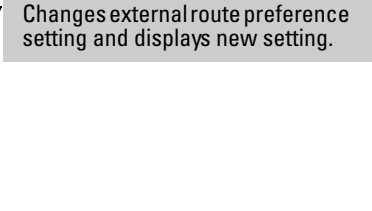


Figure 5-15. Example of Changing External Route Preference Compatibility from RFC 1583 to RFC 2328

4. Assign the Routing Switch to OSPF Areas

After you globally enable OSPF on the routing switch (in the previous step), use this command to assign one or more OSPF areas within your autonomous system (AS). A routing switch can belong to one area or to multiple areas. (Participation in a given, assigned area requires configuring one or more VLANs or subnets and assigning each to the desired area. Refer to page 5-64.)

- If you want the VLANs and any subnets configured on the routing switch to all reside in the same area, then you need to configure only that one area. (In this case, the routing switch would operate as an internal router for the area.)
- If you want to put different VLANs or subnets on the routing switch into different areas, then you need to re-execute this command for each area. (In this case, the routing switch will operate as an ABR for each of the configured areas.)

Note

Each ABR must either be directly connected to the backbone area (0) or be configured with a virtual link to the backbone area through another ABR that is directly connected to the backbone area. For information on this step, refer to “13. Configuring an ABR To Use a Virtual Link to the Backbone” on page 5-82.

Configuring an OSPF Backbone or Normal Area.

Syntax: area < ospf-area-id | backbone > [normal]
no area < ospf-area-id | backbone >

*After using **router ospf** to globally enable OSPF and enter the global OSPF context, execute this command to assign the routing switch to a backbone or other normal area.*

*The **no** form of the command removes the routing switch from the specified area.*

Default: No areas. Range: 1-16 areas (of all types)

*< ospf-area-id >: Specifies a normal area to which you are assigning the routing switch. You can assign the routing switch to one or more areas, depending on the area in which you want each configured VLAN or subnet to reside. You can enter area IDs in either whole number or dotted decimal format. (The routing switch automatically converts whole numbers to the dotted decimal format.) For example, if you enter an area-ID of **1**, it appears in the switch's configuration as **0.0.0.1** and an area-ID of 256 appears in the switch configuration as **0.0.1.0**. An area ID can be a value selected to match the IP address of a VLAN belonging to the area, or a value corresponding to a numbering system you devise for the areas in a given AS. Entering an area ID of **0** or **0.0.0.0** automatically joins the routing switch to the Backbone area. The maximum area ID value is 255.255.255.254 (4,294,967,294).*

backbone: *Assigns the routing switch to the backbone area and automatically assigns an area ID of **0.0.0.0** and an area type of **normal**. Using **0** or **0.0.0.0** with the above **ospf-area-id** option achieves the same result. The backbone area is automatically configured as a "normal" area type.*

For example, to configure a backbone and a normal area with an ID of "1" (0.0.0.1) on a routing switch:

```
ProCurve(ospf)# area backbone  
ProCurve(ospf)# area 1
```

Configuring a Stub or NSSA Area.

Syntax: area < ospf-area-id > stub < 0-16777215 > [no-summary]
area < ospf-area-id > nssa < 0-16777215 > [no-summary]
[metric-type < type1 | type2 >]
no area < ospf-area-id >

*After using **router ospf** to globally enable OSPF and enter the global OSPF context, execute this command to assign the routing switch to a stub area or NSSA. (Does not apply to backbone and normal OSPF area ABRs.)*

*The **no** form of the command removes the routing switch from the specified area.*

Default: No areas. Range: 1-16 areas (of all types)

< ospf-area-id >: Same area ID as on page 5-60 except you cannot assign a backbone area number (**0** or **0.0.0.0**) to a stub or NSSA area.

< stub | nssa > *Designates the area identified by < ospf-area-id > as a stub area or NSSA.*

< 0-16777215 >: *If the routing switch is used as an ABR for the designated area, assigns the cost of the default route (to the backbone) that is injected into the area.*

Notes: *If the routing switch is not an ABR for the stub area or NSSA, the above cost setting is still required by the CLI, but is not used.*

In the default configuration, a routing switch acting as an ABR for a stub area or NSSA injects type-3 summary routes into the area. For an NSSA, the routing switch also injects a type-7 default route into the area.

[no-summary]: Where the routing switch is an ABR for a stub area or an NSSA, this option reduces the amount of LSA traffic entering the area from the backbone by replacing the injection of type-3 summary routes with injection of a type-3 default summary route. For NSSAs, this command also disables injection of the type-7 default external route from the backbone into the area (included in the **metric-type** operation described below).

(Default: Disabled)

For more on this topic, refer to “Not-So-Stubby-Area (NSSA)” on page 5-49, “Stub Area” on page 5-50, and “Replacing Type-3 Summary LSAs and Type-7 Default External LSAs with a Type-3 Default Route LSA” on page 5-52.

[metric-type < type1 | type2 >]: Used in NSSA ABRs only. Enables injection of the type-7 default external route and type-3 summary routes into the area instead of a type 3 default route. Also specifies the type of internal cost metric to include in type-7 LSAs advertised for redistribution of external routes in the NSSA. (The redistribution—or external—cost metric is a global setting on the routing switch set by the **default-metric** command.) The **metric-type** command specifies whether to include the redistribution cost in the cost metric calculation for a type-7 default LSA injected into the area.

type1: Calculate external route cost for a type-7 default LSA as the sum of (1) the external route cost assigned by the ASBR plus (2) the internal cost from the router with traffic for the external route to the ASBR advertising the route.

type2: Calculate external route cost for a type-7 default LSA as being only the cost from the router with traffic for the external route to the ASBR advertising the route.

Using the **area < ospf-area-id > nssa < 0-16777215 >** without entering either **no-summary** or **metric-type** resets the routing switch to the state where injection of type-3 summary routes and the type-7 default external routes is enabled with **metric-type** set to **type2**.

(Default: Enabled with **metric-type type2**.)

Note: Different routers in the NSSA can be configured with different **metric-type** values.

The following examples of configuring a stub area and an NSSA on a routing switch use an (arbitrary) cost of “10”.

```
ProCurve(ospf)# area 2 stub 10
ProCurve(ospf)# area 3 nssa 10
ProCurve(ospf)# area 4 nssa 10 no-summary
ProCurve(ospf)# area 5 nssa 10 metric-type type1
```

Assigns a stub area with a cost of 10.

Assigns an NSSA with a cost of 10 and, by default, uses a Type2 default cost metric for Type-7 (external) route LSAs received from the backbone.

Assigns an NSSA with a cost of 10, blocks injection of type-3 summary routes, and starts injection of type-3 default routes from the backbone.

Sets the cost metric type for type-7 (default) LSAs injected into the NSSA.

Figure 5-16. Examples of Creating Stub Area and NSSA Assignments

5. Assign VLANs and/or Subnets to Each Area

After you define an OSPF area (page 5-59), you can assign one or more VLANs and/or subnets to it. When a VLAN is assigned to an area, all currently configured IP addresses in that VLAN are automatically included in the assignment unless you enter a specific IP address.

Note

All static VLANs configured on a routing switch configured for OSPF must be assigned to one of the defined areas in the AS.

Syntax: `vlan < vid ># ip ospf [ip-addr | all] area < ospf-area-id >`

Executed in a specific VLAN context to assign the VLAN or individual subnets in the VLAN to the specified area. Requires that the area is already configured on the routing switch (page 5-59). When executed without specifying an IP address or using the **all** keyword, this command assigns all configured networks in the VLAN to the specified OSPF area.

vlan < vid >: Defines the VLAN context for executing the area assignment.

[ip-addr]: Defines a specific subnet on the VLAN to assign to a configured OSPF area.

[all]: Assigns all subnets configured on the VLAN to a configured OSPF area.

area < ospf-area-id >: Identifies the OSPF area to which the VLAN or selected subnet should be assigned.

Notes: *If you add a new subnet IP address to a VLAN after assigning the VLAN to an OSPF area, you must also assign the new subnet to an area. If all subnets in the VLAN should be assigned to the same area, just execute **ip ospf area < ospf-area-id >**. But if different subnets belong in different areas, you must explicitly assign the new subnet to the desired area. Also, to assign a VLAN to an OSPF area, the VLAN must be configured with at least one IP address. Otherwise, executing this command results in the following CLI message:*

OSPF can not be configured on this VLAN.

Example: To assign VLAN 8 on a routing switch to area 3 and include *all* IP addresses configured in the VLAN, enter the following commands:

```
ProCurve(ospf)# vlan 8
ProCurve(vlan-8)# ip ospf area 3
```


Example. Suppose that a system operator wants to assign the three subnets configured in VLAN 10 as shown below:

- 10.10.10.1 to OSPF area 5
- 10.10.11.1 to OSPF area 5
- 10.10.12.1 to OSPF area 6

The operator could use the following commands to configure the above assignments:

```
ProCurve(ospf)# vlan 10
ProCurve(vlan-10)# ip ospf 10.10.10.1 area 5
ProCurve(vlan-10)# ip ospf 10.10.11.1 area 5
ProCurve(vlan-10)# ip ospf 10.10.12.1 area 6
```

6. Optional: Assigning Loopback Addresses to an Area

After you define the OSPF areas to which the switch belongs, you can assign a user-defined loopback address to an OSPF area. A loopback interface is a virtual interface configured with an IP address and is always reachable as long as at least one of the IP interfaces on the switch is operational. Because the loopback interface is always up, you ensure that the switch's router ID remains constant and that an OSPF network is protected from changes caused by downed interfaces.

For more information about how to configure a loopback interface, refer to “Configuring a Loopback Interface” in the chapter titled, “Configuring IP Addressing”, in the *Management and Configuration Guide* for your routing switch.

Syntax: interface loopback <0-7> ip ospf <lo-ip-address> area <ospf-area-id>

Executed in a specific loopback context to assign a loopback interface to the specified OSPF area. Requires that the specified loopback interface is already configured with an IP address on the switch.

loopback interface <0-7>: Defines the loopback context for executing the area assignment.

ip ospf <lo-ip-address>: Specifies the loopback interface by its IP address to assign to a configured OSPF area.

area <ospf-area-id>: Identifies the OSPF area to which the loopback interface is assigned. You can enter a value for the OSPF area in the format of an IP address or a number in the range 0 to 4,294,967,295.

Example: To assign user-defined loopback interface 3 on the switch to area 192.5.0.0 and include the loopback IP address 172.16.112.2 in the OSPF broadcast area, enter the following commands:

```
ProCurve(config)# interface loopback 3
ProCurve(lo-3)# ip ospf 172.16.112.2 area 192.5.0.0
```

Syntax: interface loopback <0-7># ip ospf < lo-ip-address > cost < number >

Executed in a specific loopback context to modify the cost used to advertise the loopback address (and subnet) to the area border router (ABR). Requires that the specified loopback interface is already configured with an IP address on the switch.

loopback interface <0-7>: Defines the loopback context for executing the cost assignment.

ip ospf < lo-ip-address >: Specifies the loopback interface by its IP address.

cost < number >: Specifies a number that represents the administrative metric associated with the loopback interface. Valid values are from 1 to 65535. Default: 1.

Example: To configure a cost of 10 for advertising the IP address 172.16.112.2 configured for loopback interface 3 in an OSPF area 192.5.0.0, enter the following commands:

```
ProCurve(config)# interface loopback 3
ProCurve(lo-3)# ip ospf 172.16.112.2 area 192.5.0.0
ProCurve(lo-3)# ip ospf 172.16.112.2 cost 10
```

OSPF Redistribution of Loopback Addresses: When you assign a loopback address to an OSPF area, the route redistribution of the loopback address is limited to the specified area.

When route redistribution is enabled:

- The switch advertises a loopback IP address that is not assigned to an OSPF area as an OSPF *external* route to its OSPF neighbors, and handles it as a connected route.
- The switch advertises a loopback address that is assigned to an OSPF area as an OSPF *internal* route.

To enable redistribution of loopback IP addresses in OSPF, enter the **redistribution connected** command as described in “2. Enable Route Redistribution” on page 5-69.

Example: In the following configuration, the loopback IP address 13.3.4.5 of loopback 2 is advertised only in OSPF area 0.0.0.111. The IP addresses 14.2.3.4 and 15.2.3.4 of loopback 1 are advertised in all OSPF areas.

```
ProCurve(config)# interface loopback 1
ProCurve(lo-1)# ip address 14.2.3.4
ProCurve(lo-1)# ip address 15.2.3.4
ProCurve(lo-1)# exit
ProCurve(config)# interface loopback 2
ProCurve(lo-2)# ip address 13.3.4.5
ProCurve(lo-2)# ip ospf 15.2.3.4 area 0.0.0.111
ProCurve(lo-2)# exit
```

Assigns the IP address of loopback interface 2 to OSPF area 111.

Figure 5-17. Examples of Assigning Loopback IP Addresses to OSPF Areas

To verify the OSPF redistribution of loopback interfaces, enter the **show ip route** command from any context level to display IP route table entries.

Note that in the following example, a loopback address assigned to an area is displayed as an **ospf intra-area** (internal) route to its neighbor; a loopback address not assigned to a specific area is displayed as an **ospf external** route:

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

IP Route Entries						
Destination	Gateway	VLAN	Type	Sub-Type	Metric	Dist
-----	-----	----	----	-----	-----	----
20.0.15.1/32	25.0.67.131	25	ospf	external2	10	110
20.0.16.2/32	25.0.67.131	25	ospf	intra-area	2	110

Figure 5-18. Example of Verifying OSPF Redistribution of Loopback Interfaces

7. Optional: Configure for External Route Redistribution in an OSPF Domain

Configuring route redistribution for OSPF establishes the routing switch as an ASBR (residing in a backbone, normal, or NSSA) for importing and translating different protocol routes from other IGP domains into an OSPF domain. The switches covered by this guide support redistribution for static routes, RIP routes, and directly connected routes from RIP domains into OSPF

domains. When you configure redistribution for OSPF, you can specify that static, connected, or RIP routes external to the OSPF domain are imported as OSPF routes. (Likewise, RIP redistribution supports the import of static, connected, and OSPF routes into RIP routes.) The steps for configuring external route redistribution to support ASBR operation include the following:

1. Configure redistribution filters to exclude external routes that you do not want redistributed in your OSPF domain.
2. Enable route redistribution.
3. Modify the default metric for redistribution (optional).
4. Modify the redistribution metric type (optional).
5. Change the administrative distance setting (optional).

Note

Do not enable redistribution until you have used **restrict** to configure the redistribution filters. Otherwise, your network might get overloaded with routes that you did not intend to redistribute.

1. Configure Redistribution Filters.

Syntax: `router ospf restrict < ip-addr/mask-length >`

This command prevents distribution of the specified range of external routes through an ASBR from sources external to the OSPF domain.

Default: Allow all supported, external route sources.

Note: *Use this command to block unwanted, external routes before enabling route redistribution on the ASBR.*

Example: To configure a routing switch operating as an ASBR to filter out redistribution of static, connected, or RIP routes on network 10.0.0.0, enter the following commands:

```
ProCurve(config)# router ospf restrict 10.0.0.0/8
```

Note

In the default configuration, redistribution is permitted for all routes from supported sources.

2. Enable Route Redistribution. This step enables ASBR operation on a routing switch, and must be executed on each routing switch connected to external routes you want to redistribute in your OSPF domain.

Note

Do not enable redistribution until you have configured the redistribution “restrict” filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

Syntax: [no] router ospf redistribute < connected | static | rip >

Executed on an ASBR to globally enable redistribution of the specified route type to the OSPF domain through the area in which the ASBR resides.

static: *Redistribute from manually configured routes.*

connected: *Redistribute from locally connected network(s).*

rip: *Redistribute from RIP routes.*

*The **no** form of the command disables redistribution for the specified route type.*

For example, to enable redistribution of all supported external route types through a given ASBR, execute the following commands.

```
ProCurve(config)# router ospf redistribution connected
ProCurve(config)# router ospf redistribution static
ProCurve(config)# router ospf redistribution rip
```

3. Modify the Default Metric for Redistribution. The default metric is a global parameter that specifies the cost applied to all OSPF routes by default

Syntax: router ospf default-metric < 0-16777215 >

Globally assigns the cost metric to apply to all external routes redistributed by the ASBR. By using different cost metrics for different ASBRs, you can prioritize the ASBRs in your AS.

Default: 10; Range: 0-16777215

Example: To assign a default metric of 4 to all routes imported into OSPF on an ASBR, enter the following commands:

```
ProCurve()#
ProCurve(config)# router ospf default-metric 4
```

4. Modifying the Redistribution Metric Type. The redistribution metric type is used by default for all routes imported into OSPF. Type 1 metrics are the same “units” as internal OSPF metrics and can be compared directly. Type 2 metrics are not directly comparable, and are treated as larger than the largest internal OSPF metric.

Syntax: `router ospf metric-type < type1 | type2 >`

Globally reconfigures the redistribution metric type on an ASBR.

type1: Specifies the OSPF metric plus the external metric for an external route.

type2: Specifies the external metric for an external route.

Default: type2

For example, to change from the default setting on an ASBR to type 1, enter the following command:

```
ProCurve(config)# router ospf metric-type type1
```

8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone

Configuring ranges does the following to reduce inter-area advertising:

- **Summarizing Routes:** Enable a routing switch operating as an ABR to use a specific IP address and mask to summarize a range of IP addresses into a single route advertisement for injection into the backbone. This results in only one address being advertised to the network instead of all the addresses within that range. This reduces LSA traffic and the resources needed to maintain routing tables.
- **Blocking Routes:** Prevent an ABR from advertising specific networks or subnets to the backbone area.

Each OSPF area supports up to 8 range configurations.

Syntax: `area < ospf-area-id > range < ip-addr/mask-length > [no-advertise]
[type < summary | nssa >]`

Use this command on a routing switch intended to operate as an ABR for the specified area to do either of the following:

- *Simultaneously create the area and corresponding range setting for routes to summarize or block.*
- *For an existing area, specify a range setting for routes to summarize or block.*

<ospf-area-id>: *Same area ID as on page 5-60 except you cannot use a backbone area number (0 or 0.0.0.0) for a stub area or NSSA.*

range < ip-addr/mask-length >: *Defines the range of route advertisements to either summarize for injection into the backbone area or to prevent from being injected into the backbone area.*

*The **ip-addr** value specifies the IP address portion of the range, and **mask-length** specifies the leftmost significant bits in the address. The ABR for the specified area compares the IP address of each outbound route advertisement with the address and significant bits in the mask to determine which routes to select for either summarizing or blocking. For example, a range of 10.10.32.1/14 specifies all routes in the range of 10.10.32.1 - 10.10.35.254.*

[no-advertise]: *Use this keyword only if you want to configure the ABR to prevent advertisement to the backbone of a specified range of routes. (This has the effect of “hiding” the specified range from the backbone area.) If you do not use this option, the ABR advertises the specified range of routes according to the **type < summary | nssa >** selection described below.*

[type < summary | nssa >]: *Configures the type of route summaries to advertise or block. If **type** is not used in the command, then the ABR defaults this setting to **summary**.*

summary: *Specifies internal routes in the configured range of route advertisements. If **no-advertise** (above) is used in the command, then the ABR prevents the selected internal routes from being summarized in a type-3 LSA and advertised to the backbone. If **no-advertise** is not used in the command, then the selected routes are summarized to the backbone in a type-3 LSA.*

nssa: *Specifies external routes (type-7 LSAs) in the configured range of route advertisements. If **no-advertise** (above) is used in the command, then the ABR prevents the selected external routes from being summarized in a type-5 LSA and advertised to the backbone. (Configure this option where an ABR for an NSSA advertises external routes that you do not want propagated to the backbone.) If **no-advertise** is not used in the command, then the selected routes learned from type-7 LSAs in the area are summarized to the backbone in a type-5 LSA.*

Examples of an ABR Allowing or Blocking Advertisement of a Range of Internal Routes Available in an Area. Both of the following commands define the same range of internal routes in area 30 to summarize for injection into the backbone area. (In this example, area 30 can be a normal or stub area, or an NSSA.)

```
ProCurve(ospf)# area 30 range 10.0.0.0/8  
ProCurve(ospf)# area 30 range 10.0.0.0/8 type summary
```

Figure 5-19. Example of Defining a Range of Internal Routes To Advertise to the Backbone

For the same range of routes, you can use either of the following commands to block injection of a range of summary routes (type-3 LSAs) from area 30 into the backbone.

```
ProCurve(config)# area 30 range 10.0.0.0/8 type no-advertise  
ProCurve(config)# area 30 range 10.0.0.0/8 type no-advertise summary
```

Figure 5-20. Example of Defining a Range of Internal Routes To Block from Advertising to the Backbone

Example of Allowing or Blocking a Range of External Routes Available Through an ASBR in an NSSA. This example applies only to external routes that can be advertised from an NSSA to the backbone.


```
ProCurve(config)# area 7 range 192.51.0.0/16 type nssa
```

Defines the range of external routes in the Area 7 NSSA to advertise to the backbone.

```
ProCurve(config)# area 7 range 192.51.0.0/16 no-advertise type nssa
```

Defines the range of external routes in the Area 7 NSSA to block from advertising to the backbone.

Figure 5-21. Example of Allowing or Blocking a Range of External Route Advertisements to the Backbone

9. Optional: Influence Route Choices by Changing the Administrative Distance Default

The administrative distance value can be left in its default configuration setting unless a change is needed to improve OSPF performance for a specific network configuration.

The switch can learn about networks from various protocols, including RIP, and OSPF. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned. For the switches, covered in this guide the administrative distance for OSPF routes is set at 110 for all route types (external, inter-area, and intra-area).

The switch selects one route over another based on the source of the route information. To do so, the switch can use the administrative distances assigned to the sources to influence route choices. You can change the distance settings in the OSPF global context to enable preference of one route type over another.

Syntax: distance < external | inter-area | intra-area > < 1 - 255 >

Used in the OSPF configuration context to globally reconfigure the administrative distance priority for the specified route type. 1 is the highest priority; 255 is the lowest priority.

external < 1 - 255 >: Changes the administrative distance for routes between the OSPF domain and other EGP domains.

inter-area < 1 - 255 >: Changes the administrative distance for routes between areas within the same OSPF domain.

intra-area < 1 - 255 >: Changes the administrative distance for routes within OSPF areas.

Default: 110; Range: 1 - 255

10: Optional: Change OSPF Trap Generation Choices

OSPF traps (defined by RFC 1850) are supported on the routing switches covered by this guide. OSPF trap generation is disabled by default, but you can use the following command to enable generation of any or all of the supported OSPF traps.

Syntax: [no] trap < trap-name | all >

*Used in the OSPF configuration context to enable or disable OSPF traps. The **no** form disables the specified trap.*

Default: All OSPF traps disabled.

all: *Enables or disables all OSPF traps available on the routing switch.*

trap-name: *Specifies a trap from table 5-8 to enable or disable.*

Table 5-8 summarizes OSPF traps supported on the switches covered in this guide, and their associated MIB objects from RFC 1850:

Table 5-8. OSPF Traps and Associated MIB Objects

OSPF Trap Name	MIB Object
interface-authentication-failure	ospflfAuthFailure
interface-config-error	ospflfConfigError
interface-receive-bad-packet	ospflfrxBadPacket
interface-retransmit-packet	ospfTxRetransmit
interface-state-change	
neighbor-state-change	ospfNbrStateChange
originate-lsa	ospfOriginateLsa
originate-maxage-lsa	ospfMaxAgeLsa
virtual-interface-authentication-failure	
virtual-interface-config-error	ospfVirtIfConfigError
virtual-interface-state-change	ospfVirtIfStateChange
virtual-neighbor-state-change	ospfVirtNbrStateChange
virtual-interface-receive-bad-packet	ospfVirtIfRxBad Packet
virtual-interface-retransmit-packet	ospfVirtIfTxRetransmit

For example, if you wanted to monitor the neighbor-state-change and interface-receive-bad-packet traps, you would use the following commands to configure the routing switch to enable the desired trap. The **show** command verifies the resulting OSPF trap configuration.

```
ProCurve(ospf)# trap neighbor-state-change
ProCurve(ospf)# trap interface-receive-bad-packet
ProCurve(ospf)# show ip ospf traps

OSPF Traps Enabled
=====

Neighbor State Change
Interface Receive Bad Packet
```

Figure 5-22. Example of Enabling OSPF Traps

11. Optional: Adjust Performance by Changing the VLAN or Subnet Interface Settings

The following OSPF interface parameters are automatically set to their default values. No change to the defaults is usually required unless needed for specific network configurations.

Parameter	Default	Page
cost	1	
dead-interval	40 seconds	
hello-interval	10 seconds	
priority	1	
retransmit-interval	5 seconds	
transit-delay	1 second	

A setting described in this section can be configured with the same value across all subnets in a VLAN or be configured on a per-interface basis with different values.

Note

Most of the parameters in this section also apply to virtual link configurations. However, when used on a virtual link configuration, the OSPF context requirement is different and the parameters are applied only to the interfaces included in the virtual link. Refer to “Optional: Adjust Virtual Link Performance by Changing the Interface Settings” on page 5-85.

Cost Per-Interface.

Syntax: ip ospf [ip-address | all] cost < 1 - 65535 >

Used in the VLAN context to indicate the overhead required to send a packet across an interface. You can modify the cost to differentiate between 100 Mbps and 1000 Mbps (1 Gbps) links. Allows different costs for different subnets in the VLAN.

ip ospf cost < 1 - 65535 >: Assigns the specified cost to all networks configured on the VLAN.

ip ospf < ip-address > cost < 1 - 65535 >: Assigns the specified cost to the specified subnet on the VLAN.

ip ospf all cost < 1 - 65535 >: Assigns the specified cost to all networks configured on the VLAN. (Operates the same as the **ip ospf cost** option, above.)

Default: 1; Range 1 - 65535

Dead Interval Per-Interface.

Syntax: ip ospf [ip-address | all] dead-interval < 1 - 65535 >

Used in the VLAN context to indicate the number of seconds that a neighbor router waits for a hello packet from the specified interface before declaring the interface “down”. Allows different settings for different subnet interfaces in the VLAN.

ip ospf dead-interval < 1 - 65535 >: Assigns the specified dead interval to all networks configured on the VLAN.

ip ospf < ip-address > dead-interval < 1 - 65535 >: Assigns the specified dead interval to the specified subnet on the VLAN.

ip ospf all dead-interval < 1 - 65535 >: Assigns the specified dead interval to all networks configured on the VLAN. (Operates the same as the **ip ospf dead-interval** option, above.)

Default: 40 seconds; Range: 1 - 65535 seconds.

Hello Interval Per Interface.

Syntax: ip ospf [*ip-address* | all] hello-interval < 1 - 65535 >

Used in the VLAN context to indicate the length of time between the transmission of hello packets from the routing switch to adjacent neighbors. The value can be from 1 – 65535 seconds. The default is 10 seconds.. Allows different settings for different subnet interfaces in the VLAN.

ip ospf hello-interval < 1 - 65535 >: *Assigns the specified Hello interval to all networks configured on the VLAN.*

ip ospf < ip-address > hello-interval < 1 - 65535 >: *Assigns the specified Hello interval to the specified subnet on the VLAN.*

ip ospf all hello-interval < 1 - 65535 >: *Assigns the specified Hello interval to all networks configured on the VLAN. (Operates the same as the ip ospf hello-interval option, above.)*

Default: 10 seconds; Range: 1 - 65535 seconds.

Priority Per-Interface.

Syntax: ip ospf [*ip-address* | all] priority < 1 - 255 >

Used in the VLAN context to enable changing the priority of an OSPF router. The priority is used when selecting the designated router (DR) and backup designated routers (BDRs). The value can be from 0 – 255 (with 255 as the highest priority). The default is 1. If you set the priority to 0, the routing switch does not participate in DR and BDR election. Allows different settings for different subnet interfaces in the VLAN.

ip ospf priority < 1 - 255 >: *Assigns the specified priority to all networks configured on the VLAN.*

ip ospf < ip-address > priority < 1 - 255 >: *Assigns the specified priority to the specified subnet on the VLAN.*

ip ospf all priority < 1 - 255 >: *Assigns the specified priority to all networks configured on the VLAN. (Operates the same as the ip ospf priority option, above.)*

Default: 1; Range: 0 - 255

Retransmit Interval Per-Interface.

Syntax: ip ospf [*ip-address* | all] retransmit-interval < 0 - 3600 >

Used in the VLAN context to enable changing the retransmission interval for link-state advertisements (LSAs) on an interface. The default is 5 seconds. Allows different settings for different subnet interfaces in the VLAN.

ip ospf priority < 1 - 255 >: Assigns the specified retransmit interval to all networks configured on the VLAN.

ip ospf < ip-address > priority < 1 - 255 >: Assigns the specified retransmit interval to the specified subnet on the VLAN.

ip ospf all priority < 1 - 255 >: Assigns the specified retransmit interval to all networks configured on the VLAN. (Operates the same as the **ip ospf priority** option, above.)

Default: 5 seconds; Range: 1 - 3600 seconds

Transit-Delay Per-Interface.

Syntax: ip ospf [*ip-address* | all] transit-delay < 1 - 3600 >

Used in the VLAN context to enable changing the time it takes to transmit Link State Update packets on this interface. Allows different settings for different subnet interfaces in the VLAN.

ip ospf transit-delay < 1 - 3600 >: Reconfigures the estimated number of seconds it takes to transmit a link state update packet to all networks configured on the VLAN.

ip ospf < ip-address > transit-delay < 1 - 3600 >: Reconfigures the estimated number of seconds it takes to transmit a link state update packet to all networks configured on the specified subnet on the VLAN.

ip ospf all transit-delay < 1 - 3600 >: Reconfigures the estimated number of seconds it takes to transmit a link state update packet to all networks configured on the VLAN. (Operates the same as the **ip ospf transit-delay** option, above.)

Default: 1 second; Range: 1 - 3600 seconds

Examples of Changing Per-Interface Settings. Suppose that VLAN 30 is multinetted, with two subnets in area 1 and one subnet in area 5:

```
vlan 30
 ip ospf 10.10.30.1 area 0.0.0.1
 ip ospf 10.10.31.1 area 0.0.0.1
 ip ospf 10.10.32.1 area 0.0.0.5
```

If you wanted to quickly reconfigure per-interface OSPF settings for VLAN 30, such as those listed below, you could use the commands shown in Figure 5-23.

- Assign a cost of “5” to the two subnets in area 1 and a cost of “10” to the subnet in area 5.
- Assign a dead interval of 45 seconds to the subnets in area 1 and retain the default setting (40 seconds) for the subnet in area 5.

```
ProCurve(vlan-30)# ip ospf cost 5
ProCurve(vlan-30)# ip ospf 10.10.32.1 cost 10

ProCurve(vlan-30)# ip ospf dead-interval 45
ProCurve(vlan-30)# ip ospf 10.10.32.1 dead-interval 40
```

Callout boxes:

- Configures a cost of “5” on all subnets in VLAN 30.
- Configures a cost of “10” on the 10.10.32.0 subnet in VLAN 30.
- Configures a dead interval of 45 seconds on all subnets in VLAN 30.
- Returns the 10.10.32.0 subnet in VLAN 30 to a default dead interval of 40 seconds.

Figure 5-23. Example of Reconfiguring Per-Interface Settings in a Multinetted VLAN

12. Optional: Configuring OSPF Interface Authentication

OSPF supports two methods of authentication for each VLAN or subnet—simple password and MD5. In addition, the value can be disabled, meaning no authentication is performed. Only one method of authentication can be active on a VLAN or subnet at a time, and if one method is configured on an interface, then configuring the alternative method on the same interface automatically overwrites the first method used. In the default configuration, OSPF authentication is disabled. All interfaces in the same network or subnet must have the same authentication method (password or MD5 key chain) and credentials.

OSPF Password Authentication.

Syntax: ip ospf [*ip-address*] authentication-key < *octet-string* >
no ip ospf [*ip-address*] authentication

*Used in the VLAN interface context to configure password authentication for all interfaces in the VLAN or for a specific subnet. The password takes effect immediately, and all OSPF packets transmitted on the interface contain this password. All OSPF packets received on the interface are checked for the password. If it is not present, then the packet is dropped. To disable password authentication on an interface, use the **no** form of the command.*

[*ip-address*]: *Used in subnetted VLAN contexts where you want to assign or remove a password associated with a specific subnet. Omit this option when you want the command to apply to all interfaces configured in the VLAN.*

< *octet-string* >: *An alphanumeric string of one to eight characters. (Spaces are not allowed.) To change the password, re-execute the command with the new password.*

Use **show ip ospf interface < *ip-address* >** to view the current authentication setting. (Refer to pages 5-95 and 5-97.)

Note: *To replace the password method with the MD5 method on a given interface, overwrite the password configuration by using the MD5 form of the command shown in the next syntax description. (It is not necessary to disable the currently configured OSPF password.)*

Default: Disabled

OSPF MD5 Authentication.

Syntax: ip ospf md5-auth-key-chain < chain-name-string >
no ip ospf [ip-address] authentication

*Used in the VLAN interface context to configure MD5 authentication for all interfaces in the VLAN or for a specific subnet. The MD5 authentication takes effect immediately, and all OSPF packets transmitted on the interface contain the designated key. All OSPF packets received on the interface are checked for the key. If it is not present, then the packet is dropped. To disable MD5 authentication on an interface, use the **no** form of the command.*

Note: Before using this authentication option, you must configure one or more key chains on the routing switch by using the Key Management System (KMS) described in the chapter titled “Key Management System” in the [Access Security Guide](#) for your routing switch

[ip-address]: Used in subnetted VLAN contexts where you want to assign or remove MD5 authentication associated with a specific subnet. Omit this option when you want the command to apply to all interfaces configured in the VLAN.

< chain-name-string >: The name of a key generated using the **key-chain < chain_name > key < key_id >** command. To change the MD5 authentication configured on an interface, re-execute the command with the new MD5 key.

Use **show ip ospf interface < ip-address >** to view the current authentication setting. (Refer to pages 5-95 and 5-97.)

Note: To replace the MD5 method with the password method on a given interface, overwrite the MD5 configuration by using the password form of the command shown in the next syntax description. (It is not necessary to disable the currently configured OSPF MD5 authentication.)

Default: Disabled

13. Configuring an ABR To Use a Virtual Link to the Backbone

All ABRs (area border routers) must have either a direct, physical or indirect, virtual link to the OSPF backbone area (0.0.0.0 or 0). If an ABR does not have a physical link to the area backbone, the ABR can use a *virtual link* to provide a logical connection to another ABR having a direct physical connection to the area backbone. Both ABRs must belong to the same area, and this area becomes a transit area for traffic to and from the indirectly connected ABR.

Note

A backbone area can be purely virtual with no physical backbone links. Also note that virtual links can be “daisy chained”. If so, it may not have one end physically connected to the backbone.

Because both ABRs in a virtual link connection are in the same OSPF area, they use the same *transit area ID*. This setting is automatically determined by the ABRs and should match the area ID value configured on both ABRs in the virtual link.

The ABRs in a virtual link connection also identify each other with a *neighbor router* setting:

- On the ABR having the direct connection to the backbone area, the *neighbor router* is the IP address of the router interface needing a logical connection to the backbone.
- On the opposite ABR (the one needing a logical connection to the backbone), the *neighbor router* is the IP address of the ABR that is directly connected to the backbone.

Notes

By default, the router ID is the lowest numbered IP address or (user-configured) loopback interface configured on the device. For more information or to change the router ID, see “Changing the Router ID” on page 5-13.

When you establish an area virtual link, you must configure it on both of the ABRs (both ends of the virtual link).

Configuring a Virtual Link

Syntax: ip ospf area < area-id > virtual-link < ip-address >

Used on a pair of ABRs at opposite ends of a virtual link in the same area to configure the virtual link connection.

< area-id >: *This must be the same for both ABRs in the link, and is the area number of the virtual link transit area in either decimal or dotted decimal format.*

< ip-address >: *On an ABR directly connected to the backbone area, this value must be the IP address of an ABR (in the same area) needing a virtual link to the backbone area as a substitute for a direct physical connection. On the ABR that needs the virtual link to the backbone area, this value must be the IP address of the ABR (in the same area) having a direct physical connection to the backbone area.*

Example. Figure 5-24 shows an OSPF ABR, routing switch “A”, that lacks a direct connection to the backbone area (area 0). To provide backbone access to routing switch “A”, you can add a virtual link between routing switch “A” and routing switch “C”, using area 1 as a transit area. To configure the virtual link, define it on the routers that are at each end of the link. No configuration for the virtual link is required on the other routers on the path through the transit area (such as routing switch “B” in this example).

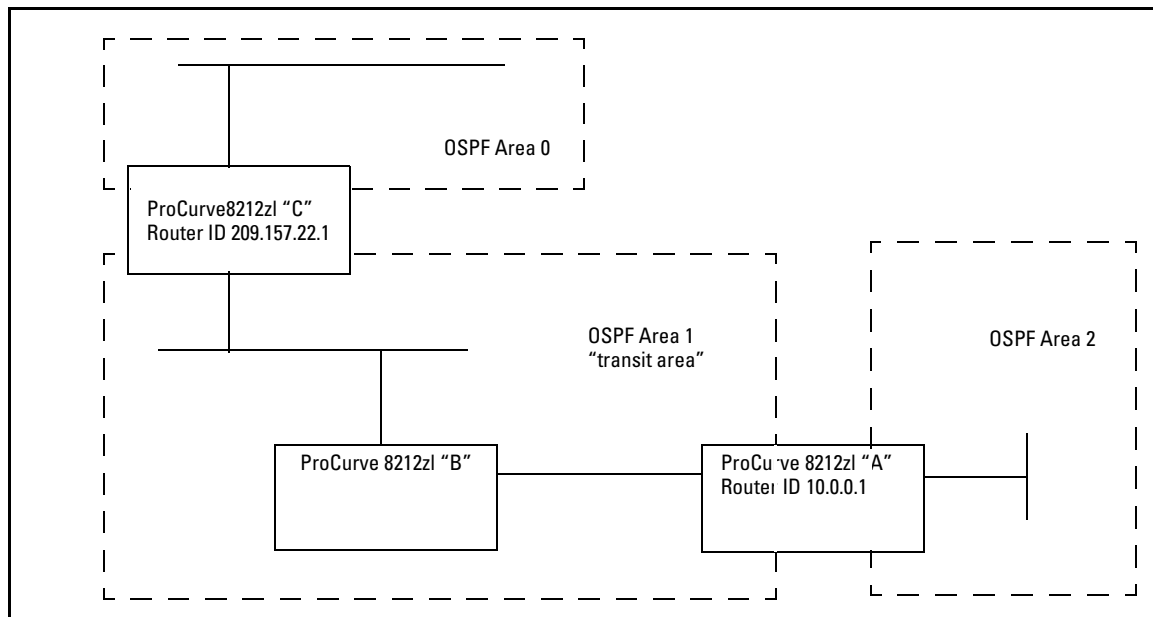


Figure 5-24. Defining OSPF virtual links within a network

To configure the virtual link on routing switch “A”, enter the following command specifying the area 1 interface on routing switch “C”:

```
ProCurve(ospf)# area 1 virtual-link 209.157.22.1
```

To configure the virtual link on routing switch “C”, enter the following command specifying the area 1 interface on routing switch “A”:

```
ProCurve(ospf)# area 1 virtual-link 10.0.0.1
```

Refer to “Optional: Adjust Virtual Link Performance by Changing the Interface Settings” on page 5-85 below for descriptions of virtual link interface parameters you can either use in their default settings or reconfigure as needed.

Optional: Adjust Virtual Link Performance by Changing the Interface Settings

The following OSPF interface parameters are automatically set to their default values for virtual links. No change to the defaults is usually required unless needed for specific network conditions. This is a subset of the parameters described under “11. Optional: Adjust Performance by Changing the VLAN or Subnet Interface Settings” on page 5-75. (The **cost** and **priority** settings are not configurable for a virtual link, and the commands for reconfiguring the settings are accessed in the router OSPF context instead of the VLAN context.)

Note

The parameter settings described in this section for virtual links must be the same on the ABRs at both ends of a given link.

Parameter	Default	Page
dead-interval	40 seconds	below
hello-interval	10 seconds	5-86
retransmit-interval	5 seconds	5-86
transit-delay	1 second	5-87

Dead Interval on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > dead-interval < 1 - 65535 >

*Used in the router OSPF context on both ABRs in a virtual link to change the number of seconds that a neighbor router waits for a hello packet from the specified interface before declaring the interface “down”. This should be some multiple of the Hello interval. The **dead-interval** setting must be the same on both ABRs on a given virtual link.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-address >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-83.)*

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-106.)*

Default: 40 seconds; Range: 1 - 65535 seconds.

Hello Interval on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > hello-interval < 1 - 65535 >

*Used in the router OSPF context on both ABRs in a virtual link to indicate the length of time between the transmission of hello packets between the ABRs on opposite ends of the virtual link. The value can be from 1 – 65535 seconds. The default is 10 seconds. The **hello-interval** setting must be the same on both ABRs on a given virtual link.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-address >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-83.)*

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-106.)*

Default: 10 seconds; Range: 1 - 65535 seconds.

Retransmit Interval on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > retransmit-interval < 1 - 3600 >

*Used in the router OSPF context on both ABRs in a virtual link to change the number of seconds between link-state advertisement (LSA) retransmissions on the virtual link. The default is 5 seconds. The **retransmit-interval** setting must be the same on both ABRs on a given virtual link. This value is also used when retransmitting database description and link-state request packets.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-address >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-83.)*

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-106.)*

Default: 5 seconds; Range: 1 - 3600 seconds

Transit-Delay on a Virtual Link.

Syntax: area < area-id > virtual link < ip-address > transit-delay < 0 - 3600 >

*Used in the router OSPF context on both ABRs in a virtual link to change the estimated number of seconds it takes to transmit a link state update packet over a virtual link. The **transit-delay** setting must be the same on both ABRs on a given virtual link.*

< area-id >: Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.

< ip-address >: For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of < ip-address > in the syntax description under “Configuring a Virtual Link” on page 5-83.)

*Use **show ip ospf virtual-link < ip-address >** to view the current setting. (Refer to the example on page 5-106.)*

Default: 1 second; Range: 1 - 3600 seconds

Example. To change the hello-interval on the virtual link configured for the network in figure 5-24 on page 5-84 to 60 seconds:

- On routing switch “A” (IP address 10.0.0.1) you would use the following command to reconfigure the current hello-interval to 60 seconds:

```
ProCurve(ospf)# area 1 virtual-link 209.157.22.1  
hello-interval 60
```

- On routing switch “C” (IP address 209.157.22.1) you would use the following command to reconfigure the current hello-interval to 60 seconds

```
ProCurve(ospf)# area 1 virtual-link 10.0.0.1  
hello-interval 60
```

Configuring OSPF Authentication on a Virtual Link

OSPF supports the same two methods of authentication for virtual links as it does for VLANs and subnets in an area—password and MD5. In the default configuration, OSPF authentication is disabled. Only one method of authentication can be active on a virtual link at a time, and if one method is configured on a virtual link, then configuring the alternative method on the same link automatically replaces the first method with the second. Both ends of a virtual link must use the same authentication method (none, password, or MD5 key chain) and related credentials. (Any interfaces that share a VLAN or subnet with the interface used on an ABR for a virtual link, including intermediate routing switches, must be configured with the same OSPF authentication.)

OSPF Password Authentication on a Virtual Link.

Syntax: `area < area-id > virtual-link < ip-addr > authentication-key < octet-string >`
`no area 1 virtual-link < ip-address > authentication`

*Used to configure password authentication in the router OSPF context on both ABRs in a virtual link. The password takes effect immediately, and all OSPF packets transmitted on the link contain this password. Every OSPF packet received on the interface for the virtual link on each ABR is checked for the password. If it is not present, then the packet is dropped. To disable password authentication on an ABR interface used for a virtual link, use the **no** form of the command. The password must be the same on both ABRs on a given virtual link.*

< area-id >: *Specifies the OSPF area in which both ABRs in a given virtual link operate. In this use, the area ID is sometimes termed “transit area ID”. This value must be the same for both ABRs in the virtual link.*

< ip-addr >: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of **< ip-address >** in the syntax description under “Configuring a Virtual Link” on page 5-83.)*

< octet-string >: *An alphanumeric string of one to eight characters. (Spaces are not allowed.) To change the password, re-execute the command with the new password.*

Note: *To replace the password method with the MD5 method on a given interface, overwrite the password configuration by using the MD5 form of the command shown in the next syntax description. (It is not necessary to disable the currently configured OSPF password.)*

Default: Disabled

OSPF MD5 Authentication on a Virtual Link.

Syntax: ip ospf md5-auth-key-chain < chain-name-string >
no ip ospf [ip-address] authentication

*Used to configure MD5 authentication in the router OSPF context on both ABRs in a virtual link . The MD5 authentication takes effect immediately, and all OSPF packets transmitted on the link contain the designated key. Every OSPF packet received on the interface for the virtual link on each ABR is checked for the key. If it is not present, then the packet is dropped. To disable MD5 authentication on an ABR interface used for a virtual link, use the **no** form of the command. The password must be the same on both ABRs on a given virtual link.*

Note: *Before using this authentication option, you must configure one or more key chains on the routing switch by using the Key Management System (KMS) described in the chapter titled “Key Management System” in the [Access Security Guide](#) for your routing switch.*

[ip-address]: *For an ABR in a given virtual link, this is the IP address used to create the link on that ABR. (This IP address matches the IP address of the interface on the opposite end of the virtual link. Refer to the description of < ip-address > in the syntax description under “Configuring a Virtual Link” on page 5-83.)*

< chain-name-string >: *The name of a key generated using the **key-chain < chain_name > key < key_id >** command. To change the MD5 authentication configured on a virtual link, re-execute the command with the new MD5 key.*

Note: *To replace the MD5 method with the password method on a virtual link, overwrite the MD5 configuration by using the password form of the command shown under “OSPF Password Authentication on a Virtual Link” on page 5-88. (It is not necessary to disable the currently configured OSPF MD5 authentication.)*

Default: Disabled

Displaying OSPF Information

You can use CLI commands to display the following OSPF information:

OSPF Information Type	Page
General Information	5-91
Area information	5-92
External link state information	5-93
Interface information	5-95
Link state information	5-98
Neighbor information	5-101
Route information	5-107
Virtual Neighbor information	5-104
Virtual Link information	5-105
OSPF Traps enabled	5-109

Displaying General OSPF Configuration Information

To display general OSPF configuration information, enter **show ip ospf general** at any CLI level:

```
ProCurve# show ip ospf general

OSPF General Status

  OSPF protocol           : enabled
  Router ID               : 10.0.8.36
  RFC 1583 compatibility  : compatible

  Intra-area distance     : 110
  Inter-area distance     : 110
  AS-external distance    : 110

  Default import metric   : 1
  Default import metric type : external type 2

  Area Border             : yes
  AS Border               : yes
  External LSA Count      : 9
  External LSA Checksum Sum : 408218
  Originate New LSA Count : 24814
  Receive New LSA Count   : 14889
```

Figure 5-25. Example of Show IP OSPF General Output

Syntax: show ip ospf general

The following fields are shown in the OSPF general status display:

Table 5-9. CLI Display of OSPF General Information

Field	Content
OSPF protocol	whether OSPF is currently enabled.
Router ID	the Router ID that this routing switch is currently using to identify itself
RFC 1583 compatibility	whether the routing switch is currently using RFC 1583 (compatible) or RFC 2328 (non-compatible rules for calculating external routes).
Intra-area distance	the administrative distance for routes within OSPF areas

Field	Content
Inter-area distance	the administrative distance for routes between areas within the same OSPF domain
AS-external	the administrative distance for routes between the OSPF domain and other, Exterior Gateway Protocol domains
Default import metric	the default metric that will be used for any routes redistributed into OSPF by this routing switch
Default import metric type	the metric type (type 1 or type 2) that will be used for any routes redistributed into OSPF by this routing switch
Area Border	whether this routing switch is currently acting as an area border router
AS Border	whether this routing switch is currently acting as an autonomous system border router (redistributing routes)
External LSA Count	the total number of external LSAs currently in the routing switch's link state database
External LSA Checksum Sum	the sum of the checksums of all external LSAs currently in the routing switch's link state database (quick check for whether database is in sync with other routers in the routing domain)
Originate New LSA Count	count of the number of times this switch has originated a new LSA
Receive New LSA Count	count of the number of times this switch has received a new LSA

Displaying OSPF Area Information

To display OSPF area information, enter **show ip ospf area** at any CLI level:

```
ProCurve(config)# show ip ospf area

OSPF Area Information

Area ID          Type    Cost  SPFR  ABR  ASBR  LSA  Checksum
-----
0.0.0.0          normal  0     1     0    0     1    0x0000781f
192.147.60.0    normal  0     1     0    0     1    0x0000fee6
192.147.80.0    stub    1     1     0    0     2    0x000181cd
```

Figure 5-26. Example of Show IP OSPF Area Output

Syntax: show ip ospf area [*ospf-area-id*]

The `[ospf-area-id]` parameter shows information for the specified area. If no area is specified, information for all the OSPF areas configured is displayed.

The OSPF area display shows the following information:

Table 5-10. CLI Display of OSPF Area Information

Field	Content
Area ID	The identifier for this area.
Type	The area type, which can be either “normal” or “stub”.
Cost	The metric for the default route that the routing switch will inject into a stub area if the routing switch is an ABR for the area. This value only applies to stub areas.
SPFR	The number of times the routing switch has run the shortest path first route calculation for this area.
ABR	The number of area border routers in this area.
ASBR	The number of autonomous system border routers in this area.
LSA	The number of LSAs in the link state database for this area.
Chksum(Hex)	The sum of the checksums of all LSAs currently in the area’s link state database. This value can be compared to the value for other routers in the area to verify database synchronization.

Displaying OSPF External Link State Information

To display external link state information, enter **show ip ospf external-link-state** at any CLI level. When you enter this command, an output similar to the following is displayed:

```
ProCurve# show ip ospf external-link-state

OSPF External LSAs

Link State ID   Router ID      Age   Sequence #   Checksum
-----
10.3.7.0        10.0.8.37     232  0x80000005  0x0000d99f
10.3.8.0        10.0.8.37     232  0x80000005  0x0000cea9
10.3.9.0        10.0.8.37     232  0x80000005  0x0000c3b3
10.3.10.0       10.0.8.37     232  0x80000005  0x0000b8bd
10.3.33.0       10.0.8.36     1098 0x800009cd  0x0000b9dd
```

Figure 5-27. Example of Show IP OSPF External-Link-State Output

Syntax: show ip ospf external-link-state

The OSPF external link state display shows the following information:

Table 5-11. CLI Display of OSPF External Link State Information

Field	Content
Link State ID	LSA ID for this LSA. Normally, the destination of the external route, but may have some "host" bits set.
Router ID	Router ID of the router that originated this external LSA.
Age	Current age (in seconds) of this LSA.
Sequence #	Sequence number of the current instance of this LSA.
Chksum(Hex)	LSA checksum value.

Syntax show ip ospf external-link-state [status] [*subset-options*]
:

router-id < *ip-addr* >

*Subset option to filter displayed external-link-state data to show LSAs with the specified router ID only. Can also be filtered by using the **link-state-id** or **sequence-number** options.*

sequence-number < *integer* >

*Subset option to filter displayed external-link-state data to show LSAs with the specified sequence number. Can also be filtered by using the **link-state-id** or **router-id** options.*

link-state-id < *ip-addr* >

*Subset option to filter displayed external-link-state data to show LSAs with the specified ID only. Can also be filtered by using the **sequence-number** or **router-id** options.*

Syntax: show ip ospf external-link-state [status] advertise

*Displays the hexadecimal data in the specified LSA packet, the actual contents of the LSAs. Can also be filtered by using the **link-state-id**, **router-id**, or **sequence-number** options.*

```

ProCurve# show ip ospf external-link-state advertise
OSPF External LSAs
Advertisements
-----
000302050a0307000a00082580000005d99f0024ffffff008000000a0000000000000000
000302050a0308000a00082580000005cea90024ffffff008000000a0000000000000000
000302050a0309000a00082580000005c3b30024ffffff008000000a0000000000000000
000302050a030a000a00082580000005b8bd0024ffffff008000000a0000000000000000
000002050a0321000a000824800009cdb9dd0024ffffff00800000010000000000000000

```

Figure 5-28. Example of the Output for Show IP OSPF External-Link-State Advertise

Displaying OSPF Interface Information

To display OSPF interface information, enter **show ip ospf interface** at any CLI level:

```

ProCurve# show ip ospf interface
OSPF Interface Status
IP Address      Status   Area ID      State   Auth-type  Cost   Priority
-----
10.3.18.36     enabled  10.3.16.0    BDR    none       1      1
10.3.53.36     enabled  10.3.48.0    BDR    none       1      1

```

Figure 5-29. Example of the Output for Show IP OSPF Interface

Syntax: show ip ospf interface [vlan <vlan-id> | <ip-addr>]

The OSPF interface display shows the following information:

Table 5-12. CLI Display of OSPF Interface Information

Field	Content
IP Address	The local IP address for this interface.
Status	enabled or disabled - whether OSPF is currently enabled on this interface.
Area ID	The ID of the area that this interface is in.

Field	Content
State	The current state of the interface. The value will be one of the following: <ul style="list-style-type: none">• DOWN - the underlying VLAN is down• WAIT - the underlying VLAN is up, but we are waiting to hear hellos from other routers on this interface before we run designated router election• DR - this switch is the designated router for this interface• BDR - this switch is the backup designated router for this interface• DROTHER - this router is not the designated router or backup designated router for this interface
Auth-type	none or simple - will be none if no authentication key is configured, simple if an authentication key is configured. All routers running OSPF on the same link must be using the same authentication type and key.
Chain	The name of the key chain configured for the specified interface. (Refer to the chapter titled "Key Management System" in the <i>Access Security Guide</i> for your routing switch.
Cost	The OSPF's metric for this interface.
Priority	This routing switch's priority on this interface for use in the designated router election algorithm.

The < **ip-addr** > parameter displays the OSPF interface information for the specified IP address.

The < **vlan-id** > parameter displays the OSPF interface information for the specified IP address.

Displaying OSPF Interface Information for a Specific VLAN or IP Address

To display OSPF interface information for a specific VLAN or IP address, enter **show ip ospf interface < ip-addr >** at any CLI level. For example:

```
ProCurve(ospf)# sho ip ospf int 10.10.50.1

OSPF Interface Status for 10.3.1836

IP Address       : 10.3.18.36           Status  : enabled
Area ID         : 10.3.16.0

State  : BDR                          Auth-type : none
Cost   : 1                             Chain     :
Type   : BCAST                         Priority  : 1

Transit Delay   : 1                    Retrans Interval : 5
Hello Interval  : 10                   Rtr Dead Interval : 40
Designated Router : 10.3.18.34         Events           : 3
Backup Desig. Rtr : 10.3.18.36
```

Figure 5-30. Example of Show IP OSPF Interface < ip-addr > Output

Syntax: show ip ospf interface [vlan < vlan-id > | < ip-addr >]

The OSPF interface display for a specific VLAN or IP address has the same information as the non-specific **show ip ospf interface** command for the **IP Address, Area ID, Status, State, Auth-type, Cost, and Priority** fields. See the information for the general command in table 5-12 on page 5-95 for definitions of these fields.

The **show ip ospf interface** command for a specific VLAN or IP address shows the following additional information:

Table 5-13. CLI Display of OSPF Interface Information – VLAN or IP Address

Field	Content
Type	Will always be BCAST for interfaces on this routing switch. Point-to-point or NBMA (frame relay or ATM) type interfaces are not supported on the switches covered in this guide.
Transit Delay	Configured transit delay for this interface.

Field	Content
Retrans Interval	Configured retransmit interval for this interface.
Hello Interval	Configured hello interval for this interface.
Rtr Dead Interval	Configured router dead interval for this interface.
Designated Router	IP address of the router that has been elected designated router on this interface.
Backup Desig. Rtr	IP address of the router that has been elected backup designated router on this interface.
Events	Number of times the interface state has changed.

If you use **show ip ospf interface vlan < vlan-id >**, the output will be the same as shown in the previous table, but for the IP address on the indicated VLAN.

Displaying OSPF Link State Information

To display OSPF link state information, enter **show ip ospf link-state** at any CLI level. When you enter this command, the switch displays an output similar to the following for all configured areas:

```

OSPF Link State Database for Area 0.0.0.0
      Advertising
LSA Type  Link State ID  Router ID      Age  Sequence #  Checksum
-----
Router    10.0.8.32             10.0.8.32      65   0x80000281  0x0000a7b6
Router    10.0.8.33             10.0.8.33     1638 0x80000005  0x0000a7c8
Network   10.3.2.37             10.0.8.37     1695 0x80000006  0x00000443
Summary   10.3.16.0             10.0.8.33     1638 0x80000007  0x0000c242
Summary   10.3.16.0             10.0.8.35     1316 0x80000008  0x0000aa58
Summary   10.3.17.0             10.0.8.33     1638 0x8000027b  0x0000becf
Summary   10.3.17.0             10.0.8.35     1316 0x80000008  0x0000a957
AsbSummary 10.0.8.36             10.0.8.33     1412 0x80000002  0x00002cba

OSPF Link State Database for Area 10.3.16.0
      Advertising
LSA Type  Link State ID  Router ID      Age  Sequence #  Checksum
-----
Router    10.0.8.33             10.0.8.33     1727 0x8000027e  0x0000d53c
Router    10.0.8.34             10.0.8.34     1420 0x80000283  0x0000de4f
Network   10.3.16.34          10.0.8.34     1735 0x80000005  0x00001465

```

Figure 5-31. Example of Show IP OSPF Link-State Output

The OSPF link state display shows the following contents of the LSA database; one table for each area:

Table 5-14. CLI Display of OSPF Link State Information

Field	Content
LSA Type	Type of LSA. The possible types are: <ul style="list-style-type: none"> • Router • Summary • Network • AsbSummary
Link State ID	LSA ID for this LSA. The meaning depends on the LSA type.
Advertised Router ID	Router ID of the router that originated this LSA.
Age	Current age (in seconds) of this LSA.
Sequence #	Sequence number of the current instance of this LSA.
Chksum(Hex)	LSA checksum value.

Syntax show ip ospf link-state [status] [subset-options] [< advertise [subset-options] >]:

advertise: *Displays the hexadecimal data in LSA packets (advertisements) for the OSPF area(s) configured on the routing switch. The output can also be filtered by area (**area-id**), **link-state-id**, **router-id**, **sequence-number**, and/or **type**.*

Default: All OSPF areas configured on the routing switch.

ospf-area-id: *Used to restrict display of LSA database or advertisements to show only the data from a specific OSPF area. Can also be used with other subset options (**router-id**, **sequence-number**, **external link-state-id**, and/or **type**) to further define the source of displayed information.*

link-state-id < ip-addr >

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified IP address as a link-state ID. Can also be used with other subset options (**ospf-area-id**, **router-id**, **sequence-number**, **external link-state-id**, and **type**) to further define the source of displayed information.*

router-id < ip-addr >

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified router ID. Can also be used with other subset options (**ospf-area-id**, **link-state-id**, **sequence-number**, and **type**) to further define the source of displayed information.*

sequence-number < integer >

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified sequence number. Can also be used with other subset options (**ospf-area-id**, **link-state-id**, **router-id**, and **type**) to further define the source of displayed information.*

type < router | network | summary | as-summary | external | multicast | nssa >

*Used to restrict display of LSA database or advertisements to show only the data from sources having the specified type. Can also be used with other subset options (**ospf-area-id**, **link-state-id**, **router-id**, and **sequence-number**) to further define the source of displayed information.*

An example of **show ip ospf link-state advertise** output is:

```
ProCurve_8212(config)# show ip ospf link-state advertise
OSPF Link State Database for Area 0.0.0.0
  Advertisements
-----
000202010a0008200a00082080000281a7b60054000000050a030e00ffffff0003000001...
000202010a0008210a00082180000006a5c90024010000010a0008230a03112104000002
000102010a0008230a00082380000015755d006c010000070a030600ffffff0003000001...
000202020a0302250a0008258000000702440024ffffff000a0008250a0008230a000820
000202030a0310000a00082180000008c043001cffffff0000000002
000102030a0310000a00082380000009a859001cffffff0000000001
000002030a0310000a00082480000009ac53001cffffff0000000002
000202040a0008240a000821800000032abb001c000000000000000b
000102040a0008240a00082380000004c12a001c0000000000000002

OSPF Link State Database for Area 10.3.16.0
  Advertisements
-----
000202010a0008210a0008218000027fd33d00540500000050a031900ffffff0003000001...
000102010a0008220a00082280000284dc500060000000060a031500ffffff0003000001...
000102020a0311220a0008228000027bf9080020ffffff000a0008220a000821
```

Figure 5-32. Example of the Output for Show IP OSPF Link-State Advertise

Displaying OSPF Neighbor Information

To display OSPF information for all neighbors, enter **show ip ospf neighbor** at any CLI level:

OSPF Neighbor Information							
Router ID	Pri	IP Address	NbIfState	State	Rxmt	QLen	Events
10.0.8.34	1	10.3.18.34	DR	FULL	0		6
10.3.53.38	1	10.3.53.38	DR	FULL	0		6

Figure 5-33. Example of Show IP OSPF Neighbor Output

Syntax: show ip ospf neighbor [*ip-addr*]

The [*ip-addr*] can be specified to retrieve detailed information for the specific neighbor only. This is the IP address of the neighbor, not the router ID.

This display shows the following information.

Table 5-15. CLI Display of OSPF Neighbor Information

Field	Description
Router ID	The router ID of the neighbor.
Pri	The OSPF priority of the neighbor. The priority is used during election of the Designated Router (DR) and Backup designated Router (BDR).
IP Address	The IP address of this routing switch's interface with the neighbor.
NbIfState	The neighbor interface state. The possible values are: <ul style="list-style-type: none"> DR – this neighbor is the elected designated router for the interface. BDR – this neighbor is the elected backup designated router for the interface. blank – this neighbor is neither the DR or the BDR for the interface.

Field	Description
State	<p>The state of the conversation (the adjacency) between your routing switch and the neighbor. The possible values are:</p> <ul style="list-style-type: none">• INIT – A Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor. (The switch itself did not appear in the neighbor's Hello packet.) All neighbors in this state (or higher) are listed in the Hello packets sent from the associated interface.• 2WAY – Communication between the two routers is bidirectional. This is the most advanced state before beginning adjacency establishment. The Designated Router and Backup Designated Router are selected from the set of neighbors in the 2Way state or greater.• EXSTART – The first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Description (DD) sequence number. Neighbor conversations in this state or greater are called adjacencies.• EXCHANGE – The switch is describing its entire link state database by sending Database Description packets to the neighbor. Each Database Description packet has a DD sequence number, and is explicitly acknowledged. Only one Database Description packet can be outstanding at any time. In this state, Link State Request packets can also be sent asking for the neighbor's more recent advertisements. All adjacencies in Exchange state or greater are used by the flooding procedure. In fact, these adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets.• LOADING – Link State Request packets are sent to the neighbor asking for the more recent advertisements that have been discovered (but not yet received) in the Exchange state.• FULL – The neighboring routers are fully adjacent. These adjacencies will now appear in router links and network link advertisements.
Rxmt QLen	Remote transmit queue length – the number of LSAs that the routing switch has sent to this neighbor and for which the routing switch is awaiting acknowledgements.
Events	The number of times the neighbor's state has changed.

Displaying OSPF Redistribution Information

As described under “2. Enable Route Redistribution” on page 5-69, you can configure the routing switch to redistribute connected, static, and RIP routes into OSPF. When you redistribute a route into OSPF, the routing switch can use OSPF to advertise the route to its OSPF neighbors.

To display the status of the OSPF redistribution, enter **show ip ospf redistribute** at any CLI context level:

```
ProCurve# show ip ospf redistribute

OSPF redistributing
Route type Status
-----
connected  enabled
static      enabled
rip         enabled
```

Figure 5-34. Example of Output for Show IP OSPF Redistribute

The display shows whether redistribution of each of the route types, connected, static, and RIP is enabled.

Displaying OSPF Redistribution Filter (restrict) Information

As described under “7. Optional: Configure for External Route Redistribution in an OSPF Domain” on page 5-67, you can configure the redistribution filters on the routing switch to restrict route redistribution by OSPF.

To display the status of the OSPF redistribution filters, enter **show ip ospf restrict** at any CLI context level.

```
ProCurve# show ip ospf restrict

OSPF restrict list

IP Address      Mask
-----
10.0.8.0        255.255.248.0
15.0.0.0        255.0.0.0
```

Figure 5-35. Example of Output for Show IP OSPF Restrict

This display shows the configured restrict entries.

Displaying OSPF Virtual Neighbor Information

If virtual links are configured on the routing switch, you can display OSPF virtual neighbor information by entering **show ip ospf virtual-neighbor** at any CLI level.

```

OSPF Virtual Interface Neighbor Information

Router ID          Area ID           State    IP Address      Events
-----
10.0.8.33         10.3.16.0        FULL    10.3.17.33      5
10.0.8.36         10.3.16.0        FULL    10.3.18.36      5
  
```

Figure 5-36. Example of Output for Show IP OSPF Virtual-Neighbor

Syntax: `show ip ospf virtual-neighbor [area < area-id > | < ip-address >]`

This display shows the following information.

Table 5-16. CLI Display of OSPF Virtual Neighbor Information

Field	Description
Router ID	The router ID of this virtual neighbor (configured).
Area ID	The area ID of the transit area for the virtual link to this neighbor (configured).
State	The state of the adjacency with this virtual neighbor. The possible values are the same as the OSPF neighbor states. See the State parameter definition in table 5-15 on page 5-101. Note that virtual neighbors should never stay in the 2WAY state.
IP Address	IP address of the virtual neighbor that the routing switch is using to communicate to that virtual neighbor.
Events	The number of times the virtual neighbor's state has changed.

Notice from the syntax statement that **ip-address** can be specified to display detailed information for a particular virtual neighbor. If an **area-id** is specified only virtual neighbors belonging to that area are shown.

Displaying OSPF Virtual Link Information

If virtual links are configured on a routing switch, you can display OSPF virtual link information by entering **show ip ospf virtual-link** at any CLI level.

```
ProCurve# show ip ospf virtual-link

OSPF Virtual Interface Status

  Transit AreaID  Neighbor Router  Authentication  Interface State
  -----
  10.3.16.0       10.0.8.33        none            P2P
  10.3.16.0       10.0.8.36        none            P2P
```

Figure 5-37. Example of Output for Show IP OSPF Virtual-Link

Syntax: show ip ospf virtual-link [area < *area-id* > | < *ip-address* >]

This display shows the following information.

Table 5-17. CLI Display of OSPF Virtual Link Information

Field	Description
Transit Area ID	Area ID of transit area for the virtual link.
Neighbor Router	Router ID of the virtual neighbor.
Authentication	none or simple (same as for normal interface).
Interface State	The state of the virtual link to the virtual neighbor. The possible values are: <ul style="list-style-type: none"> DOWN – the routing switch has not yet found a route to the virtual neighbor. P2P – (point-to-point) the routing switch has found a route to the virtual neighbor. Virtual links are “virtual” serial links, hence the point-to-point terminology.

Notice from the syntax statement that **ip-address** can be specified to display detailed information for a particular virtual neighbor. If an **area-id** is specified only virtual links belonging to that area are shown.

Example:. To get OSPF virtual link information for IP address 10.0.8.33, enter **show ip ospf virtual-link 10.0.8.33**. A display similar to the following is shown.

```
ProCurve# show ip ospf virtual-link 10.0.8.33

OSPF Virtual Interface Status for interface 10.0.8.33
  Transit AreaID   : 10.3.16.0
  Neighbor Router  : 10.0.8.33

  Authentication   : none                Chain           :
  Interface State  : P2P                 Transit Delay   : 1
  Events           : 1                   Rtr Interval    : 5
  Dead Interval    : 40                   Hello Interval  : 10
```

Figure 5-38. Example of Output for Show IP OSPF Virtual-Link < ip-addr >

In this display, these fields show the same type of information as described for the general OSPF virtual link display: **Transit Area ID**, **Neighbor Router**, **Authentication**, and **Interface State**. This display shows the following additional information:

Table 5-18. CLI Display of OSPF Virtual Link Information – Specific IP Address

Field	Description
Events	The number of times the virtual link interface state has changed.
Transit delay	The configured transit delay for the virtual link.
Rtr Interval	The configured retransmit interval for the virtual link.
Hello Interval	The configured hello interval for the virtual link.
Dead Interval	The configured router dead interval for the virtual link

Displaying OSPF Route Information

To display OSPF route and other OSPF configuration information, enter **show ip ospf** at any CLI level:

```
ProCurve# show ip ospf

OSPF Configuration Information

  OSPF protocol   : enabled
  Router ID      : 10.0.8.35

Currently defined areas:
```

Area ID	Type	Stub Default Cost	Stub Summary LSA	Stub Metric Type
backbone	normal	1	don't send	ospf metric
10.3.16.0	normal	1	don't send	ospf metric
10.3.32.0	normal	1	don't send	ospf metric

```
Currently defined address ranges:
```

Area ID	LSA Type	IP Network	Network Mask	Advertise
10.3.16.0	Summary	10.3.16.0	255.255.255.0	yes

```
OSPF interface configuration:
```

IP Address	Area ID	Admin Status	Type	Authen Type	Cost	Pri
10.3.2.35	backbone	enabled	BCAST	none	1	1
10.3.3.35	backbone	enabled	BCAST	none	1	1
10.3.16.35	10.3.16.0	enabled	BCAST	none	1	1
10.3.32.35	10.3.32.0	enabled	BCAST	none	1	1

```
OSPF configured interface timers:
```

IP Address	Transit Delay	Retransmit Interval	Hello Interval	Dead Interval
10.3.2.35	1	5	10	40
10.3.3.35	1	5	10	40
10.3.16.35	1	5	10	40
10.3.32.35	1	5	10	40

```
OSPF configured virtual interfaces:
```

Area ID	Router ID	Authen Type	Xmit Delay	Rxmt Intvl	Hello Intvl	Dead Interval
10.3.16.0	10.0.8.33	none	1	5	10	40
10.3.16.0	10.0.8.36	none	1	5	10	40

Figure 5-39. Example of Output for Show IP OSPF

Syntax: show ip ospf

This screen has a lot of information, most of it already covered in other show commands. The following table shows definitions for the fields:

Table 5-19. CLI Display of OSPF Route and Status Information

Field	Description
OSPF protocol	enabled or disabled – indicates if OSPF is currently enabled.
Router ID	The Router ID that this routing switch is currently using to identify itself.
Currently Defined Areas:	
Area ID	The identifier for this area.
Type	The type of OSPF area (normal or stub).
Stub Default Cost	The metric for any default route we injected into a stub area if the routing switch is an ABR for the area. This value only applies to stub areas.
Stub Summary LSA	send or don't send – indicates the state of the no-summary option for the stub area. The value indicates if the area is “totally stubby” (no summaries sent from other areas) or just “stub” (summaries sent). Only applies to stub areas, and only takes effect if the routing switch is the ABR for the area.
Stub Metric Type	This value is always ospf metric .
Currently defined address ranges:	
Area ID	The area where the address range is configured.
LSA Type	This value is always Summary .
IP Network	The address part of the address range specification.
Network Mask	The mask part of the address range specification.
Advertise	Whether we are advertising (yes) or suppressing (no) this address range.

Note

The remaining interface and virtual link information is the same as for the previously described OSPF show commands. Refer to table 5-12 (page 5-95) and table 5-13 (page 5-97).

Displaying OSPF Traps Enabled

In the default configuration, OSPF traps are disabled. Use this command to view which OSPF traps have been enabled.

Syntax: show ip ospf traps

Lists the OSPF traps currently enabled on the routing switch.

For more information on OSPF trap use, refer to “10: Optional: Change OSPF Trap Generation Choices” on page 5-74.

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

The switches covered by this guide support optional load-sharing across redundant links where the network offers two, three, or four equal-cost next-hop routes for traffic to different subnets. (All traffic for different hosts in the same subnet goes through the same next-hop router.)

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

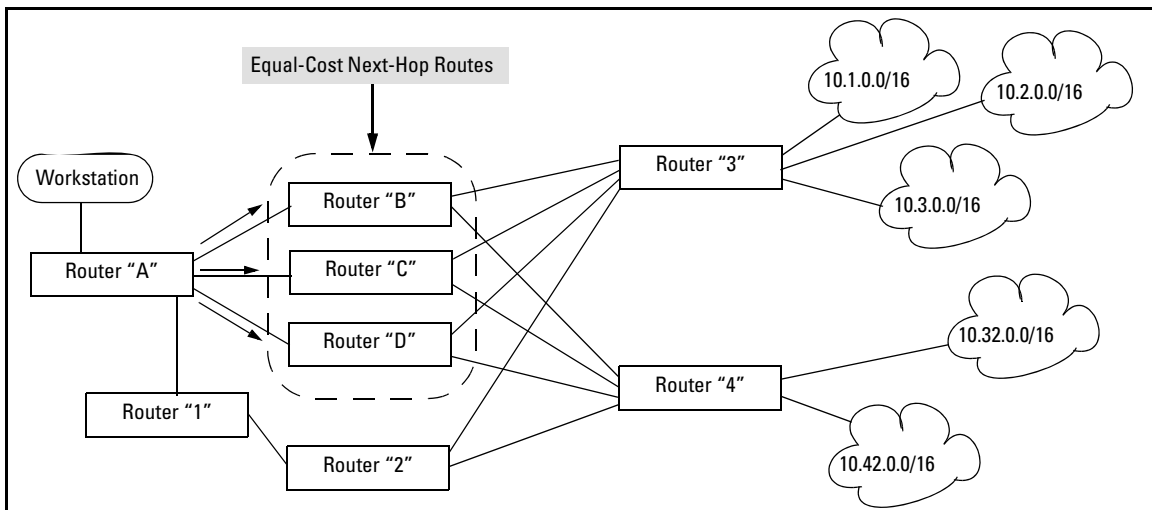


Figure 5-40. Example of Load-Sharing Traffic to Different Subnets Through Equal-Cost Next-Hop Routers

Figure 5-41. Example of a Routing Table for the Network in Figure 5-40

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

Note that IP load-sharing does not affect routed traffic to different hosts on the same subnet. That is, all traffic for different hosts on the same subnet will go through the same next-hop router. For example, if subnet 10.32.0.0 includes two servers at 10.32.0.11 and 10.32.0.22, then all traffic from router "A" to these servers will go through router "B".

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

*When OSPF is enabled and multiple, equal-cost, next-hop routes are available for traffic destinations on different subnets, this feature, by default, enables load-sharing among up to four next-hop routes. The **no** form of the command disables this load-sharing so that only one route in a group of multiple, equal-cost, next-hop routes is used for traffic that could otherwise be load-shared across multiple routes. For example, in figure 5-40 on page 5-109, the next-hop routers "B", "C", and "D" are available for equal-cost load-sharing of eligible traffic.*

Disabling IP load-sharing means that router "A" selects only one next-hop router for traffic that is actually eligible for load-sharing through different next-hop routers. (Default: Enabled with four equal-cost, next-hop routes allowed)

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

< 1 - 4 >

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

Displaying the Current IP Load-Sharing Configuration

Use the **show running** command to view the currently active IP load-sharing configuration, and **show config** to view the IP load-sharing configuration in the startup-config file. (While in its default configuration, IP load-sharing does not appear in the command output.) If IP load sharing is configured with non-

default settings (disabled or configured for either two or three equal-cost next-hop paths), then the current settings are displayed in the command output.

```
ProCurve(config)# show running
Running configuration:
; J9091A Configuration Editor; Created on release
#K.12.XX
hostname "ProCurve"
module 1 type J8702A
snmp-server community "public" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24
    ip address dhcp-bootp
    exit
[ip_load-sharing_3]
access-controller vlan-base 2000
```

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

Figure 5-42. Displaying a Non-Default IP Load-Sharing Configuration

Configuring IRDP

The ICMP Router Discovery Protocol (IRDP) is used by ProCurve routing switches to advertise the IP addresses of their router interfaces to directly attached hosts. IRDP is disabled 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 multicast 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 5-43. 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 user intervention. The protocol is composed of three components:

- DHCP client
- DHCP server
- 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.

ProCurve routing switches provide the DHCP relay agent to enable communication from a DHCP server to DHCP clients on subnets other than the one the server resides on. The DHCP relay agent transfers DHCP messages from DHCP clients located on a subnet without a DHCP server to other subnets. It also relays answers from DHCP servers to DHCP clients.

The DHCP relay agent is transparent to both the client and the server. Neither side is aware of the communications that pass through the DHCP relay agent. As DHCP clients broadcast requests, the DHCP relay agent receives the packets and forwards them to the DHCP server. During this process, the DHCP relay agent increases the hop count by one before forwarding the DHCP message to the server. A DHCP server includes the hop count from the DHCP request that it receives in the response that it returns to the client.

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).

Prerequisites for DHCP Relay Operation

For the DHCP Relay agent to work on the switch, you must complete the following steps:

1. Enable DHCP Relay on the routing switch (the default setting).
2. Ensure that a DHCP server is servicing the routing switch.
3. Enable IP Routing on the routing switch.
4. Ensure that there is a route from the DHCP server to the routing switch and back.
5. Add the IP address of the DHCP server to the list of VLANs connected to the routing switch by configuring an IP helper address.

Enabling DHCP Relay

The DHCP Relay function is enabled by default on a ProCurve routing switch. However, if DHCP has been disabled, you can re-enable it by entering the following command at the global configuration level:

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

To disable the DHCP Relay function, enter the **no** form of the command:

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

Configuring an IP Helper Address

To add the IP address of a DHCP server for a specified VLAN on a routing switch, enter the **ip helper-address** command at the VLAN configuration level as in the following example:

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

To remove the DHCP server helper address, enter the **no** form of the command:

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

Operating Notes

- You can configure up to 2048 IP helper addresses on a routing switch. The helper addresses are shared between the DHCP relay agent and UDP forwarder feature.
- A maximum of sixteen IP helper addresses is supported in each VLAN.

Hop Count in DHCP Requests

When a DHCP client broadcasts requests, the DHCP relay agent in the routing switch receives the packets and forwards them to the DHCP server (on a different subnet, if necessary). During this process the DHCP relay agent increments the hop count before forwarding DHCP packets to the server. The DHCP server, in turn, includes the hop count from the received DHCP request in the response sent back to a DHCP client.

As a result, the DHCP client receives a non-zero hop count in the DHCP response packet. Because some legacy DHCP/BootP clients discard DHCP responses which contain a hop count greater than one, they may fail to boot up properly. Although this behavior is in compliance with RFC 1542, it prevents a legacy DHCP/BootP client from being automatically configured with a network IP address.

Disabling the Hop Count in DHCP Requests

To disable the default behavior of a DHCP relay agent so that the hop count in a DHCP client request is not increased by one at each hop when it is forwarded to a DHCP server, enter the **no dhcp-relay hop-count-increment** command at the global configuration level:

```
ProCurve(config)# no dhcp-relay hop-count-increment
```

To reset the default function which increases the hop count in each DHCP request forwarded to a DHCP server, enter the following command:

```
ProCurve(config)# dhcp-relay hop-count-increment
```

Operating Notes

- By default, the DHCP relay agent increases the hop count in each DHCP request by one. You must enter the **no dhcp-relay hop-count-increment** command to disable this function.
- You enter the **no dhcp-relay hop-count-increment** command at the global configuration level. The command is applied to all interfaces on the routing switch that are configured to forward DHCP requests.
- This DHCP Relay enhancement only applies to DHCP requests forwarded to a DHCP server. The server does not change the hop count included in the DHCP response sent to DHCP clients.
- When you disable or re-enable the DHCP hop count function, no other behavior of the relay agent is affected.
- You can configure the DHCP Relay hop count function only from the CLI; you cannot configure this software feature from the drop-down menus.
- A new MIB variable, `hpDhcpRelayHopCount`, is introduced to support SNMP management of the hop count increment by the DHCP relay agent in a switch.

Verifying the DHCP Relay Configuration

Displaying the DHCP Relay Setting

Use the **show config** command (or **show running** for the running-config file) to display the current DHCP Relay setting.

Note

The DHCP relay and hop count increment settings appear in the **show config** command output only if the non-default values are configured. For more information about the DHCP hop count increment, see “Hop Count in DHCP Requests” on page 5-117.

```
ProCurve# show config
Startup configuration:
; J9091A Configuration Editor; Created on release #K.12.XX
hostname "ProCurve"
cdp run
module 1 type J8702A
ip default-gateway 18.30.240.1
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1
  ip address 18.30.240.180 255.255.248.0
  no untagged A2-A24
  exit
no dhcp-relay
no dhcp-relay hop-count-increment
```

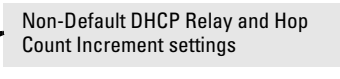


Figure 5-44. Displaying Startup Configuration with DHCP Relay and Hop Count Increment Disabled

Displaying DHCP Helper Addresses

To display the list of currently configured IP Helper addresses for a specified VLAN on the switch, enter the **show ip helper-address vlan** command.

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

*Displays the IP helper addresses of DHCP servers configured for all static VLANs in the switch or on a specified VLAN, regardless of whether the DHCP Relay feature is enabled. The **vlan <vlan-id>** parameter specifies a VLAN ID number.*

The following command lists the currently configured IP Helper addresses for VLAN 1.

```
ProCurve(config)# show ip helper-address vlan 1
IP Helper Addresses
IP Helper Address
-----
10.28.227.97
10.29.227.53
```

Figure 5-45. Displaying IP Helper Addresses

Displaying the Hop Count Setting

To verify the current setting for increasing the hop count in DHCP requests, enter the **show dhcp-relay** command. Note that the current setting is displayed next to DHCP Request Hop Count Increment.

```
ProCurve# show dhcp-relay

Status and Counters - DHCP Relay Agent
DHCP Relay Agent Enabled           : Yes
DHCP Request Hop Count Increment: Disabled
Option 82 Handle Policy           : Replace
Remote ID                          : MAC Address

Client Requests      Server Responses
Valid      Dropped   Valid      Dropped
-----
1425          2       1425        0
```

DHCP Option 82

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.

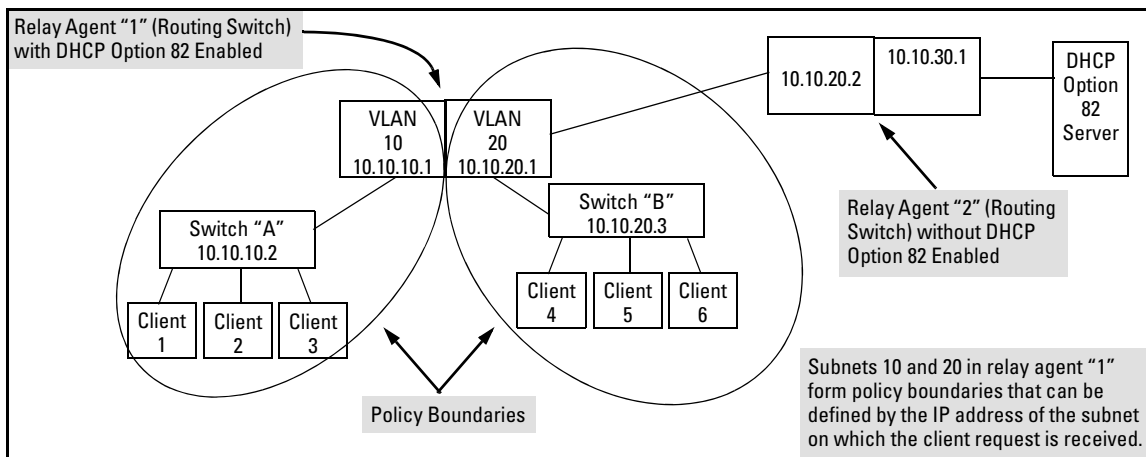


Figure 5-46. Example of a DHCP Option 82 Application

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 list on page 5-125.)

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 or the (optional) Management VLAN 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 5-124.)

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.

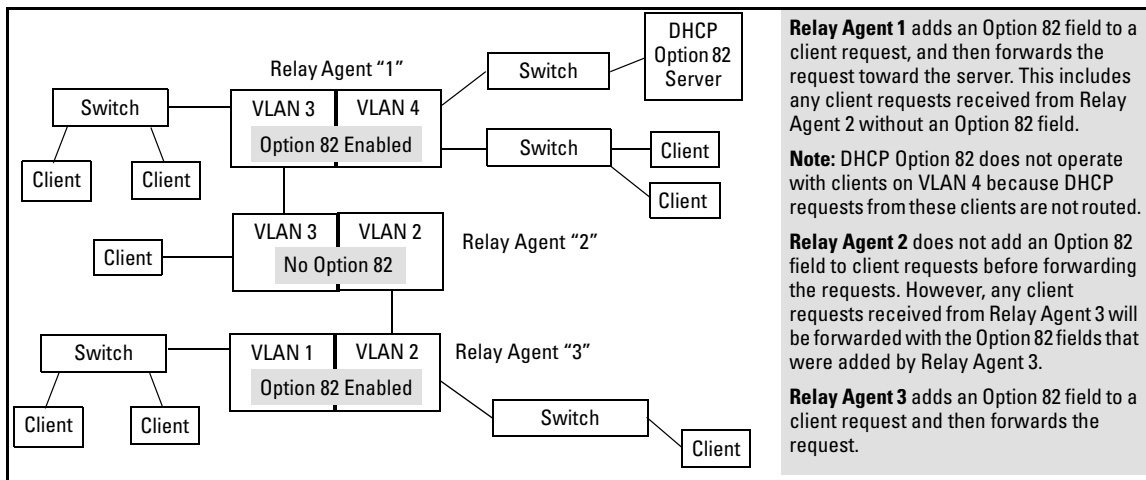


Figure 5-47. Example of DHCP Option 82 Operation in a Network with a Non-Compliant Relay Agent

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 Management VLAN option if a Management VLAN is configured and you want all DHCP clients on the routing switch to use the same IP address. (This is useful if you are applying the same IP addressing policy to DHCP client requests from ports in different VLANs on the same routing switch.) Configuring this option means the Management VLAN's IP address appears in the remote ID subfield of all DHCP requests originating with clients connected to the routing switch, regardless of the VLAN on which the requests originate.
- 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(config)# show system-information
Status and Counters - General System Information
System Name       : ProCurve
System Contact    :
System Location   :
MAC Age Time (sec) : 300
Time Zone         : 0
Daylight Time Rule : None

Firmware revision : K.11.00   Base MAC Addr   : 00110a-a50c20
ROM Version        : K.11.00   Serial Number    : SG426NB048

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

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

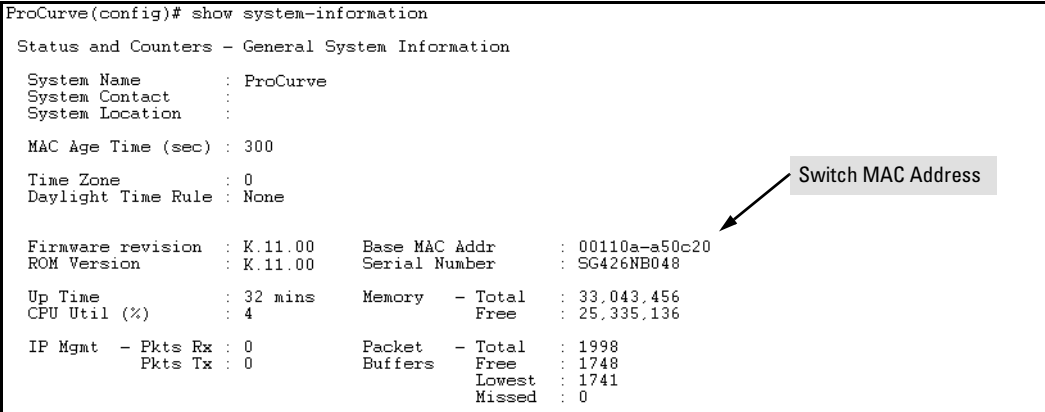


Figure 5-48. Using the CLI To View the Switch MAC Address

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

the sequential index number for that port position in the slot. (To view the Index number assignments for ports in the routing switch, use the **walkmib ifname** command.)

For example, the Circuit ID for port B11 on a ProCurve switch is “35”. (See Figure 5-49, below.)

```
ProCurve# walkmib ifname

ifName.1 = A1
ifName.2 = A2
ifName.3 = A3
ifName.4 = A4
ifName.25 = B1
ifName.26 = B2
ifName.27 = B3
ifName.28 = B4
ifName.29 = B5
ifName.30 = B6
ifName.31 = B7
ifName.32 = B8
ifName.33 = B9
ifName.34 = B10
ifName.35 = B11 ←
ifName.36 = B12
ifName.37 = B13
ifName.38 = B14
ifName.39 = B15
ifName.40 = B16
ifName.41 = B17
ifName.42 = B18
ifName.43 = B19

-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

In this example, the switch has a 4-port module installed in slot “A” and a 24-port module installed in slot “B”. Thus, the first port numbers in the listing are the Index numbers reserved for slot “A”. The first Index port number for slot “B” is “25”, and the Index port number for port B11 (and therefore the Circuit ID number) is “35”.

The Index (and Circuit ID) number for port B11 on the routing switch.

Figure 5-49. 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**).

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.

Option 82 Configuration	DHCP Client Request Packet Inbound to the Routing Switch	
	Packet Has No Option 82 Field	Packet Includes an Option 82 Field
Drop	Append an Option 82 Field	Drop causes the routing switch to drop an inbound client request with an Option 82 field already appended. If no Option 82 fields are present, drop causes the routing switch to add an Option 82 field and forward the request. As a general guideline, configure drop on relay agents at the edge of a network, where an inbound client request with an appended Option 82 field may be unauthorized, a security risk, or for some other reason, should not be allowed.

Multiple Option 82 Relay Agents in a Client Request Path

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




Figure 5-50. 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 5-51. 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 5-52. 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 5-127.) Enabling validation on the routing switch can enhance protection against DHCP server responses that are either from untrusted sources or are carrying invalid Option 82 information.

With validation enabled, the relay agent applies stricter rules to variations in the Option 82 field(s) of incoming server responses to determine whether to forward the response to a downstream device or to drop the response due to invalid (or missing) Option 82 information. Table 5-20, below, describes relay agent management of DHCP server responses with optional validation enabled and disabled

Table 5-20. 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

To configure DHCP Option 82 on a routing switch, enter the **dhcp-relay option 82** command.

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

append: *Configures the switch to append an Option 82 field to the client DHCP packet. If the client packet has existing Option 82 field(s) assigned by another device, 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 or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the **ip** or **mgmt-vlan** option (below).*

replace: *Configures the switch to replace existing Option 82 field(s) in an inbound client DHCP packet with an Option 82 field for the 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 or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the **ip** or **mgmt-vlan** 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 or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the **ip** or **mgmt-vlan** 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 5-129.*

[ip | mac | mgmt-vlan]

*This option specifies the remote ID suboption that the switch uses in Option 82 fields added or appended to DHCP client packets. The type of remote ID defines DHCP policy areas in the client requests sent to the DHCP server. If a remote ID suboption is not configured, then the routing switch defaults to the **mac** option. (Refer to “Option 82 Field Content” on page 5-124.)*

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.*

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

*If you enter the **dhcp-relay option 82** command without specifying either **ip** or **mac**, the MAC address of the switch on which the packet was received from the client is configured as the remote ID. For information about the Remote ID values used in the Option 82 field appended to client requests, see “Option 82 Field Content” on page 5-124.*

Example of Option 82 Configuration

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

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

The resulting effect on DHCP operation for clients X, Y, and Z is shown in table 5-21.

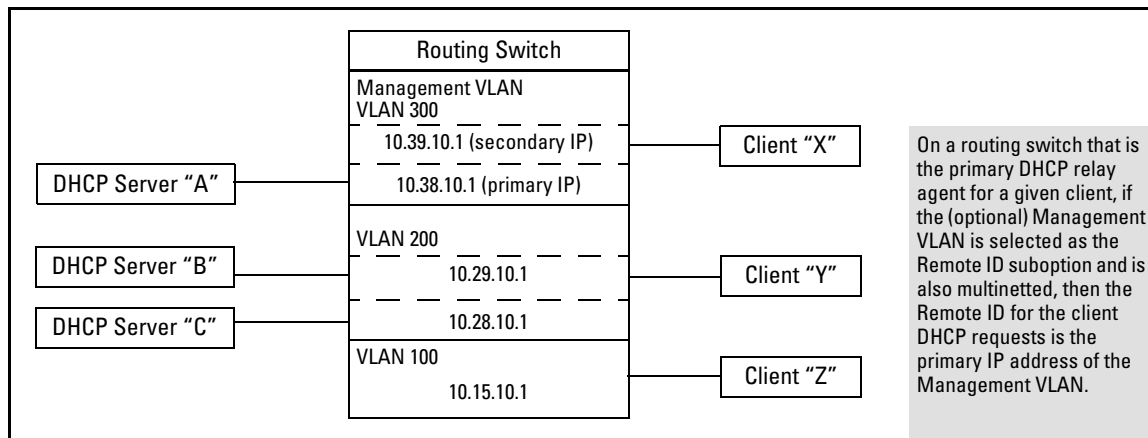


Figure 5-53. DHCP Option 82 When Using the Management VLAN as the Remote ID Suboption

Table 5-21. DHCP Operation for the Topology in Figure 5-53

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

*The IP address of the primary DHCP relay agent receiving a client request packet is automatically added to the packet, and is identified as the giaddr (*gateway interface address*). This is the IP address of the VLAN on which the request packet was received from the client. For more information, refer to RFC 2131 and RFC 3046.

Operating Notes

- 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 cannot 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 data and an error message is logged in the switch's Event Log.
- Because routing is not allowed between the Management VLAN and other VLANs, a DHCP server must be available in the Management VLAN if clients in the Management VLAN require a DHCP server.
- If the Management VLAN IP address configuration changes after **mgmt-vlan** has been configured as the remote ID suboption, the routing switch dynamically adjusts to the new IP addressing for all future DHCP requests.
- The Management VLAN and all other VLANs on the routing switch use the same MAC address.

UDP Broadcast Forwarding

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 5-142.

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 5-22:

Table 5-22. 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 routing switch. 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 5-142.

< 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)

rip: Routing Information Protocol (520)

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 routing switch 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.

```
WorkingConfig(config)# show ip forward-protocol

IP Forwarder Addresses
  [UDP Broadcast Forwarding: Disabled]
[VLAN: 1
 | IP Forward Addresses UDP Port
 |-----|
 | 15.75.11.43           37
 | 15.75.11.255         53
 | 15.75.12.255         1813
 |
 | VLAN: 2
 | IP Forward Addresses UDP Port
 |-----|
 | 15.75.12.255         1812
 |
 | VLAN: 3
 | IP Forward Addresses UDP Port
 |-----|
 | 15.75.10.155         162
 |-----|
```

Global Display Showing UDP Broadcast Forwarding Status and Configured Forwarding Addresses for Inbound UDP Broadcast Traffic for All VLANs Configured on the routing switch.

Figure 5-54. Displaying Global IP Forward-Protocol Status and Configuration

```
ProCurve(config)# show ip forward-protocol [vlan 1]

IP Forwarder Addresses
  [UDP Broadcast Forwarding: Disabled]
[IP Forward Addresses UDP Port
 |-----|
 | 15.75.11.43           37
 | 15.75.11.255         53
 | 15.75.12.255         1813
 |-----|
```

Display Showing UDP Broadcast Forwarding Status and the Configured Forwarding Addresses for inbound UDP Broadcast Traffic on VLAN 1

Figure 5-55. 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 2048 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 5-115.) 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
udp-bcast-forward: IP Routing support must be enabled first.	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.
UDP broadcast forwarder feature enabled	UDP broadcast forwarding has been globally enabled on the router. Appears in the Event Log and, if configured, in SNMP traps.
UDP broadcast forwarder feature disabled	UDP broadcast forwarding has been globally disabled on the routing switch. 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.
UDP broadcast forwarder must be disabled first.	Appears in the CLI if you attempt to disable routing while UDP forwarding is enabled on the switch.

Virtual Router Redundancy Protocol (VRRP)

Contents

Overview	6-3
Terminology	6-4
General Operation	6-5
Virtual Router (VR)	6-8
Virtual IP Address	6-8
Master Router	6-9
Owner Router	6-9
Backup Router	6-10
Virtual Router MAC Address	6-10
VRRP and ARP	6-11
General Operating Rules	6-11
Steps for Provisioning VRRP Operation	6-13
Basic Configuration Process	6-13
Example Configuration	6-15
Associating More Than One Virtual IP Address With a VR	6-17
Configuring VRRP	6-18
Enabling VRRP in the Global Configuration Context	6-18
Creating a VR and Entering the VR Context	6-19
Configuring a VR Instance on a VLAN Interface	6-20
Changing VR Advertisement Interval and Source IP Address ..	6-22
Preempt Mode on VRRP Backup Routers	6-24
Enabling or Disabling VRRP Operation on a VR	6-24
Displaying VRRP Configuration and Statistics Data	6-25
VRRP Configuration Data	6-25
Displaying the VRRP Global Configuration	6-25

Virtual Router Redundancy Protocol (VRRP)

Contents

Displaying All VR Configurations on the Router	6-25
Displaying a Specific VR Configuration	6-27
VRRP Statistics Data	6-28
Displaying Global VRRP Statistics Only	6-28
Displaying Statistics for All VRRP Instances on the Router	6-29
Displaying Statistics for All VRRP Instances in a VLAN	6-32
Displaying Statistics for a Specific VRRP Instance	6-33
Standards Compliance	6-33
Operating Notes	6-34
Event Log Messages	6-35

Overview

In many networks, edge devices are often configured to send packets to a statically configured default router. If this router becomes unavailable, the devices that use it as their first-hop router become isolated from the network.

VRRP uses dynamic failover to ensure the availability of an end node's default router. This is done by assigning the IP address used as the default route to a "virtual router", or VR. The VR includes:

- an Owner router assigned to forward traffic designated for the virtual router (If the Owner is forwarding traffic for the VR, it is the *Master* router for that VR.)
- one or more prioritized Backup routers (If a Backup is forwarding traffic for the VR, it has replaced the Owner as the Master router for that VR.)

This redundancy provides a backup for gateway IP addresses (first-hop routers) so that if a VR's Master router becomes unavailable, the traffic it supports will be transferred to a Backup router without major delays or operator intervention. This operation can eliminate single-point-of-failure problems and provide dynamic failover (and failback) support.

As long as one physical router in a VR configuration is available, the IP addresses assigned to the VR are always available, and the edge devices can send packets to these IP addresses without interruption.

Advantages to using VRRP include:

- minimizing failover time and bandwidth overhead if a primary router becomes unavailable
- minimizing service disruptions during a failover
- providing backup for a load-balanced routing solution
- addressing failover problems at the router level instead of on the network edge
- avoiding the need to make configuration changes in the end nodes if a gateway router fails
- eliminating the need for router discovery protocols to support failover operation

Terminology

Backup: A router configured in a VR as a Backup to the Owner configured for the same VR. There must be a minimum of one Backup in a VR to support VRRP operation if the Owner fails. Every backup is created with a configurable priority (default: 100) that determines the precedence for becoming the Master of the VR if the Owner or another Backup operating as the Master becomes unavailable.

Master: The Owner or Backup router that is currently the physical forwarding agent for routed traffic using the VR as a gateway. There can be only one router operating as the Master for a network or (in the case of a multinetted VLAN) a subnet. If the router configured as the Owner for a VR is available to the network, it will also be the Master. If the Owner fails or loses availability to the network, the highest-priority Backup becomes the Master.

Owner: The router configured in a VR to “own” the “virtual” IP address associated with the VR. (The virtual IP address for the VR must be configured as a real IP address on the VLAN on which the VR is configured. The Owner is automatically configured with the highest VRRP router priority in the VR (255) and operates as the Master router for the VR unless it becomes unavailable to the network.

VR (Virtual Router): Consists of one Owner router and one or more Backup routers, all of which belong to the same network or (in the case of a multinetted VLAN, the same subnet). The Owner is the router that owns the IP address(es) associated with the VR. The VR has one virtual IP address (or, in the case of a multinetted VLAN, multiple, virtual IP addresses) that corresponds to a real IP address on the Owner, and is assigned an identification number termed the VRID.

VRID: The identifier for a specific VR configured on a specific VLAN interface. On a given router, a VRID can be used for only one VR in a given VLAN, but can be used again for a different VR in a different VLAN.

General Operation

VRRP supports router redundancy through a prioritized election process among routers configured as members of the same virtual router (VR).

On a given VLAN, a VR includes two or more member routers configured with a virtual IP address that is also configured as a real IP address on one of the routers, plus a virtual router MAC address. The router that owns the IP address is configured to operate as the Owner of the VR for traffic-forwarding purposes, and by default has the highest VRRP priority in the VR. The other router(s) in the VR have a lower priority and are configured to operate as Backups in case the Owner router becomes unavailable.

The Owner normally operates as the Master for a VR. But if it becomes unavailable, then a failover to a Backup router belonging to the same VR occurs, and this Backup becomes the current Master. If the Owner recovers, a failback occurs, and “Master” status reverts to the Owner. (Note that using more than one Backup provides additional redundancy, meaning that if both the Owner and the highest-priority Backup fail, then another, lower-priority Backup can take over as Master.)

Note

- The virtual IP address used by all VRRP routers in a VR instance is a real IP address that is also configured on the applicable VLAN interface on the VR’s Owner router.
 - The same MAC and virtual IP addresses are included in the VRRP configuration for the Owner and all Backup routers belonging to the same VR, and are used as the source addresses for all traffic forwarded by the VR.
-

Figure 6-1 illustrates a virtual router on VLAN 100 supported by Router 1 (R1) and Router 2 (R2).

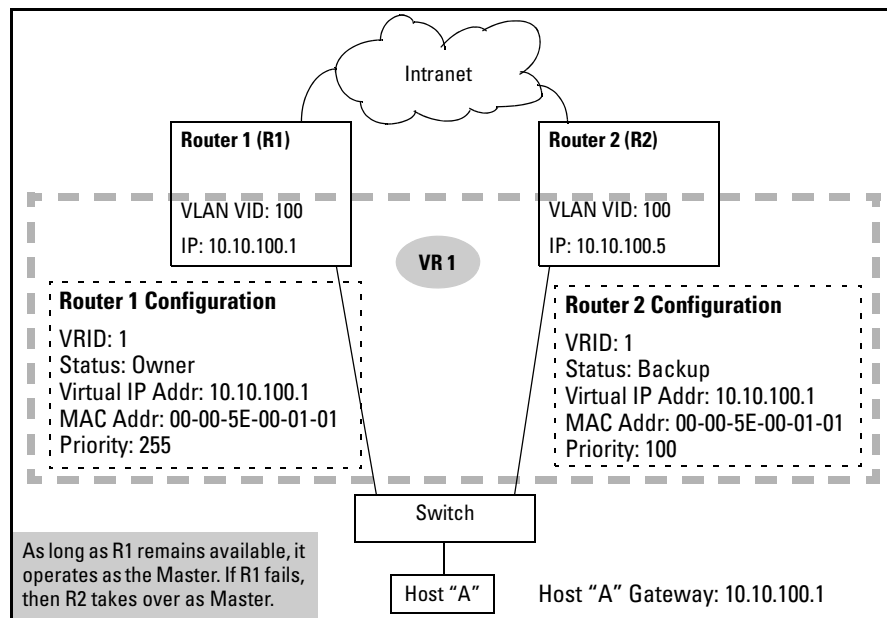


Figure 6-1. Example of Using VRRP To Provide Redundant Network Access

VR Parameter	Router 1 VR Configuration	Router 2 VR Configuration	Operation
VRID (Virtual Router ID)	1	1	All routers in the same VR have the same VRID.
Status	Owner	Backup	One Owner and one or more Backups are allowed in a given VR.
Virtual IP Address	10.10.100.1	10.10.100.1	The IP address configured for VLAN 100 in R1 (the Owner) is also configured as the Virtual IP Address for VRRP in both R1 and R2.
VR Source MAC Address	00-00-5E-00-01-01		For any VR in any VLAN, this is always defined as 00-00-5E-00-01-< VRID >, and is not configurable.
Priority	255 (Default)	100 (Default)	The router configured as Owner in any VR is automatically assigned the highest priority (255). Backup routers are assigned a default priority of 100, which can be reconfigured.

In figure 6-1:

1. Host “A” uses 10.10.100.1 as its next-hop gateway out of the subnet, as represented by the virtual router (VR 1).
 - Router 1 (the configured Owner) advertises itself as the Master in the VR supporting the gateway and:
 - “owns” the VR’s (virtual) IP address
 - transmits ARP responses that associate the VR’s virtual IP address with the (shared) source MAC address for VR 1.
 - During normal operation, Router 1 forwards the routed traffic for host “A”.
2. If Router 1 fails or otherwise becomes unavailable:
 - a. Router 1 advertisements of its Master status for VR 1 fail to reach Router 2 (which is the only configured backup).
 - b. After the time-out period for receiving Master advertisements expires on Router 2, the VR initiates a failover to Router 2 and it becomes the new Master of the VR.
 - c. Router 2 advertises itself as the Master of the VR supporting the gateway and:
 - takes control of the VR’s (virtual) IP address
 - begins transmitting ARP responses that associate the VR’s virtual IP address with the (shared) source MAC address for VR 1
 - d. Host “A” routed traffic then moves through Router 2.
3. If Router 1 again becomes available:
 - a. Router 1 resumes advertising itself as the Master for the VR and sends ARP responses that associate the VR’s virtual IP address with the (shared) source MAC address for VR 1.
 - b. Router 2 receives the advertisement from Router 1 and ceases to operate as the VR’s Master, and halts further transmission of its own VRRP advertisements and ARP responses related to VR 1.
 - c. The VR executes a failback to Router 1 as Master, and Host “A” traffic again moves through Router 1.

Virtual Router (VR)

A Virtual Router (VR) instance consists of one Owner router and one or more Backup routers belonging to the same network. Any VR instance exists within a specific VLAN, and all members of a given VR must belong to the same subnet. In a multinetted VLAN, multiple VRs can be configured. The Owner operates as the VR's Master unless it becomes unavailable, in which case the highest-priority backup becomes the VR's Master.

A VR includes the following:

- a virtual router identification (*VRID*) configured on all VRRP routers in the same network or, in the case of a multinetted VLAN, on all routers in the same subnet
- the same virtual IP address configured on each instance of the same VR
- a status of either Owner or Backup configured on each instance of the same VR (On a given VR there can be one Owner and One or more Backups.)
- a priority level configured on each instance of the VR (On the Owner router the highest priority setting, 255, is automatically fixed. On Backups, the default priority setting is 100 and is configurable.)
- a VR MAC address (not configurable)

Where a VLAN is configured with only one network (IP address), one VR is allowed in that VLAN. In a multinetted VLAN, there can be one VR per subnet, with a maximum of 32 VRs in any combination of Masters and Backups.

Note

All routers in a given VR must belong to the same network (or subnet, in the case of a multinetted VLAN).

Virtual IP Address

The virtual IP address associated with a VR must be a real IP address already configured in the associated VLAN interface on the Owner router in the VR. Also, the Owner and all other (Backup) routers belonging to the VR have this IP address configured in their VRID contexts as the *virtual IP address*. In figure 6-1 on page 6-6, 10.10.100.1 is a real IP address configured on VLAN 100 in Router 1, and is the virtual IP address associated with VR 1.

Note that if the configured Owner in a VR becomes unavailable, then it is no longer the Master for the VR and a Backup router in the VR is elected to assume the role of Master, as described under “Backup Router” on page 6-10.

A subnetted VLAN allows multiple, virtual IP addresses. However, if there are 32 or fewer IP addresses in a VLAN interface and you want VRRP support on multiple subnets, then the recommended approach is to configure a separate VR instance for each IP address in the VLAN. In cases where VRRP support is needed for more than 32 IP addresses in the same VLAN, refer to “Associating More Than One Virtual IP Address With a VR” on page 6-17.

Master Router

The current Master router in a VR operates as the “real”, or physical gateway router for the network or subnet for which a virtual IP address is configured .

Control of Master Selection. Selection of the Master is controlled by the VRRP priority value configured in the VRID context of each router in the VR. The router configured as the Owner in the VR is automatically assigned the highest VRRP priority (255) and, as long as it remains available, operates as the Master router for the VR. (The other routers belonging to the VR as Backups are assigned the default priority value (100) and can be reconfigured to any priority value between 1 and 254, inclusive.) If the current Master becomes unavailable, the protocol uses the priority values configured on the other, available routers in the VR to select another router in the VR to take over the Master function.

Function of the VRRP Advertisement. The current Master router sends periodic advertisements to inform the other router(s) in the VR of its operational status. If the backup VR(s) fail to receive a Master advertisement within the timeout interval, the current Master is assumed to be unavailable and a new Master is elected from the existing Backups. The timeout interval for a VR is three times the advertisement interval configured on the VR(s) in the network or subnet. In the default VRRP configuration, the advertisement interval is one second and the resulting timeout interval is three seconds.

Note

All VRRP routers belonging to the same VR must be configured with the same advertisement interval. As required in RFC 3768, if a locally configured advertisement interval does not match the interval received in an inbound VRRP packet, then the VR drops that packet.

Owner Router

An Owner router for a VR is the default Master router for the VR, and operates as the Owner for all subnets included in the VR. As mentioned earlier, the VRRP priority on an Owner router is always 255 (the highest).

Note

On a multinetted VLAN where multiple subnets are configured in the same VR, the router must be either the Owner for all subnets in the VR or a Backup for all subnets in the VR.

Backup Router

There must be at least one Backup router. A given VR instance on a Backup router must be configured with the same *virtual IP address* as the Owner for that VR (and both routers must belong to the same network or subnet). Router 2 in figure 6-1 on page 6-6 illustrates this point.

VR Priority Operation. In a Backup router's VR configuration, the virtual router priority defaults to 100. (The priority for the configured Owner is automatically set to the highest value; 255.) In a VR where there are two or more Backup routers, the priority settings can be reconfigured to define the order in which Backups will be reassigned as Master in the event of a failover from the Owner.

Preempt Mode. Where multiple Backup routers exist in a VR, if the current Master fails and the highest-priority Backup is not available, then VRRP selects the next-highest priority Backup to operate as Master. If the highest-priority Backup later becomes available, it pre-empts the lower-priority Backup and takes over the Master function. If you don't want a Backup router to have this preemptive ability on a particular VR, you can disable this operation with the **no preempt-mode** command. (Note that Preempt Mode applies only to VRRP routers configured as Backups.) Refer to "Preempt Mode on VRRP Backup Routers" on page 6-24.

Virtual Router MAC Address

When a VR instance is configured, the protocol automatically assigns a MAC address based on the standard MAC prefix for VRRP packets, plus the VRID number (as described in RFC 3768). The first five octets form the standard MAC prefix for VRRP, and the last octet is the configured VRID. That is:

00-00-5E-00-01-< VRid >

For example, the virtual router MAC address for the VR in figure 6-1 on page 6-6 is 00-00-5E-00-01-01.

VRRP and ARP

The Master for a given VR responds to ARP requests for the virtual IP addresses with the VR's assigned MAC address. The virtual MAC address is also used as the source MAC address for the periodic advertisements sent by the current Master.

The VRRP router responds to ARP requests for non-virtual IP addresses (IP addresses on a VLAN interface that are not configured as virtual IP addresses for any VR on that VLAN) with the system MAC address.

General Operating Rules

- IP routing must be enabled on the router before enabling VRRP.
- IP must be enabled on a VLAN before creating a VR instance on the VLAN.
- virtual IP address:
 - On an Owner: The virtual IP address configured in a VR instance must match one of the IP addresses configured in the VLAN interface on which the VR is configured.
 - On a Backup: The virtual IP address configured in a VR instance cannot be a “real” IP address configured in a VLAN interface on that router.

Note

The virtual IP address configured for one VR cannot be configured on another VR.

- Before changing a router from Owner to Backup, or the reverse, the virtual IP address must be removed from the configuration.
- The priority configuration on an Owner can only be 255. The priority configuration on a Backup must be 254 or lower; the default being 100.
- advertisement intervals:
 - A VRRP router must be configured as an Owner or Backup before configuring the advertisement interval.
 - If a VRRP router has a different advertisement interval than a VRRP packet it receives, the router drops the packet. For this reason, the advertisement interval must be the same for the Owner and all Backups in the same VR.

- When a VR is active you cannot change any of the following on that VR:
 - priority
 - advertisement interval
 - preempt mode
 - virtual IP address
- A VR exists within a single VLAN interface. If the VLAN is multinetted, then a separate VR can be configured within the VLAN for each subnet. A VLAN allows up to 32 VRs and the switch allows up to 2048 VRs.
- All routers in the same VR must belong to the same network or subnet.
- The router supports the following maximums:
 - 32 VRs per VLAN in any combination of Masters and Backups
 - 2048 VRs per router
 - 32 IP addresses per VR
- Each VR uses one MAC address as described under “Virtual Router MAC Address” on page 6-10.
- If an IP address is deleted on a VLAN interface, one of the following occurs:
 - VR Owner: If the VR uses the same IP address as a virtual IP address, then that IP address is deleted from the VR.
 - VR Backup: If the VR has a virtual IP address in the same subnet as that of the deleted IP address, then that virtual IP address will be deleted from the VR.

If the deleted virtual IP address was the last virtual IP address of an active VR, then the VR will be deactivated. (For more on multiple, virtual IP addresses on a VR, refer to “Associating More Than One Virtual IP Address With a VR” on page 6-17.

Steps for Provisioning VRRP Operation

Basic Configuration Process

This process assumes the following for VRRP operation:

- VLANs on the selected routers are already configured and IP-enabled.
- IP routing is enabled
- The network topology allows multiple paths for routed traffic between edge devices.

1. Configure the Owner for VRRP operation and a VR instance.

- a. On the router intended as the Owner for a particular network or subnet, enter the global configuration context and enable VRRP.

```
router vrrp
```

- b. Enter the desired VLAN context and configure a VR instance.

```
vlan < vid >
```

```
vrrp vrid < 1 - 255 >
```

Note that this step places the CLI in the context of the specified VR.

- c. Configure the router as the Owner of the VR instance.

```
owner
```

Note that this step automatically fixes the router's priority as 255 (the highest) for this VR instance. (The Owner priority cannot change.)

- d. Configure the router's real IP address and subnet mask for the current VLAN interface as the virtual IP address for the VR instance. You can use either of the following methods:

```
virtual-ip-address < ip-mask >
```

```
virtual-ip-address/mask-bits
```

- e. Activate the Owner VR instance.

```
enable
```

- f. Inspect the configuration for the Owner VR.

```
show vrrp vlan < vid > vrid < vrid-# > config
```

Leave the Owner's advertisement interval at its default (1 second). (For more on this topic, refer to "Changing VR Advertisement Interval and Source IP Address" on page 6-22.)

Virtual Router Redundancy Protocol (VRRP)

Steps for Provisioning VRRP Operation

2. Configure a Backup for the same VR instance as for the Owner in step 1.
 - a. On another router with an interface in the same network or subnet as is the Owner (configured in step 1), enter the global configuration context and enable VRRP.

router vrrp

- b. Configure (and enter) the same VR instance as was configured for the Owner in step 1.

vlan < vid >

vrrp vrid < 1 - 255 >

- c. Configure the router as a Backup for the VR instance.

backup

Note that this step sets the Backup router's priority as 100 for this VR instance.

- d. Optional: If there is only one Backup router, or if you want the priority among backups to be determined by the lowest IP address among the Backups, leave the VR instance priority for the current backup router at the default of 100. If you want to control Backup router priority by creating a numeric hierarchy among the Backup routers in the VR, then set the priority on each accordingly.

priority < 1 - 254 >

- e. Configure the virtual IP address for the current VR. Use the same address as you used for the Owner router's instance of the VR. As mentioned earlier, you can use either of the following methods:

virtual-ip-address < ip-mask >

virtual-ip-address/mask-bits

- f. Activate the Backup VR instance.

enable

- g. Inspect the configuration for the Owner VR.

show vrrp vlan < vid > vrid < vrid-# > config

Leave the advertisement interval for Backup routers at the default (1 second). (For more on this topic, refer to "Changing VR Advertisement Interval and Source IP Address" on page 6-22.)

3. Repeat step 2 for each Backup router on the same VR.

Example Configuration

In VR 1, below, R1 is the Owner and the current Master router, and R2 is the (only) Backup in the VR. If R1 becomes unavailable, VR 1 fails over to R2.

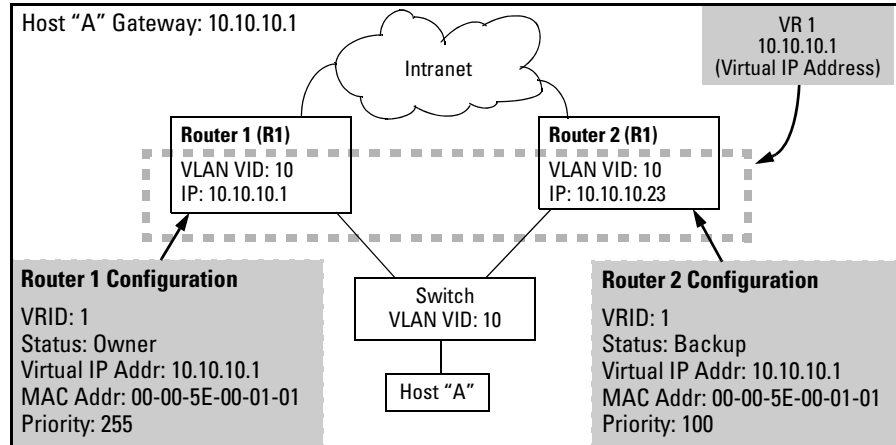


Figure 6-2. Example of a Basic VRRP Configuration

	VLAN 10 IP	VR 1 IP	Status
Router 1	10.10.10.1	10.10.10.1	Owner
Router 2	10.10.10.23	10.10.10.1	Backup

Virtual Router Redundancy Protocol (VRRP) Steps for Provisioning VRRP Operation

```

ProCurve(config)# router vrrp
ProCurve(config)# vlan 10
ProCurve(vlan-10)# vrrp vrid 1
ProCurve(vlan-10-vrid-1)# owner
ProCurve(vlan-10-vrid-1)# virtual-ip-address 10.10.10.1 255.255.255.0
ProCurve(vlan-10-vrid-1)# enable
ProCurve(vlan-10-vrid-1)# show vrrp vlan 10 vrid 1 config
VRRP Virtual Router Configuration Information
  Vlan ID : 10
  Virtual Router ID : 1

  Administrative Status [Disabled] : Enabled
  Mode [Uninitialized] : Owner
  Priority [100] : 255
  Advertisement Interval [1] : 1
  Preempt Mode [True] : True
  Primary IP Address : Lowest
  IP Address      Subnet Mask
  -----
  10.10.10.1     255.255.255.0
  
```

This router is the Owner for VR 1 in VLAN 10.

Because this router is the Owner, the priority is fixed at 255 and cannot be changed.

For the same reason, the Preempt mode cannot be changed.

Because there is only one virtual IP address configured on the VR, the source address included with advertisements from this VR is the same as the virtual IP address for the VR.

Figure 6-3. VRRP Configuration for Router 1 (R1) in Figure 6-2, Above

```

ProCurve(config)# router vrrp
ProCurve(config)# vlan 10
ProCurve(vlan-10)# vrrp vrid 1
ProCurve(vlan-10-vrid-1)# backup
ProCurve(vlan-10-vrid-1)# virtual-ip-address 10.10.10.1/24
ProCurve(vlan-10-vrid-1)# enable
ProCurve(vlan-10-vrid-1)# show vrrp vlan 10 vrid 1 config
VRRP Virtual Router Configuration Information
  Vlan ID : 10
  Virtual Router ID : 1

  Administrative Status [Disabled] : Enabled
  Mode [Uninitialized] : Backup
  Priority [100] : 100
  Advertisement Interval [1] : 1
  Preempt Mode [True] : True
  Primary IP Address : Lowest
  IP Address      Subnet Mask
  -----
  10.10.10.1     255.255.255.0
  
```

This router is a Backup in VR 1 for VLAN 10.

Because this router is a Backup, the priority is set by default to 100 and can be changed to manipulate the precedence for Backup routers in the VR.

On a Backup router, the Preempt mode can be changed. However, in a VR having only one backup, Preempt mode has no effect.

Figure 6-4. VRRP Configuration for Router 2 (R2) in Figure 6-2 on Page 6-15

Associating More Than One Virtual IP Address With a VR

This need arises if a VLAN is configured with more than 32 subnets *and* it is necessary to apply VRRP to all of these subnets.

Because a VLAN on the routers covered by this Guide supports up to 32 VRs, applying VRRP to a higher number of subnets in the VLAN requires multiple virtual IP addresses in one or more VRs.

If the Owner of a VR is associated with multiple virtual IP addresses, then the Backup router(s) belonging to the same VR must also be associated with the same set of virtual IP addresses. If the virtual IP addresses on the Owner are not also on the Backup(s), a misconfiguration exists. VRRP advertisement packets sent by the VR Master will be dropped by the VR Backup(s) on account of a mismatch among virtual IP addresses.

Configuring VRRP

Enabling VRRP in the Global Configuration Context

VRRP can be configured regardless of the global VRRP configuration status. However, enabling a VR and running VRRP requires enabling it in the global configuration context.

Syntax: [no] router vrrp

*Enables or disables VRRP operation in the global configuration context. IP routing must be enabled before enabling VRRP on the router. Disabling global VRRP halts VRRP operation on the router, but does not affect the current VRRP configuration. Enabling or disabling VRRP generates an Event Log message. To display the current global VRRP configuration, use **show vrrp config global**.*

(Default: Disabled)

Syntax: [no] router vrrp traps]

Enables or disables SNMP trap generation for the following events:

New Master — *Indicates that the sending router has transitioned to 'Master' state.*

Authentication Failure - *Indicates that a VRRP packet has been received from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type.*

Notes: *This feature assumes the **snmp-server host** command has been used to configure a a trap receiver. If a VRRP packet is received with an authentication type other than 0 (zero; that is, no authentication), then the packet is dropped. (Refer to “Operating Notes” on page 6-34.)*

(Default: Enabled)

For example, the following commands enable VRRP at the global configuration level and then display the current global VRRP configuration:

```
ProCurve(config)# router vrrp
ProCurve(config)# show vrrp config global

VRRP Global Configuration Information

VRRP Enabled      : Yes
Traps Enabled     : Yes
```

Figure 6-5. Example of Enabling and Displaying the Global VRRP Configuration

Creating a VR and Entering the VR Context

This command is used to create (or delete) a VR instance and to enter a VR context to do further configuration steps.

Syntax: [no] vrrp vrid < 1 - 255 >

Used in the VLAN interface context to create a virtual router (VR) instance and to enter the context of the new VR. It is also used to enter the context of an existing VR, and is the method used for accessing a VR for configuration purposes. You can configure up to 32 VRs on a multinetted VLAN.

The VLAN interface must be IP enabled.

For example, to create VR 1 in VLAN 10 and enter the VR context, you would execute the following command:

```
ProCurve(vlan-10)# vrrp vrid 1
ProCurve(vlan-10-vrid-1)#
```

Configuring a VR Instance on a VLAN Interface

The preceding section describes the command for creating and entering a VR context. This section describes the configuration and activation commands available in the VR context.

Assigning Owner and Backup Status. Each VRRP router must be configured as either the Owner of the VR instance or a Backup for the instance.

Syntax: < owner | backup >

Used in a VR context of a VLAN to set the router as either the Owner of the VR on that interface or as a Backup. There can be one Owner per network or subnet for a given VR. If the VLAN is multinetted and multiple subnets are configured in the same VR, the router must be either the Owner for all subnets in the VR or a Backup for all subnets in the VR. The VR instance must be disabled (the default VR state) when using this command.

(Default: None)

These commands configure and display the Owner status in VR 1 on VLAN 10:

```
ProCurve(vlan-10-vrid-1)# owner
ProCurve(vlan-10-vrid-1)# show vrrp config

VRRP Global Configuration Information
  VRRP Enabled      : Yes
  Traps Enabled     : Yes

VRRP Virtual Router Configuration Information
  Vlan ID           : 10
  Virtual Router ID : 10

  Administrative Status [Disabled] : Disabled
  Mode [Uninitialized] : Owner
  Priority [100]       : 255
  Advertisement Interval [1] : 1
  Preempt Mode [True] : True
  Primary IP Address  : Lowest

IP Address      Subnet Mask
-----
10.10.10.1     255.255.255.0
```

Executing the **owner** or **backup** command must be done in the VR context of the VLAN in which the VR exists.

Mode and Priority settings for the configured Owner on a VR.

Figure 6-6. Example of Owner Configuration on a VR

Configuring a Virtual IP address in a VR. The virtual IP address must be the same for the Owner and all Backups on the same network or subnet in a VR.

Syntax: virtual-ip-address < owner-ip-addr >/mask-length >
virtual-ip-address < owner-ip-addr > < mask >

Used in a VR context of a VLAN to assign an IP address/mask combination to a VR instance.

For an Owner: The virtual IP address must be one of the IP addresses configured on the VLAN interface for that VR.

For a Backup: The virtual IP address must match the virtual IP address for the Owner.

The Owner and the Backup(s) using a given virtual IP address must all belong to the same network or subnet. Also, the VR instance must be disabled (the default VR state) when using this command.

(Default: None)

For example, if VLAN 10 on router “A” is configured with an IP address of 10.10.10.1/24 and VR 1, and you want router “A” to operate as the Owner for this VR, then the virtual IP address of the Owner in VR 1 on router “A” is also 10.10.10.1/24. On router “B”, which will operate as a Backup for VR 1, VLAN 10 is configured (in the same network) with an IP address of 10.10.10.15/24. However, because the Backup must use the same virtual IP address as the Owner, the virtual IP address for the Backup configured on router “B” for VR 1 is also 10.10.10.1/24.




Figure 6-7. Example of Virtual IP Address Assignment for Owner and Backup

Reconfiguring the Priority for a Backup. When you configure a Backup in a VR, it is given a default priority of 100. This command is intended for use where it is necessary to establish a precedence among the Backup routers on the same network or subnet in a given VR.

Syntax: priority < 1 - 254 >

Used in a VR context of a VLAN where the router is configured as a Backup. This command changes the Backup's priority and is used to establish the precedence of a Backup where there are multiple Backups belonging to the same network or subnet. Also, the VR instance must be disabled (the default VR state) when using this command.

Note: An Owner is automatically assigned the highest priority, 255, which cannot be changed unless the Owner status is reconfigured to Backup.

(Range: 1 - 254, where 1 is the lowest precedence; Default: 100)

Changing VR Advertisement Interval and Source IP Address

The advertisement interval is used in one of two ways, depending on whether a VRRP router is operating as a Master or a Backup.

Syntax: advertise-interval < 1 - 255 >

- *When a VRRP router is operating as Master, this value specifies the interval at which the router sends an advertisement notifying the other VRRP routers on the network or subnet that a Master is active.*
- *When a VRRP router is operating as a Backup, it uses this value to calculate a timeout interval (3 x advt-interval).*

The VR instance must be disabled (the default VR state) when using this command.

(Range: 1 - 255 seconds; Default: 1 second)

For information on advertisements and advertisement intervals, see "Function of the VRRP Advertisement" on page 6-9

Note: All VRRP routers belonging to the same VR must be configured with the same advertisement interval. As required in RFC 3768, if a locally configured advertisement interval does not match the interval received in an inbound VRRP packet, then the VR drops that packet.

Syntax: primary-ip-address < ip-address | lowest >

*Specifies the virtual IP address to designate as the source for VRRP advertisements from the VR. If there is only one virtual IP address configured on the VR, the default setting (**lowest**) is sufficient. Where there are multiple virtual IP addresses in the same VR and you want to designate an advertisement source other than the lowest IP Address, use this command.*

For an Owner VR, the primary IP address must be one of the virtual IP addresses configured on the VR. For a Backup VR, the primary IP address must be in the same subnet as one of the virtual IP addresses configured on the VR. In addition, the primary IP address for a Backup VR must be one of the IP addresses configured on the VLAN on which the VR is configured.

The VR instance must be disabled (the default VR state) when using this command.

(Default: lowest)

Note: *It is common in VRRP applications to have only one virtual IP address per VR. In such cases, the protocol uses that address as the source IP address for VRRP advertisements, and it is not necessary to specify an address.*

Preempt Mode on VRRP Backup Routers

This command applies to VRRP Backup routers only, and is used to minimize network disruption due to unnecessary preemption of the Master operation among Backup routers.

Syntax: [no] preempt-mode

*Disables or re-enables Preempt mode. In the default mode, a Backup router coming up with a higher priority than another Backup that is currently operating as Master will take over the Master function. Using the **no** form of the command disables this operation, thus preventing the higher-priority Backup from taking over the Master operation from a lower-priority Backup. This command does not prevent an Owner router from resuming the Master function after recovering from being unavailable. Also, the VR instance must be disabled (the default VR state) when using this command. For more on Preempt mode, refer to “Preempt Mode” on page 6-10.*

(Default: Enabled)

Enabling or Disabling VRRP Operation on a VR

After configuring a new VR or changing the configuration on an existing VR, you must use this command to enable the VR to operate.

Syntax: [no] enable

*Enabling or disabling a VR enables or disables dynamic VRRP operation on that VR. Also, it is necessary to disable a VR before changing its configuration. Note that VRRP must be enabled (using the **router vrrp** command) in the global configuration context before enabling a VR. (Disabling a VR can be done regardless of the current, global VRRP configuration.)*

(Default: Disabled)

Displaying VRRP Configuration and Statistics Data

VRRP Configuration Data

Displaying the VRRP Global Configuration

Syntax: show vrrp config global

This command displays the configuration state for the global VRRP configuration and VRRP trap generation.

For example:

```
ProCurve(config)# show vrrp config global
VRRP Global Configuration Information
VRRP Enabled      : No
Traps Enabled     : Yes
```

Figure 6-8. Example Output Showing the Default Global VRRP Configuration

Displaying All VR Configurations on the Router

Syntax: show vrrp config

This command displays the configuration for the global VRRP configuration and all VRs configured on the router.

For example, the following figures lists output indicating two Owner VRs configured on the router:

Virtual Router Redundancy Protocol (VRRP)
Displaying VRRP Configuration and Statistics Data

```
ProCurve(config)# show vrrp config

VRRP Global Configuration Information

VRRP Enabled      : Yes
Traps Enabled     : Yes

VRRP Virtual Router Configuration Information

Vlan ID : 10
Virtual Router ID : 10

Administrative Status [Disabled] : Disabled
Mode [Uninitialized] : Owner
Priority [100] : 255
Advertisement Interval [1] : 1
Preempt Mode [True] : True
Primary IP Address : Lowest



| IP Address | Subnet Mask   |
|------------|---------------|
| 10.10.10.1 | 255.255.255.0 |



VRRP Virtual Router Configuration Information

Vlan ID : 20
Virtual Router ID : 20

Administrative Status [Disabled] : Enabled
Mode [Uninitialized] : Owner
Priority [100] : 255
Advertisement Interval [1] : 1
Preempt Mode [True] : True
Primary IP Address : Lowest



| IP Address | Subnet Mask   |
|------------|---------------|
| 10.10.20.1 | 255.255.255.0 |


```

This data shows the virtual IP address(es) configured on VR 10.

This data shows the virtual IP address(es) configured on VR 20.

Figure 6-9. Example VRRP Configuration Listing with Two Owner VRs Configured

Displaying a Specific VR Configuration

Syntax: show vrrp vlan 23 vrid 10 config

Displays the configuration for a specific VR in a specific VLAN.

For example, the following command displays the configuration of a VR identified as VR 10 in VLAN 23:

```
ProCurve(config)# show vrrp vlan 23 vrid 10 config

VRRP Virtual Router Configuration Information

Vlan ID : 23
Virtual Router ID : 10

Administrative Status [Disabled] : Disabled
Mode [Uninitialized] : Owner
Priority [100] : 255
Advertisement Interval [1] : 1
Preempt Mode [True] : True
Primary IP Address : Lowest

IP Address          Subnet Mask
-----
10.10.10.1          255.255.255.0
```

Figure 6-10. Example of Displaying the Configuration for a Specific VR

VRRP Statistics Data

All command outputs shown in this section assume that VRRP is enabled at the global configuration level. If global VRRP is disabled, these commands produce the following output:

```
VRRP Global Statistics Information

VRRP Enabled           : No
```

Figure 6-11. Statistics Command Output If Global VRRP Is Disabled

Displaying Global VRRP Statistics Only

Syntax: show vrrp statistics global

Displays the global VRRP statistics for the router.

- *VRRP Enabled*
- *Protocol Version: 2*
- *Invalid VRID Pkts Rx: VRRP packets received for a VRID that is not configured on the specific VLAN of the VRRP router.*
- *Checksum Error Pkts Rx: VRRP packets received with a bad checksum*
- *Bad Version Pkts Rx: VRRP advertisement packets received with a version number other than 2.*

```
ProCurve(config)# show vrrp statistics global

VRRP Global Statistics Information

VRRP Enabled           : Yes
Protocol Version       : 2
Invalid VRID Pkts Rx   : 0
Checksum Error Pkts Rx : 0
Bad Version Pkts Rx    : 0
```

Figure 6-12. Example of a Global VRRP Statistics Output

Displaying Statistics for All VRRP Instances on the Router

Syntax: show vrrp [statistics]

Displays the following VRRP statistics:

- *global VRRP statistics for the router*
- *VRRP statistics for all VRs configured on the router:*
 - **State:** *Indicates whether the router is a Backup or the current Master of the VR.*
 - **Uptime:** *The amount of time the router has been up since the last reboot.*
 - **Virtual MAC Address:** *The virtual MAC address for the VR instance.*
 - **Master's IP Address:** *The IP address used as the source IP address in the last advertisement packet received from the VR Master. If this VR is the Master, then this is the primary IP address of the VR. If the VR is disabled, this value appears as 0.0.0.0.*
 - **Associated IP Address Count:** *Number of virtual IP addresses.*
 - **Advertise Packets Rx:** *The number of VRRP Master advertisements the VR has received from other VRRP routers since the last reboot.*
 - **Zero Priority Tx:** *The number of VRRP advertisement packets received with the priority field set to 0 (zero).*
 - **Bad Length Pkts:** *The number of VRRP packets received with missing fields of information.*
 - **Mismatched Interval Pkts:** *The number of VRRP packets received from other routers (since the last reboot) with an advertisement interval that is different from the interval configured on the current VR. (Note that VRRP packets received with an interval mismatch are dropped.)*
 - **Mismatched IP TTL Pkts:** *The number of VRRP packets received with the IP TTL field not set to 255. Such packets are dropped.*
 - **Become Master:** *The number of times the VR has become the Master since the last reboot.*

— Continued —

— Continued from Previous Page —

- **Zero Priority Tx:** The number of VRRP advertisement packets sent with the priority field set to **0** (zero).
- **Bad Type Pkts:** The number of VRRP packets received with packet type not equal to 1 (that is, not an advertisement packet.)
- **Mismatched Addr List Pkts:** The number of VRRP packets received wherein the list of virtual IP addresses doesn't match the locally configured virtual IP addresses for a VR.
- **Mismatched Auth Type Pkts:** The number of VRRP packets received with the authentication type not equal to 0 (zero, which is no authentication).

Note that **show vrrp** and **show vrrp statistics** give the same output.

For example, the following output shows the VRRP statistics on a router having one VR (VR 1 in VLAN 10) configured.

```
ProCurve(config)# show vrrp

VRRP Global Statistics Information

VRRP Enabled           : Yes
Protocol Version       : 2
Invalid VRID Pkts Rx   : 0
Checksum Error Pkts Rx : 0
Bad Version Pkts Rx    : 0

VRRP Virtual Router Statistics Information

Vlan ID                 : 10
Virtual Router ID       : 1
State                   : Master
Up Time                 : 31 mins
Virtual MAC Address     : 00005e-000101
Master's IP Address     : 10.10.10.2
Associated IP Addr Count : 1
Advertise Pkts Rx      : 1213      Become Master           : 2
Zero Priority Rx        : 0          Zero Priority Tx         : 0
Bad Length Pkts        : 0          Bad Type Pkts           : 0
Mismatched Interval Pkts : 0      Mismatched Addr List Pkts : 0
Mismatched IP TTL Pkts  : 0      Mismatched Auth Type Pkts : 0
```

Figure 6-13. Output for Show VRRP Command Includes Global and VR Statistics

Displaying Statistics for All VRRP Instances in a VLAN

Syntax: show vrrp vlan < vid > [statistics]

This command displays the VRRP statistics for all VRs configured on the specified VLAN.

*The actual statistics data per VR is the same as for the **show vrrp [statistics]** command (pages 6-29 and 6-31).*

*Note that **show vrrp vlan < vid >** and **show vrrp vlan < vid > statistics** produce the same output.*

In the following example, there is one VR configured in VLAN 10.

```
ProCurve(config)# show vrrp vlan 10

VRRP Virtual Router Statistics Information

Vlan ID                : 10
Virtual Router ID      : 10
State                  : Master
Up Time                : 6 mins
Virtual MAC Address    : 00005e-00010a
Master's IP Address    : 10.10.10.1
Associated IP Addr Count : 1
Advertise Pkts Rx      : 1          Become Master           : 1
Zero Priority Rx       : 0          Zero Priority Tx         : 0
Bad Length Pkts       : 0          Bad Type Pkts           : 0
Mismatched Interval Pkts : 0      Mismatched Addr List Pkts : 0
Mismatched IP TTL Pkts : 0          Mismatched Auth Type Pkts : 0
```

Figure 6-14. Example of Displaying Statistics for All VRs in a VLAN

Displaying Statistics for a Specific VRRP Instance

Syntax: show vrrp vlan < vid > vrid < 1 - 255 > [statistics]

This command displays the VRRP statistics for a specific VR configured on a specific VLAN.

*The actual statistics data per VR is the same as for the **show vrrp [statistics]** command (pages 6-29 and 6-31).*

*Note that **show vrrp vlan < vid > vrid < 1 - 255 >** and **show vrrp vlan < vid > vrid < 1 - 255 > statistics** produce the same output.*

Standards Compliance

VRRP on the switches supported by this Guide includes the following:

- Complies with RFC 3768 Virtual Router Redundancy Protocol (VRRP), except for maximum number of VRs per VLAN, which is 32 on the routers covered by this Guide.
- Compatible with ProCurve Series 9300m routers, the ProCurve 9408sl router, and the ProCurve Series 8100fl switches. (VRRP on these devices is based on RFC 2338.)
- Complies with RFC 2787-- Definitions of Managed Objects for VRRP, except for support for authentication-related values.
- Applies to use on IPv4 routers.

Operating Notes

- **VRRP Advertisements Not Reaching the Backup(s):** If a Master is forwarding traffic properly, but its Backup(s) are prevented from receiving the Master's VRRP advertisements, then both routers will operate in the Master mode for the VR. If this occurs, traffic for the applicable gateway will continuously alternate between routers (sometimes termed "flapping").
- **Deleting an IP Address Used To Support a VR:** Refer to "General Operating Rules" on page 6-11.
- **VR Limits:** A VLAN allows up to 32 VRs, and a VR allows up to 32 IP addresses. This means that one VR can support up to 32 subnets. This capacity enables use of VRRP on all subnets in a VLAN that has more than 32 subnets.
- **IPv4:** The routers covered by this Guide support IPv4 IP addressing for VRRP applications.
- **Authentication Type:** As per RFC 3768, the authentication type for VRRP packets inbound on the router is 0 (zero; that is, "no authentication"). Packets with other authentication types are dropped, and authentication type is not supported in the VRRP MIB. If you are coordinating the use of VRRP on the routers covered by this manual with another vendor's implementation based on an older RFC, then you must set the authentication type to 0 (zero) on the other vendor's device.
- **Proxy-ARP requests and MAC addresses:** The following table shows which MAC address is returned in response to a proxy-ARP request.

Configured as:	Administratively:	Returns:
Owner	Enabled	VRRP MAC address
Owner	Disabled	Default VLAN MAC address
Backup	Enabled, in Master state	VRRP MAC address
Backup	Enabled, not in Master state	VRRP router does not respond to proxy-ARP request.
Backup	Disabled	Default VLAN MAC address

Event Log Messages

Message	Meaning
Failure to send out pkt for vrid <vrid-#>, vid <vid-#>	A VRRP packet could not be sent out for the indicated VR on the specific VLAN due to any system-dependent problem. If packets could not be sent out, the expected protocol operation may be hampered.
No VR with vrid <vrid-#> found on vid <vid-#>	Indicates a VRRP packet received for a VR that does not exist on the VLAN. This can indicate asymmetric configuration of VRs across VRRP routers.
Pkt recd on a non-VRRP Vlan with vid <vid-#>	A VRRP packet was received on a VLAN that does not have any VRs. This could possibly be a result of misconfiguration of VRs on VLANs.
Pkt recd with version number <ver-#>, expected <ver-#>	A VRRP packet was received with a wrong version number.
Vrid <vrid-#> on Vid <vid-#> has taken backup IP ctrl	The Owner of a VR is not available and a Backup has taken Master control of the VR.
Vrid <vrid-#> on Vid <vid-#> has taken owner IP ctrl	The Owner of a VR has taken Master control of the VR, either following a reboot or a failback from a Backup serving as Master.
Vrid <vrid-#> on Vid <vid-#> lost backup IP ctrl	The indicated VR has been preempted by either the Owner or a higher-priority Master.
Vrid <vrid-#>, Vid <vid-#> IP addr is duplicated on the network	The virtual IP address owned by the indicated VR on the indicated VLAN is duplicated on the network.
Vrid <vrid-#>, Vid <vid-#> recd pkt from a duplicate master	A VRRP packet was received from a duplicate master VR by the indicated VR on the indicated VLAN.
Vrid <vrid-#>, Vid <vid-#> recd pkt with advt int mismatch	The indicated VR on the indicated VLAN has received a VRRP Master message carrying a different advertisement interval than is configured on the receiving VR and has dropped the packet.
Vrid <vrid-#>, Vid <vid-#> recd pkt with auth type mismatch	Indicates the VR has received a packet with the authentication type set to 1 or 2. These are generally valid authentication types, but are not required by RFC 3768. Thus, the software supports only an authentication type of 0 (zero), and VRRP packets with 1 or 2 for authentication type are dropped. Refer to "Authentication Type" under "Operating Notes" on page 6-34.

— Continued —

Virtual Router Redundancy Protocol (VRRP)
Event Log Messages

Message	Meaning
<i>— Continued from Previous Page —</i>	
Vrid <vrid-#>, Vid <vid-#> recd pkt with bad IP-TTL	A VRRP packet was received by the indicated VR on the indicated VLAN with an IP TTL value not equal to 255.
Vrid <vrid-#>, Vid <vid-#> recd pkt with checksum error	The indicated VR on the indicated VLAN has received a VRRP advertisement packet with a checksum error. The VR has therefore dropped that packet.
Vrid <vrid-#>, Vid <vid-#> recd pkt with invalid auth type	Indicates the VR has received a VRRP packet with an authentication type set to a value other than the 0, 1, or 2 (allowed by RFC 3768) and has dropped the packet.
Vrid <vrid-#>, Vid <vid-#> recd pkt with IP address mismatch	A VRRP packet was received by the indicated VR on the indicated VLAN with virtual IP address(es) that did not match the virtual IP addresses configured on the receiver VR.
Vrid <vrid-#>, Vid <vid-#> recd pkt with invalid type	A VRRP packet was received by the indicated VR on the indicated VLAN with the packet type not equal to 1.
VRRP has been disabled on this router	VRRP was disabled at the global config level.
VRRP has been enabled on this router	VRRP was enabled at the global config level.

Index

Symbols

A

ABR

- definition ... 5-44
- OSPF ... 5-44

ACL

- operation with PIM ... 3-36

address

- IP ... 5-13

administrative distance, OSPF ... 5-73

advertisement, OSPF ... 5-43

- area ... 5-59
- retransmit interval ... 5-78
- retransmit interval in virtual link ... 5-86

area range, OSPF

- configuring ... 5-70

area, OSPF

- assigning VLAN to ... 5-64, 5-65
- configuring ... 5-59
- definition ... 5-47
- displaying area information ... 5-92

ARP

- cache ... 5-8
- cache table ... 5-8
- configuring parameters ... 5-15
- how it works ... 5-15
- proxy ... 5-17

assigning

- IP address ... 5-13

authentication

- OSPF
- MD5 ... 5-81

auto port setting ... 2-5

Autonomous system, OSPF ... 5-47

B

blocked port

- from IGMP operation ... 2-5

bootstrap message, defined ... 4-6

bootstrap router ... 4-6

broadcast forwarding ... 5-116

broadcast traffic

- enabling forwarding of directed ... 5-18

BSM, PIM-SM ... 4-6

BSR

- change priority setting ... 4-36
- configuration ... 4-13
- configuring a candidate ... 4-35
- display data ... 4-59
- election ... 4-13
- enable or disable operation ... 4-35
- fault recovery ... 4-14
- non-default settings ... 4-60
- operation ... 4-13
- PIM-SM domain ... 4-6

C

caches

- ARP ... 5-8
- IP forwarding ... 5-9

Candidate Rendezvous Point

- defined ... 4-6
- See C-RP ... 4-6

chain, key management ... 5-96

CIDR ... 5-13

circuit ID ... 5-122, 5-125

configuration

- ARP parameters ... 5-15
- default route ... 5-25
- DHCP Relay ... 5-115
- IP routing parameters ... 5-13
- OSPF ... 5-40
- RIP ... 5-26, 5-28
- enabling RIP globally ... 5-28
- router loop prevention ... 5-33
- router ID ... 5-13
- static IP routes ... 5-21, 5-23

C-RP

- add multicast group ... 4-40
- change hold time ... 4-40
- configuring operation ... 4-38
- defined ... 4-6
- display config ... 4-63

- display status ... 4-63
- displaying current set ... 4-61
- election priority ... 4-41
- enabling or disabling ... 4-40
- multicast groups ... 4-38
- specify VLAN interface ... 4-38
- with PIM-SM router ... 4-6

D

- default route ... 5-25
- Depending ... 2-14
- Designated Router
 - defined ... 4-7
 - election criteria ... 4-12
 - in VLAN ... 4-12
- DHCP
 - hop count, disabling ... 5-117
 - hop count, displaying ... 5-120
- DHCP Relay
 - broadcast forwarding ... 5-116
 - configuration ... 5-115
 - enabling ... 5-116
 - helper address ... 5-117
 - hop count in requests ... 5-117
 - minimum requirements ... 5-116
 - Option 82
 - circuit ID ... 5-122, 5-125
 - packet forwarding ... 5-115
 - See also* Option 82.
 - verifying configuration ... 5-118
- directed broadcasts ... 5-18
- DR (designated router)
 - defined ... 4-7
 - election criteria ... 4-12
 - OSPF
 - election ... 5-45
 - See* OSPF.

E

- edge device ... 6-3
- edge router, defined ... 4-7
- Enabling ... 6-24
- enabling OSPF ... 5-56
- enabling RIP, globally ... 5-28
- event log
 - See* log.

- Exclude Source
 - See* IGMP.
- external LSA
 - displaying ... 5-93

F

- failover, VRRP ... 6-3
- filters
 - effect of IGMP ... 2-28
 - maximum allowed ... 2-6
 - See* OSPF and RIP.
- flow, defined ... 4-7
- forwarding
 - directed broadcasts ... 5-18
 - parameters, IP routing
 - configuring ... 5-17
- forwarding port, IGMP ... 2-5

H

- helper address for DHCP Relay ... 5-117
- hop count in DHCP requests
 - disabling ... 5-117
 - displaying configuration ... 5-120

I

- IANA ... 5-142
- ICMP
 - configuring ... 5-19
 - disabling messages ... 5-19
- IGMP
 - benefits ... 2-3
 - configure per VLAN ... 2-5
 - effect on filters ... 2-28
 - Exclude Source ... 2-12
 - Fast Leave ... 2-14
 - high-priority disabled with PIM ... 3-36
 - 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 ... 2-25
 - forwarding ... 2-18
 - forwarding commands ... 2-21
 - show command ... 2-23
 - vlan context command ... 2-22
 - query ... 2-12
 - report ... 2-12
 - status ... 2-13
 - traffic ... 2-5
 - Version 3 ... 2-12
 - IGP ... 5-41
 - Include Source
 - See* IGMP.
 - IP address
 - assigning ... 5-13
 - CIDR notation ... 5-13
 - IP address, multiple ... 6-17
 - IP address, virtual ... 6-5
 - IP forwarding cache ... 5-9
 - IP global parameters ... 5-10
 - IP interface parameters ... 5-12
 - IP route exchange protocols ... 5-10
 - IP route table ... 5-8
 - IP routing
 - ARP cache table ... 5-8
 - AS ... 5-41
 - backup designated router ... 5-41
 - BDR ... 5-41
 - chain ... 5-96
 - changing ARP parameters ... 5-15
 - changing router ID ... 5-13
 - configuring static routes ... 5-21
 - default route ... 5-25
 - designated router ... 5-41
 - DHCP Relay configuration
 - See also* DHCP relay.
 - directed broadcasts ... 5-18
 - DR ... 5-41
 - forwarding cache ... 5-9
 - forwarding parameters ... 5-17
 - global parameters ... 5-10
 - helper address ... 5-117
 - helper address, UDP ... 5-12
 - interface parameters ... 5-12
 - interior gateway protocol ... 5-41
 - IP static routes
 - administrative distance ... 5-24
 - blackhole ... 5-21, 5-24
 - configuration ... 5-24
 - default route ... 5-11, 5-22
 - default route, configuring ... 5-25
 - display ... 5-25
 - maximum ... 5-6
 - null interface ... 5-22
 - null route ... 5-21, 5-23
 - one per destination ... 5-21
 - reject ... 5-24
 - VLAN state ... 5-23
 - loopback interface
 - See* OSPF.
 - null routes ... 5-21
 - overview ... 5-6
 - parameter configuring ... 5-13
 - Proxy ARP, enabling ... 5-17
 - route exchange protocols ... 5-10
 - router ID ... 5-45, 5-46, 5-47
 - routing table ... 5-8
 - See also* DHCP, ICMP, IRDP, OSPF, and UDP
 - broadcast forwarding.
 - static route configuration ... 5-23
 - static route parameters ... 5-22
 - static route types ... 5-21
 - static routes
 - administrative distance ... 5-22
 - discard traffic ... 5-23
 - discard, ICMP notification ... 5-23
 - metric ... 5-22
 - tables and caches ... 5-7
 - terminology ... 5-41
 - type-7 LSA ... 5-41
 - VLAN interface ... 5-7
- IRDP**
- configuring ... 5-112
 - displaying information ... 5-114
 - enabling globally ... 5-113
 - enabling on VLAN interface ... 5-113
- K**
- key-chain ... 5-89
- L**
- leave group

- See* IGMP.
 - log, counter ... 3-38
 - log, PIM messages
 - loopback interface ... 5-47
 - LSA
 - displaying ... 5-98
 - external
 - displaying ... 5-93
- ## M
- management VLAN
 - See* Option 82.
 - MD5 authentication, OSPF ... 5-81
 - metric
 - OSPF redistribution ... 5-70
 - multicast group
 - See* IGMP.
 - multicast source, defined ... 4-7
 - multimedia
 - See* IGMP.
 - multinetted VLANs and Option 82 ... 5-131
 - multiple relay agents ... 5-128
- ## O
- Option 82 ... 5-120
 - circuit ID ... 5-122, 5-125
 - compliance ... 5-121
 - configuring operation ... 5-131
 - field content ... 5-124
 - forwarding policy ... 5-122, 5-127
 - management VLAN ... 5-123, 5-125, 5-131
 - multinetted VLANs ... 5-131
 - multiple relay agents ... 5-128
 - operation ... 5-123
 - Option 82 field ... 5-120
 - overview ... 5-120
 - policy boundary ... 5-122
 - primary relay agent ... 5-123
 - relay agent ... 5-123
 - Relay Agent Information ... 5-120
 - remote ID ... 5-123, 5-124
 - requirements ... 5-123
 - secondary relay agent ... 5-123
 - server support ... 5-121
 - validating server response packets ... 5-129
 - OSPF
 - ABR ... 5-41, 5-42, 5-44
 - ABR, connection requirement ... 5-59
 - ABR, range configuration ... 5-53
 - administrative distance ... 5-53, 5-73
 - advertisement, blocking ... 5-72
 - advertisements ... 5-71
 - area ... 5-42, 5-47, 5-53
 - assigning VLAN to ... 5-64, 5-65
 - configuring ... 5-59
 - area border router ... 5-41, 5-44
 - area configuration ... 5-59
 - area information ... 5-92
 - area range
 - configuring ... 5-70
 - area types ... 5-47
 - ASBR ... 5-41, 5-45, 5-53
 - ASBR, advertising ... 5-51
 - ASBR, in NSSA ... 5-53
 - assigning area range ... 5-70
 - authentication ... 5-54
 - description ... 5-80, 5-88
 - interface ... 5-80
 - MD5 ... 5-81, 5-89
 - MD5, virtual link ... 5-89
 - password ... 5-80, 5-88
 - virtual link ... 5-88
 - autonomous system ... 5-41, 5-47
 - autonomous system boundary router ... 5-41, 5-45
 - backbone area ... 5-41, 5-47, 5-48
 - backbone area, configure ... 5-60
 - blocking routes ... 5-70
 - chain ... 5-96
 - changing compliance setting ... 5-57
 - changing defaults ... 5-75
 - changing port parameters ... 5-57
 - configuration rules ... 5-54
 - configuration steps ... 5-56
 - configuring ... 5-40
 - cost ... 5-53, 5-75, 5-76
 - dead-interval ... 5-53, 5-75, 5-76
 - default metric ... 5-69
 - default parameter settings ... 5-53
 - default port parameters ... 5-57
 - default route ... 5-41
 - displaying configuration and status ... 5-90
 - displaying information ... 5-90, 5-91, 5-95, 5-97
 - area ... 5-92

- external LSA ... 5-93
- LSA ... 5-98
- neighbor ... 5-101
- virtual link ... 5-105
- virtual neighbor ... 5-104
- displaying redistribution ... 5-103
- DR (designated router) ... 5-45
 - election ... 5-45
- enabling ... 5-52, 5-56
- enabling redistribution ... 5-69
- external LSA ... 5-41
- external route cost options ... 5-62
- external routes, redistribution ... 5-53
- general configuration steps ... 5-53
- general information ... 5-91
- hello-interval ... 5-53, 5-75, 5-77
- interface
 - defaults ... 5-75
- interface parameters ... 5-75
- interior router ... 5-41, 5-44
- key-chain ... 5-89
- link-state advertisement ... 5-41, 5-42
- loopback interface ... 5-47, 5-53
- loopback interface, assigning ... 5-65
- loopback interface, redistribution ... 5-66
- LSA ... 5-41, 5-42
 - displaying information ... 5-93
- LSA types ... 5-43
- LSA, external, reduction ... 5-51
- LSA, reduction ... 5-50
- MD5 authentication ... 5-81
- metric-type ... 5-62
- neighbor ... 5-82
- no-advertise ... 5-71
- normal area ... 5-42, 5-47, 5-48
- normal area, configure ... 5-60
- no-summary ... 5-62
- no-summary, effect ... 5-52
- not-so-stubby area ... 5-42
- not-so-stubby-area
 - See also* NSSA.
- NSSA ... 5-42, 5-43, 5-47, 5-49
- NSSA, configuring ... 5-61
- overview ... 5-43
- parameters, default global ... 5-54
- parameters, default interface ... 5-55
- password ... 5-80
- password, virtual link ... 5-88
- priority ... 5-53, 5-75, 5-77
- range, blocking ... 5-72
- redistribution ... 5-69
 - enabling ... 5-69
 - metric ... 5-69
 - metric type ... 5-70
- redistribution filters
 - displaying ... 5-103
- redistribution information ... 5-103
- redistribution, configuring ... 5-68
- redistribution, loopback interface ... 5-66
- restrict redistribution filters, display ... 5-103
- restrict redistribution. ... 5-67
- retransmit-interval ... 5-75, 5-78
- RFC 1583 compliance option ... 5-53
- RFC 1583, compliance setting ... 5-57
- RFC 1583, example ... 5-58
- RFC 2178 ... 5-57
- RFC 2328 ... 5-43, 5-50, 5-51, 5-57
- RFC 3101 ... 5-43, 5-50
- RFC compliance ... 5-50, 5-57
- route choice, influencing ... 5-73
- router ID ... 5-14, 5-45, 5-46, 5-47, 5-51, 5-52, 5-82, 5-98, 5-99, 5-101, 5-104, 5-105, 5-107, 5-108
- router ID, displayed ... 5-91, 5-93, 5-94
- routing table, displaying ... 5-107
- show commands ... 5-90
- stub area ... 5-42, 5-47, 5-50
- stub area, configuring ... 5-61
- summarizing routes ... 5-70
- summary link-state-advertisement ... 5-42
- topological database ... 5-42
- transit area ... 5-82
- transit area ID ... 5-85, 5-86, 5-87, 5-88
- transit-delay ... 5-75, 5-78
- traps ... 5-53, 5-74
- type-3 default summary LSA ... 5-62
- type-3 LSA ... 5-42, 5-50
- type-3 summary LSA ... 5-52, 5-62
- type-5 LSA ... 5-41, 5-42
- type-7 default external LSA ... 5-52, 5-62
- virtual link ... 5-42, 5-59
 - authentication ... 5-88
 - change settings ... 5-85
 - configuration ... 5-82
 - dead-interval ... 5-85
 - defaults ... 5-85
 - displaying information ... 5-105

- hello-interval ... 5-86
- interface parameters ... 5-85
- MD5 authentication ... 5-89
- parameters ... 5-85
- retransmit-interval ... 5-86
- transit-delay ... 5-87
- virtual neighbor
 - displaying information ... 5-104
- overview, IP routing ... 5-6

P

parameters

- IP global ... 5-10
- IP interface ... 5-12
- OSPF
 - interface ... 5-75
- See also* DHCP, ICMP, IRDP, and OSPF.

peers, RIP

- displaying information ... 5-37

PIM error message ... 3-38

PIM-DM

- age-out, multicast group entry ... 3-25
- bandwidth conservation ... 3-8
- common subnet requirement ... 3-6
- compatible draft versions ... 3-4
- configuration ... 3-11, 3-12, 3-13, 3-14, 3-21, 3-30
- configuration order ... 3-12
- configuration, general elements ... 3-9
- configuration, router ... 3-12
- default settings recommended ... 3-9
- displaying data and configuration ... 3-22
- draft versions 1 and 2 ... 3-4
- error message ... 3-38
- expiry time ... 3-26, 3-34
- extended branch ... 3-5
- features ... 3-4
- flood ... 3-6
- flood and prune ... 3-6, 3-7, 3-28
- flood and prune cycle ... 3-35
- flow ... 3-6, 3-9
- flow, bridged ... 3-37
- flow, equalizing ... 3-36, 3-38, 3-39, 3-40, 3-41
- flow, hardware ... 3-10, 3-13
- flow, multicast, limit ... 3-10, 3-38
- flow, software ... 3-10, 3-13
- flow, VLAN limit ... 3-4
- forwarding state ... 3-7

- general operation ... 3-5
- graft packets ... 3-16, 3-17
- group entry, age-out ... 3-25
- hello hold-time ... 3-15, 3-30
- hello interval, effect ... 3-15
- host ... 3-9
- IGMP required, per VLAN ... 3-9
- IGMP requirement ... 3-36
- IGMP version 1 ... 3-4
- IGMP version 2 ... 3-4
- IGMP version 3 ... 3-4
- IGMP, per VLAN ... 3-5
- IP address required ... 3-36, 3-39
- join ... 3-5, 3-6, 3-9
- limit, multicast flow ... 3-10
- log message ... 3-37, 3-38
- log message counter operation ... 3-38
- MIB support ... 3-4
- MRT ... 3-4, 3-10, 3-13, 3-31, 3-37
- MRT, explained ... 3-9
- multicast address ... 3-5, 3-10
- multicast flow, limit ... 3-10
- multicast group address
 - See* multicast address.
- multicast router, multiple ... 3-32
- Multicast Routing MIB ... 3-4
- multicast routing table
 - See* MRT. ... 3-4
- multicast routing, defined ... 3-10
- multicast server ... 3-10
- multinetted VLAN ... 3-6, 3-10, 3-19, 3-20
 - common subnet required ... 3-10, 3-15
- neighbor field, blank ... 3-27
- neighbor, PIM ... 3-10, 3-23, 3-34
- OSPF ... 3-5
- outbound VLAN limit ... 3-10
- PIM instance per VLAN ... 3-10
- prune ... 3-6, 3-10, 3-26, 3-33
- prune delay ... 3-17, 3-18
- prune state ... 3-7
- pruned branch ... 3-5
- prune-pending state ... 3-18
- pruning ... 3-7
- reverse path forwarding ... 3-5
- RFC 2932 ... 3-4
- RFC 2932 exceptions ... 3-42
- RFCs, applicable ... 3-41
- RIP ... 3-5

- route data ... 3-23
- router configuration ... 3-12
- routing protocol ... 3-5, 3-9
- routing switch 9300 ... 3-35
- RPF ... 3-5
- S/G pair ... 3-9, 3-10
- SNMP traps ... 3-13, 3-28
- source address, unicast ... 3-10
- state refresh ... 3-7, 3-8, 3-13, 3-25, 3-28, 3-31, 3-34, 3-35
- state refresh, on other routers ... 3-35
- static route ... 3-5
- subnet, common ... 3-6, 3-10
- time-to-live threshold ... 3-19, 3-24
- traps, SNMP ... 3-13, 3-28
- tree, multicast ... 3-5, 3-6
- TTL zero ... 3-38
- unicast routing ... 3-4, 3-5
- unicast source address ... 3-5
- unicast source address, server ... 3-10
- version differences ... 3-40
- VLAN support, inbound ... 3-4
- VLAN support, outbound ... 3-4
- VLAN, flow limit ... 3-4
- VLAN, multinetted ... 3-6
- VLAN, PIM instance per ... 3-10
- XRRP ... 3-4

PIM-SM

- age-out, multicast group entry ... 4-48
- border routers ... 4-12
- BSR ... 4-12, 4-13
 - candidate configuration ... 4-35
 - message interval ... 4-37
 - non-default settings ... 4-60
 - priority setting ... 4-36
 - protocol ... 4-5
- changing DR priority ... 4-33
- compatible draft versions ... 4-6
- configuration ... 4-27, 4-56
- configuring candidate-RPs ... 4-37
- Designated Router ... 4-12
- display BSR data ... 4-59
- display config ... 4-51
- display C-RP config ... 4-63
- display RP set ... 4-61
- display status ... 4-51
- displaying settings ... 4-46
- DR ... 4-7
 - priority ... 4-57
- draft versions 1 and 2 ... 4-6
- enable/disable SNMP Traps ... 4-41
- entries in routing table ... 4-52
- event log messages ... 4-66
- expire time ... 4-48, 4-57
- features ... 4-5
- flow capacity ... 4-5
- flow, defined ... 4-7
- flow, hardware ... 4-41
- flow, software ... 4-41
- flow, VLAN limit ... 4-5
- group address ... 4-47, 4-52
- group entry, age-out ... 4-48
- hello delay ... 4-31
- hello hold-time ... 4-30
- hello interval ... 4-31
- hello interval, effect ... 4-30
- IGMP link ... 4-4
- IGMP version 1 ... 4-6
- IGMP version 2 ... 4-6
- IGMP version 3 ... 4-6
- join ... 4-7
- join/prune interval ... 4-42
- lan-prune-delay ... 4-31
- list interfaces ... 4-55
- MIB support ... 4-6
- MRT ... 4-5, 4-41, 4-52
- multicast group distribution ... 4-36
- multicast router, multiple ... 4-53
- Multicast Routing MIB ... 4-6
- multicast routing protocol ... 4-48
- multicast routing table
 - See MRT. ... 4-5
- multicast source ... 4-7
- neighbor ... 4-47
- neighbor field, blank ... 4-49
- neighbor, PIM ... 4-57
- non-flooding model ... 4-9
- operating notes ... 4-65
- operation ... 4-9
- PMBR not supported ... 4-12
- propagation delay ... 4-31
- prune ... 4-7, 4-49, 4-54
 - assert ... 4-54
 - delay ... 4-32
- prune delay ... 4-33
- prune-pending state ... 4-33

- rendezvous point ... 4-7
- rendezvous point tree ... 4-7
- RFC 2932 ... 4-6
- router types ... 4-12
- RP ... 4-12
- RP mapping ... 4-5
- RPF ... 4-7
- RP-Set command ... 4-8
- RPT-bit ... 4-53
- shortest path tree ... 4-8
- show VLAN configs ... 4-55
- SNMP traps ... 4-41, 4-51
- source address ... 4-47, 4-52
- state refresh ... 4-48, 4-51, 4-57
- static rendezvous point ... 4-8
- Static-RP ... 4-12
- time-to-live threshold ... 4-50
- traps, SNMP ... 4-41, 4-51
- TTL threshold ... 4-50
- unicast routing ... 4-5
- unicast routing protocol ... 4-49
- up time ... 4-48
- using SPT controls ... 4-42
- VLAN support, inbound ... 4-5
- VLAN support, outbound ... 4-5
- VLAN, flow limit ... 4-5
- VRRP ... 4-6
- PMBR ... 4-12
- port
 - auto, IGMP ... 2-5
 - blocked, IGMP ... 2-5
 - forwarding, IGMP ... 2-5
 - state, IGMP control ... 2-5
- primary relay agent ... 5-123
- priority ... 2-5
- protocols
 - IP route exchange ... 5-10
- Proxy ARP, enabling ... 5-17
- proxy forwarding, IGMP ... 2-19
- prune, defined ... 4-7

Q

- query
 - See* IGMP.
- quick start ... 1-8

R

- redistribution ... 5-103
 - See* RIP and OSPF.
- redistribution filters
 - See* OSPF and RIP.
- relay agent ... 5-123
- remote ID ... 5-123
- Rendezvous Point Tree, defined ... 4-7
- Rendezvous Point, defined ... 4-7
- report
 - See* IGMP.
- Revers Path Forwarding, defined ... 4-7
- RFC 2178 ... 5-50
- RFC 2178 compliance
 - enabling for OSPF ... 5-57
- RFC 2328 ... 5-43, 5-50, 5-51
- RFC 2338 ... 6-33
- RFC 2787 ... 6-33
- RFC 2932 ... 3-4, 4-6
- RFC 2932 MIB exceptions ... 3-42
- RFC 3101 ... 5-43, 5-50
- RFC 3768 ... 6-9, 6-10, 6-33, 6-34
- RFCs, PIM-applicable ... 3-41
- RIP
 - changing cost of RIP routes ... 5-29
 - changing RIP type ... 5-29
 - changing route loop prevention ... 5-33
 - changing the RIP metric ... 5-29
 - configuring ... 5-26, 5-28
 - displaying configuration and status ... 5-33
 - displaying general information ... 5-34
 - displaying information ... 5-33
 - displaying interface information ... 5-35
 - displaying peer information ... 5-37
 - displaying redistribution information ... 5-38
 - displaying restrict information ... 5-39
 - enabling globally ... 5-28
 - enabling on a VLAN ... 5-29
 - enabling route redistribution ... 5-31
 - general information ... 5-34
 - global parameters ... 5-27
 - interface information ... 5-35
 - interface parameters ... 5-27
 - overview ... 5-26
 - parameters and defaults ... 5-27
 - peer information ... 5-37
 - redistribution ... 5-31
 - displaying ... 5-38

- enabling ... 5-31
- redistribution filters
 - displaying ... 5-39
- redistribution into RIP ... 5-30
- redistribution, configuring ... 5-30
- restrict filter information ... 5-39
- restrict redistribution ... 5-39
- route loop prevention, RIP configuration ... 5-33
- router ID, changing ... 5-13
- router, multicast, with IGMP ... 2-12
- routing
 - See also* DHCP, ICMP, IRDP, UDP broadcast forwarding, and OSPF.
- routing table, displaying ... 5-107
- routing, *See* IP routing.
- RP
 - defined ... 4-7
 - dynamic ... 4-6
 - with PIM-SM router ... 4-6
- RPF, defined ... 4-7
- RP-Set, defined ... 4-8
- RPT
 - traffic restricted to ... 4-11

S

- secondary relay agent ... 5-123
- setup screen ... 1-8
- Shortest Path Tree, defined ... 4-8
- SPT
 - defined ... 4-8
 - operation ... 4-10
 - PIM-SM traffic ... 4-42
- static IP routes
 - configuring ... 5-21, 5-23
 - route types ... 5-21
 - See also* IP routing.
- Static Rendezvous Point
 - defined ... 4-8
 - See* SPT
- static-RP
 - manual configuration ... 4-42
- Static-RP, defined ... 4-8
- subnet ... 2-13

T

- tables

- ARP cache ... 5-8
- IP ... 5-7
- IP route ... 5-8
- transit area
 - OSPF ... 5-82
- traps
 - OSPF ... 5-74

U

- UDP broadcast forwarding ... 5-137
 - address types ... 5-137
 - application ... 5-137
 - configure ... 5-139
 - global enable ... 5-139
 - invalid entry ... 5-138
 - IP helper address, effect ... 5-137
 - maximum entries ... 5-137
 - port-number ranges ... 5-142
 - show command ... 5-141
 - subnet address ... 5-137
 - subnet masking ... 5-138
 - UDP/TCP port number listing ... 5-142
 - unicast address ... 5-137
 - VLAN, subnetted ... 5-137

V

- virtual link
 - OSPF
 - displaying information ... 5-105
 - parameters ... 5-85
- virtual MAC address ... 6-10
- virtual neighbor
 - OSPF
 - displaying information ... 5-104
- VLAN
 - assigning OSPF area to ... 5-64, 5-65
 - IGMP configuration ... 2-5
- VLAN interface
 - description ... 5-7
 - IP routing parameters ... 5-12
 - OSPF
 - interface parameters ... 5-75
- VLAN, outbound limit ... 3-10
- VRRP
 - advantages ... 6-3
 - advertisement ... 6-7

- function ... 6-9
- interval ... 6-9, 6-11
- ARP response ... 6-7, 6-11
- authentication type ... 6-34
- Backup router ... 6-3, 6-10
 - as Master. ... 6-3
 - defined ... 6-4
 - elected as Master ... 6-8, 6-9
 - multiple ... 6-5
 - not receiving advertisements ... 6-34
 - precedence ... 6-4
 - priority
 - See* priority.
 - priority, configure ... 6-22
 - virtual IP address ... 6-10
- backup, configuring ... 6-20
- basic configuration steps ... 6-13
- configuring ... 6-3-??
- disable global ... 6-18
- disable on VR ... 6-24
- disabled during configuration ... 6-20
- display
 - all instances ... 6-29
 - configuration for all VRs ... 6-25
 - global configuration ... 6-25
 - statistics per VLAN ... 6-32
 - statistics, global ... 6-28
 - statistics, specific instance ... 6-33
 - uptime ... 6-29
 - VR, specific ... 6-27
- dropped packets ... 6-34
- election process ... 6-5
- enable global ... 6-18
- enable on VR ... 6-24
- event log ... 6-18
- event log messages ... 6-35
- example ... 6-6, 6-15, 6-16
- failback ... 6-3, 6-5, 6-7
- failover ... 6-3, 6-6, 6-7, 6-10
- IP address, deleting ... 6-12
- IP address, mismatch ... 6-17
- IP address, per VR ... 6-12
- IP address, real ... 6-5, 6-8
- IP address, virtual ... 6-5, 6-6, 6-8, 6-9, 6-10, 6-11
- IPv4 ... 6-34
- MAC address
 - shared ... 6-7
 - source ... 6-7, 6-11
 - virtual ... 6-11
- MAC address, virtual ... 6-10
- Master router ... 6-5, 6-9
 - advertisements failing ... 6-34
 - defined ... 6-4
 - election ... 6-8
 - Owner unavailable ... 6-8
 - See also* Owner router. ... 6-3
- multinetted VLAN ... 6-8, 6-10, 6-12
- overview ... 6-3
- Owner priority
 - See* priority.
- Owner router ... 6-5, 6-9
 - default Master ... 6-9
 - defined ... 6-4
 - priority
 - See* priority.
 - See also* Master router. ... 6-3
- owner, configuring ... 6-20
- pre-empt mode ... 6-10
- pre-empt mode, configure ... 6-24
- priority ... 6-6, 6-9
 - Backup ... 6-8, 6-11, 6-14
 - Backup default ... 6-9
 - Owner ... 6-5, 6-9, 6-11, 6-13
 - Owner default ... 6-10
 - range for Backup router ... 6-9
 - VR ... 6-10
- priority, Owner ... 6-8
- real gateway ... 6-9
- RFC
 - See* RFC.
- source address for VR ... 6-5
- standards compliance ... 6-33
- traps, disable ... 6-18
- traps, enable ... 6-18
- virtual router
 - See* VR.
- virtual router ID
 - See* VRID.
- VLAN, subnetted ... 6-9
- VR
 - advertisement interval ... 6-11
 - advertisement interval, change ... 6-22
 - changes ... 6-12
 - configure an instance ... 6-20
 - deactivate ... 6-12
 - defined ... 6-4


- IP address ... 6-12
- IP address limit ... 6-34
- IP address, delete ... 6-34
- MAC address ... 6-5, 6-8
- MAC address, source ... 6-6
- maximum in a VLAN ... 6-8
- maximum per switch ... 6-12
- maximum per VLAN ... 6-12
- membership ... 6-8
- multiple IP addresses ... 6-17
- multiple VRs in VLAN ... 6-8
- multiple, in a VLAN ... 6-9
- operation ... 6-8
- owner IP address ... 6-21
- priority
 - See* priority.
- subnet limit per VLAN ... 6-34
- virtual IP address ... 6-21
- virtual IP address, configure ... 6-23
- virtual IP address, default ... 6-23

VRID ... 6-6, 6-8, 6-10

- configure ... 6-19
- defined ... 6-4
- maximum per VLAN ... 6-19

W

warranty ... 1-ii



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

© Copyright 2007 Hewlett-Packard
Development Company, L.P.

Reproduction, adaptation, or translation
without prior written permission is prohibited
except as allowed under the copyright laws.

September 2007

Manual Part Number
5991-8586