PIM-SM (Sparse Mode)

Contents

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	4-61
Display Current RP-Set	n/a	4-63
Display Candidate-RP Data	n/a	4-65

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.

License Requirements

In the 3500yl and 5400zl switches, PIM-SM is included with the Premium License. In the 6200yl and 8200zl switches, this feature is included with the base feature set.

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.)

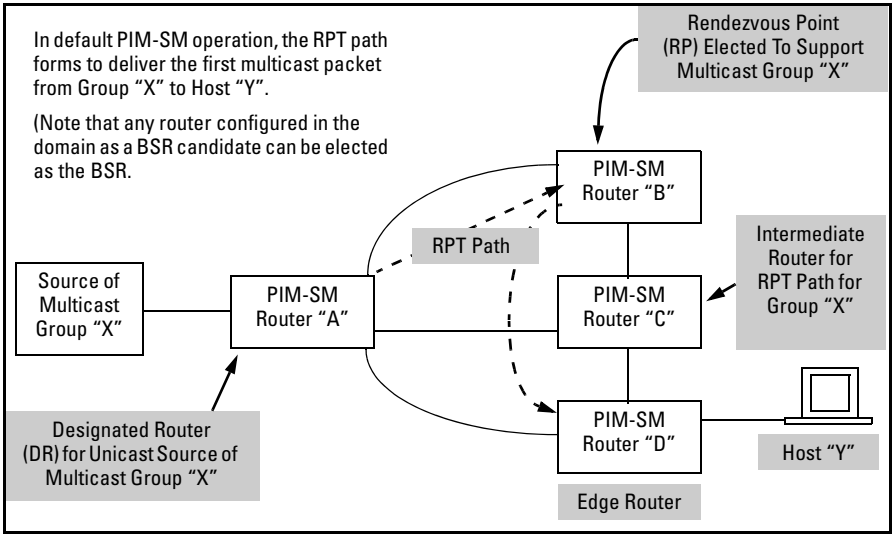

Figure 4-1. Example PIM-SM Domain with RPT Active To Support a Host Joining a Multicast Group

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 4-67.) 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.

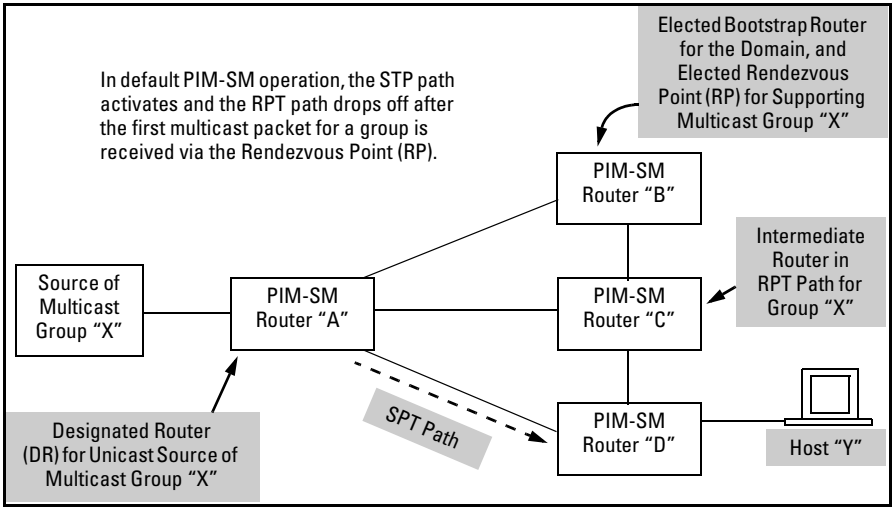

Figure 4-2. Example PIM-SM Domain with SPT Active To Support a Host that Has Joined a Multicast Group

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)

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

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

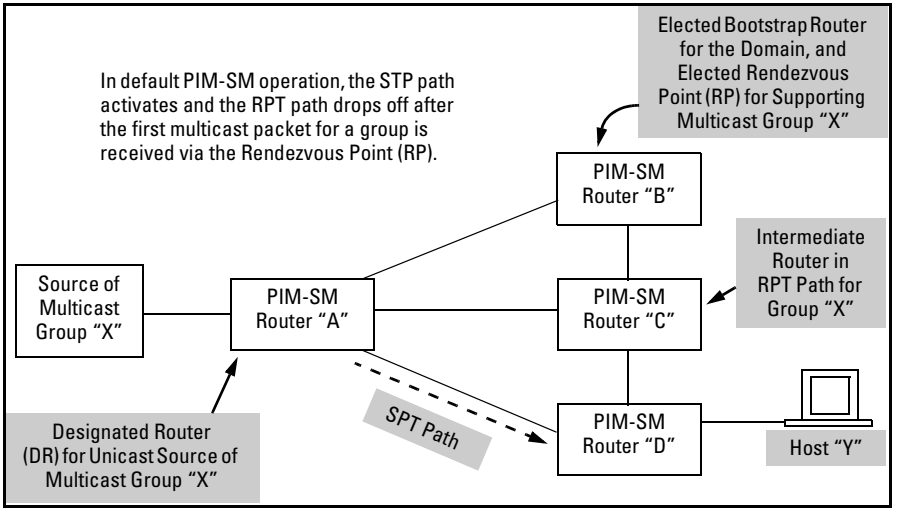

Figure 4-3. Example PIM-SM Domain with SPT Active To Support a Host that Has Joined a Multicast Group

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:

; J8693A Configuration Editor; Created on release #K.11.XX

hostname "ProCurve"
module 2 type J8705A
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
```

Global Routing Configuration for PIM-SM Support

Note: Either RIP, OSPF, or static routing can be used for a routing protocol.

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.

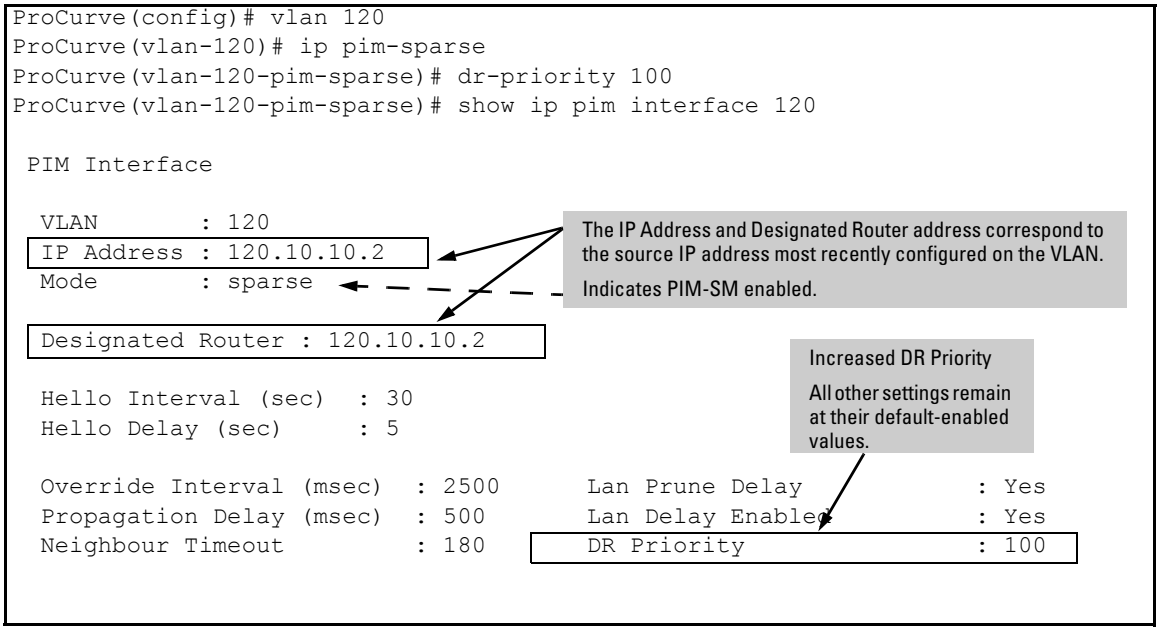

Figure 4-6. Example of Enabling PIM-SM in a VLAN

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.

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.

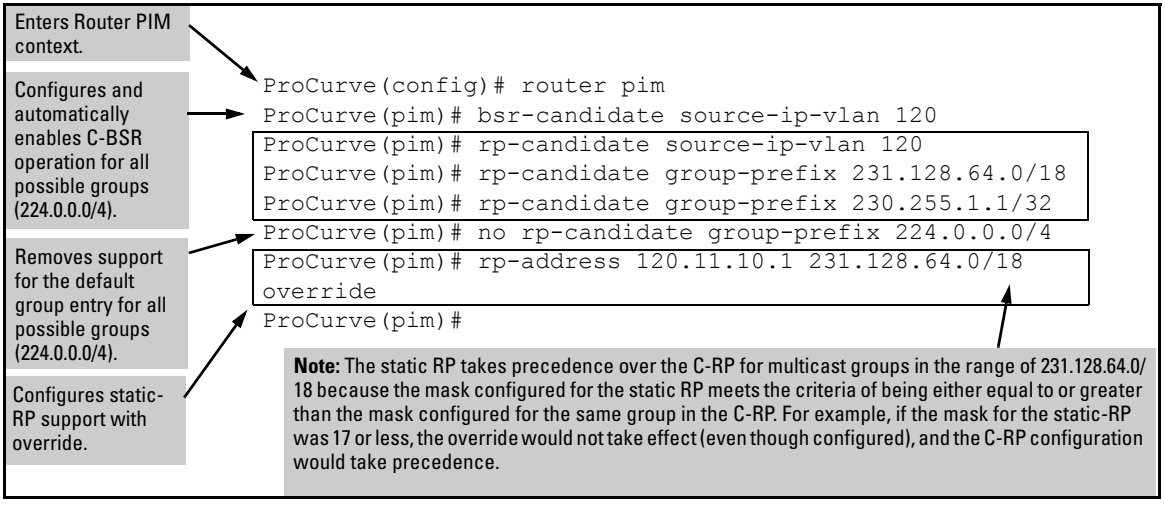

Figure 4-7. Example of Enabling PIM-SM in the Router PIM Context

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

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

 IP Multicast Route Entries

 Total number of entries : 2

 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 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 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.*

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.

```
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

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 Pending Join Requests

Use the **show ip pim pending** and **show ip pim rp-pending** commands to display the pending join requests received on the switch.

Syntax: show ip pim pending [< *ip-address* >]

Displays the joins received on the switch from downstream devices that want to join a specified (,G) or (S,G) multicast group (flow) address or all multicast groups known on the switch.*
A join remains in a pending state until traffic is received for the flow. The VLAN (PIM interface) on which each join was received is also displayed.

Incoming VLAN: *VLAN ID on which a join request is received.*

Source IPv4 Address: *IP address of the source of multicast traffic in an (S,G) group.*

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

Join Pending

  Group 224.0.3.4
    (*,G) Pending
      Incoming VLAN: 5
      Incoming VLAN: 3
    (S,G) Pending
     Incoming VLAN: 8         Source IPv4 Address: 10.0.3.9
     Incoming VLAN: 23       Source IPv4 Address: 10.0.3.10

  Group 224.0.11.8
    (*,G) Pending
      Incoming VLAN: 19
      Incoming VLAN: 88
```

Figure 4-20. Sample Output of show ip pim pending Command

Syntax: show ip pim rp-pending [< *ip-address* >]

Displays the joins received on the switch from downstream devices that want to listen to the multicast traffic in all (,G) or (S,G) multicast groups (flows) that a specified Rendezvous Point (RP) address or all RPs in the domain are responsible for.*

A join remains in a pending state until traffic is received for the flow. The VLAN (PIM interface) on which each join was received is also displayed.

Incoming VLAN: *VLAN ID from which a join request is received.*

Source IPv4 Address: *IP address of the source of multicast traffic in an (S,G) group.*

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

 (*,*,RP) Join Pending

  RP 10.0.4.4
    Incoming VLAN: 17

  RP 10.0.7.8
    Incoming VLAN: 2
    Incoming VLAN: 9
```

Figure 4-21. Sample Output of show ip pim rp-pending Command

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.

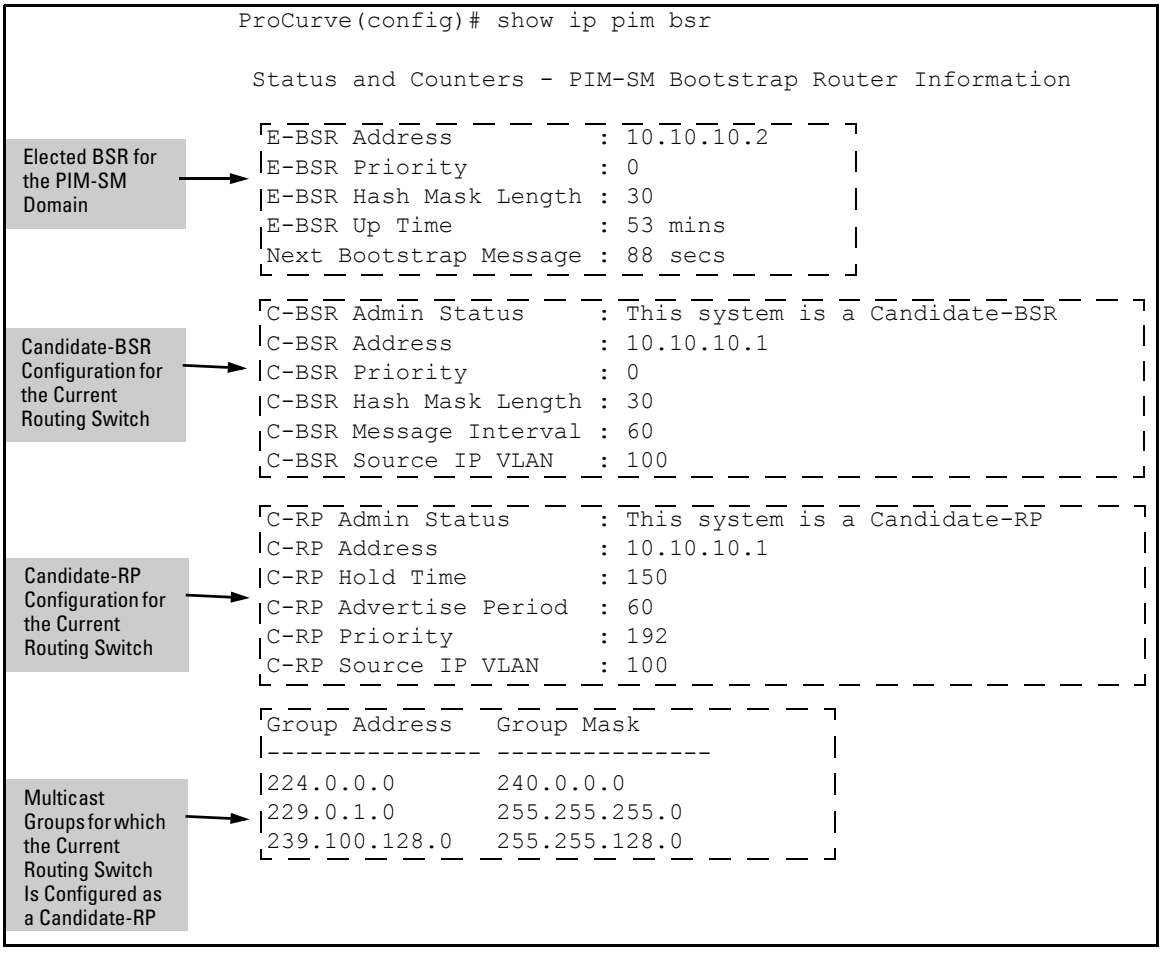

Figure 4-22. Example of Listing BSR Data for the Domain and the Immediate Router

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
.
.
.
```

Example of Non-Default BSR Candidate Configuration in the Router's Running Configuration

Note: priority appears only if it is configured to a non-default value.

Figure 4-23. 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.*

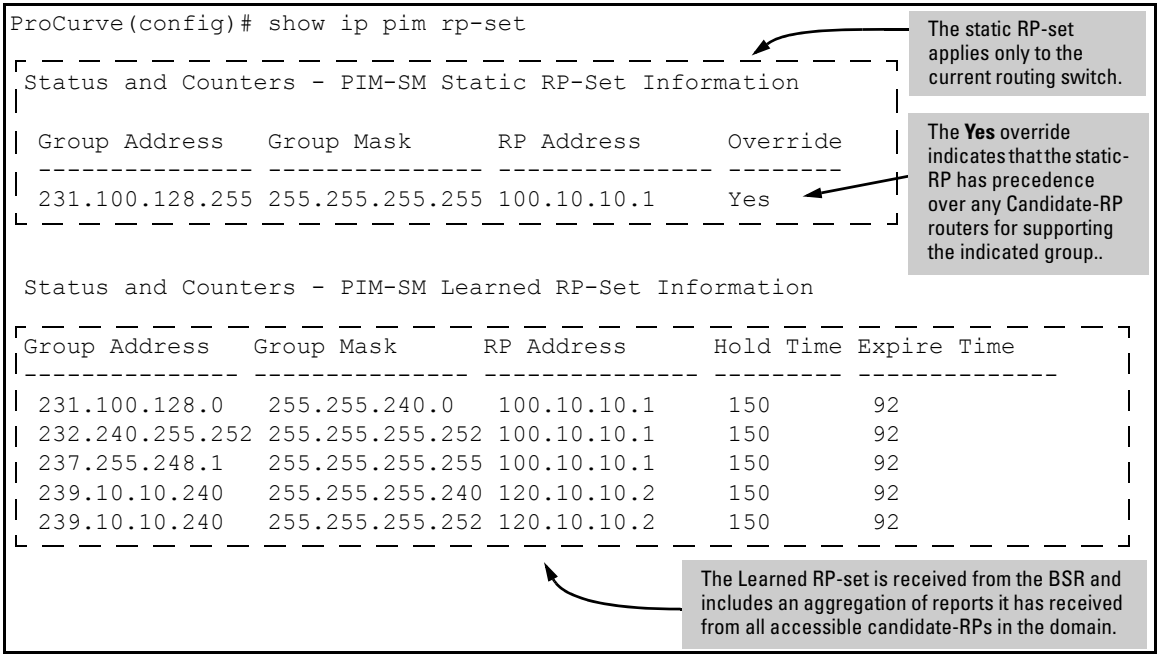

Figure 4-24. Listing Both the Learned and Static RP-Set Data

```
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 4-25. 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 4-26. 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 4-27. Example Listing for a Router that is Not Configured as a C-RP

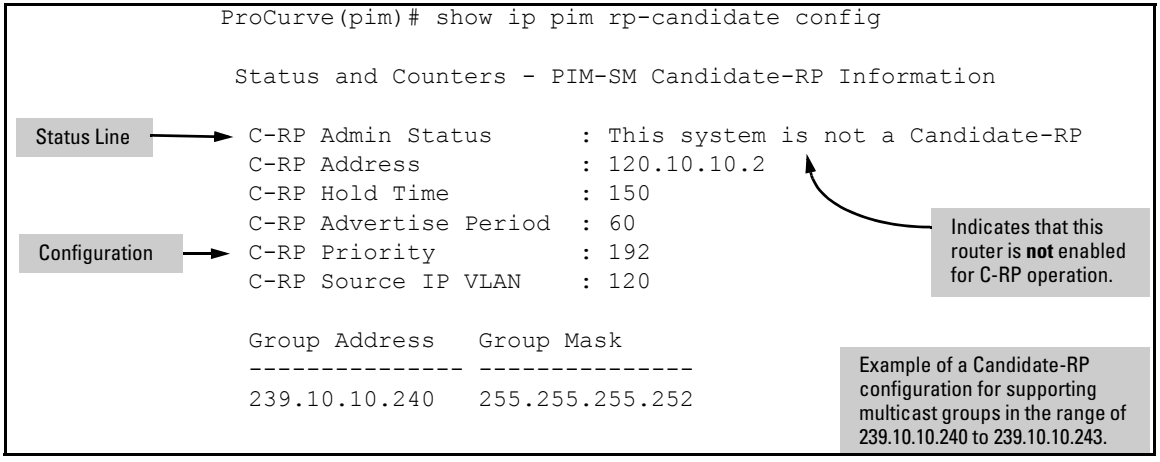

Figure 4-28. Example of the Full Candidate-RP Configuration Listing

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
      .
      .
      .
```

Example of Non-Default Candidate-RP Configuration in the Router's Running Configuration

Figure 4-29. 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
< *multicast-addr* >/< *mask* > 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.
<*pkt-type*> pkt, src IP < *ip-addr* > vid < *vid-#* > (not a nbr)	A PIM packet was received that doesn't have a neighbor.
Bad <parameter-name> in <*pkt-type*> pkt from IP < *ip-addr* >	The PIM packet was dropped due to a bad parameter in the packet from the IP address shown.
BSM send to < *ip-addr* > failed	A BSM (Bootstrap Message) send failed. The IP address shown is the BSM destination address.
Candidate BSR functionality disabled <*pkt-type*>	Candidate BSR functionality has been disabled.
Candidate RP functionality disabled	Candidate RP functionality has been disabled.
C-RP advertisement send to < *ip-addr* > failed	A C-RP advertisement send failed. The IP address shown is the destination address of the message.
Enabled as Candidate BSR using address: < *ip-addr* >	Candidate BSR functionality has been enabled at the indicated IP address.
Enabled as Candidate RP using address: < *ip-addr* >	Candidate RP functionality has been enabled at the indicated IP address.
Failed alloc of HW <*flow*> for flow < *src-ip-addr* >, < *multicast-addr* >	Hardware resources are consumed and software routing is being done for the flow.
Failed to initialize <*pkt-type*> as a call back routine	The IP address manager PIM callback routine failed to initialize.
Failed to alloc a <*pkt-type*> pkt (vid < *vid-#* >)	Allocation of a packet buffer failed message.
I/F configured with IP < *ip-addr* > on vid < *vid-#* >	The IP address on the PIM interface has changed to the indicated address.
I/F removal with IP < *ip-addr* > on vid < *vid-#* >	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 recvfrom 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.

Message	Meaning
`Rcvd unkwn addr fmly <`*add-family*`> in <`*pkt-type*`> pkt from <` *ip-addr* `>`	A PIM packet with an unknown address family was received.
`Rcvd <`*pkt-type*`> pkt with bad len from <` *ip-addr* `>`	A PIM packet with an inconsistent length was received from the indicated IP address.
`Send error(<` *error-#* `>) on <` *packet-type* `> pkt on VID <` *vid-#* `>`	Send packet failed on the indicated VLAN.
`Static RP configuration failure: <` *src-ip-addr* `>, <` *multicast-addr* `>`	The configuration of a static RP for the indicated multicast group has failed on the indicated interface.
`Unable to alloc a buf of size <` *size* `> for <` *memory element* `>`	PIM_DM could not allocate memory for the indicated buffer.
`Unable to alloc a msg buffer for <` *system-event* `>`	Informs the user that a message buffer could not be allocated for the indicated system event.
`Unable to allocate <`*table-type*`> table`	The PIM interface has been removed due to an IP address removal or change.
`Unexpected state/event <`*state*`>/<`*event*`> in <` *statemachine* `> 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.