

Configuring Multicast Support with PIM-SM

Contents

Overview	13-3
Multicast Trees	13-4
RP Tree	13-4
SP Tree	13-5
Multicast Routing Table	13-6
Joining a Shared or RP Tree	13-8
Switching from an RP to an SP Tree	13-9
RPs	13-9
Edge Routers	13-12
A Source's DR	13-14
Building RP and SP Trees When the Source Begins Multicasting First	13-15
A Source Begins Multicasting Before Any Hosts Join Its Group	13-15
A Host Joins a Group After Routers Have Already Switched to an SP Tree	13-16
RP Selection	13-17
PIM-SM Packets	13-18
Join/Prune Packets	13-18
Register Packets	13-25
Register-Stop Packets	13-26
Bootstrap Packets	13-26
Hellos	13-26
Asserts	13-26

Configuring PIM-SM	13-28
Enabling PIM-SM	13-29
Configuring a Static RP Set	13-30
Specifying Static RPs that Support All Groups	13-31
Specifying a Static RP for a Specific Group	13-32
Specifying When the Router Switches to the SP Tree	13-35
Forcing the Router to Use the RP Tree Permanently	13-36
Changing an Interface's DR Priority	13-36
Changing PIM-SM Timers	13-37
Join/Prune Period	13-38
Hello Timer	13-39
Override and Propagation Delay Timers	13-39
Configuration Examples	13-40

Overview

In order to receive multicast packets from one network and route them to hosts in different networks, a router must implement a multicast routing protocol. The ProCurve Secure Router supports Protocol Independent Multicast-Sparse Mode (PIM-SM).

Working with whatever unicast routing protocols the routers implement, PIM-SM allows routers in a PIM domain to construct a shared, unidirectional tree for each multicast group active in the network. The shared tree is rooted at a rendezvous point (RP) router, which is the router responsible for drawing multicast traffic from a new source to receivers in the associated group.

PIM-SM then allows a router to construct a more efficient, source-specific tree after it actually begins to receive multicast traffic.

PIM-SM also regulates the selection of:

- A designated router (DR) for each multi-access network—The DR forwards join/prunes for the subnet. It also forwards multicast traffic from any sources on the subnet.
- An RP for each multicast group—The RP acts as the root of the multicast tree; it receives join/prunes for the group, as well as receives multicasts from sources and forwarding them as needed.

To construct the unidirectional multicast trees, routers send PIM join/prune packets to RPs (or, for a source-specific trees, to the sources' DRs). A join/prune packet includes joins for the groups for which a router needs to receive multicasts. The packet also includes prunes for the groups for which the router no longer needs to forward multicasts to connected hosts or downstream routers.

Routers construct their multicast routing table based on the PIM join/prune packets that they receive and based on the Internet Group Management Protocol (IGMP) joins and leaves that they receive from connected hosts. An IGMP join for a group indicates that at least one host in a subnet is a member of that group. (See *Chapter 12: Configuring Multicast Support for a Stub Network* for more information on IGMP.)

An entry in the multicast routing table lists connections to downstream routers and networks as outgoing interfaces and the connection to the upstream router as the incoming interface. A router only accepts a multicast packet if it arrives on the appropriate incoming interface. If accepted, the router matches the packet to an entry in its multicast routing table according to its destination (and, optionally, source) address. The router then forwards the multicast packet through the outgoing interfaces indicated.

PIM-SM builds the multicast routing table conservatively. Interfaces only become outgoing interfaces for a multicast group if they specifically receive a PIM or IGMP join for that group. Because PIM-SM minimizes floods, it is well-suited for a WAN environment, in which bandwidth is often limited.

Multicast Trees

As hosts join and leave a group, the networks and routers that need the multicast traffic change. A multicast tree directs the flow of multicast traffic to a group's edge routers, which are the routers that connect to hosts active in the group.

A multicast tree indicates the upstream neighbor from which a router expects to receive multicasts for a specific group and the downstream neighbors to which the router must forward these multicasts.

PIM-SM defines two different types of tree:

- a shared, or RP tree
- a shortest path (SP) tree

RP Tree

Edge routers use the RP tree to reach the RP, which is the router that initially connects the routers that need multicasts for a group to sources for that group.

Using PIM-SM, routers generate a shared, RP tree for each multicast group. The leaves of the tree are networks that contain at least one host in that group. The root of the tree is the RP for the group. The tree is unidirectional: traffic must flow from the RP to the multicast hosts.

It might seem most logical that the root of the tree would be the designated router (DR) for the source of the traffic. (Indeed, the source's DR *is* the root of the SP tree.) However, multicast sources and receivers do not initially know where to find each other. A multicast group may have multiple sources, and

these sources may change. In addition, when hosts join a multicast group, they do not know the address of the source. Sources and receivers need a common point at which to discover each other, and the RP provides this point.

The DR of each subnet forwards join/prunes toward the RP so that the RP can begin forwarding multicasts to the appropriate routers as soon as a source begins transmitting. When a source does begin sending multicasts, the source's DR encapsulates and unicasts the multicasts to the RP, the predetermined distribution point.

On ProCurve Secure Routers, which use static RPs, the network administrator specifies the router or routers that can act as an RP.

SP Tree

An RP tree applies to multicast traffic from any source. Once a specific source begins to transmit a multicast stream, PIM-SM allows routers to generate a more efficient tree to that source. This tree is called the SP tree.

When a router receives a multicast packet, it learns the address of the multicast source from the source address in the packet's IP header. From its unicast routing table, the router may know of a better connection (one with a lower metric or administrative distance) to the source address than that through the RP. The router can use the better connection as an SP connection to the source. The process of using a packet's source address to find the upstream interface for a multicast tree is called reverse path forwarding (RPF).

Only edge routers can initiate the generation of an SP tree. Intervening routers, which transit multicasts but do not directly connect to multicast hosts, cannot join an SP tree in response to multicast traffic from a source. However, they can join an SP tree when they receive a PIM-SM join from a downstream router.

A group's RP can also establish an SP tree to the source, even if it does not also act as an edge router for that group. The RP uses its SP tree to receive multicasts from a source that has already registered with it. The RP can forward multicasts it receives on the SP tree to directly connected hosts, downstream routers in its SP tree, or routers that have not switched from the RP to the SP tree.

Like the RP tree, the SP tree is unidirectional. The root of the tree is the DR for the multicast source, and edge routers that connect to hosts in the multicast group are the leaves. Traffic always flows from the source to the edge routers.

The process for switching from an RP to an SP tree will be described in more detail in “Switching from an RP to an SP Tree” on page 13-9.

Multicast Routing Table

Just as a unicast routing table has an entry for each unicast destination address to which the route can forward traffic, a multicast routing table has an entry for every multicast group for which the router must transit traffic. (The router determines which groups these are according to the IGMP reports and PIM join/prunes that it receives.)

An entry in a unicast routing table includes information such as a destination address and forwarding interface. An entry in a multicast routing table can include:

- a multicast group address and, optionally, a source address
- a list of outgoing interfaces
- an incoming interface
- an RP address
- flags such as SPT or RP bits

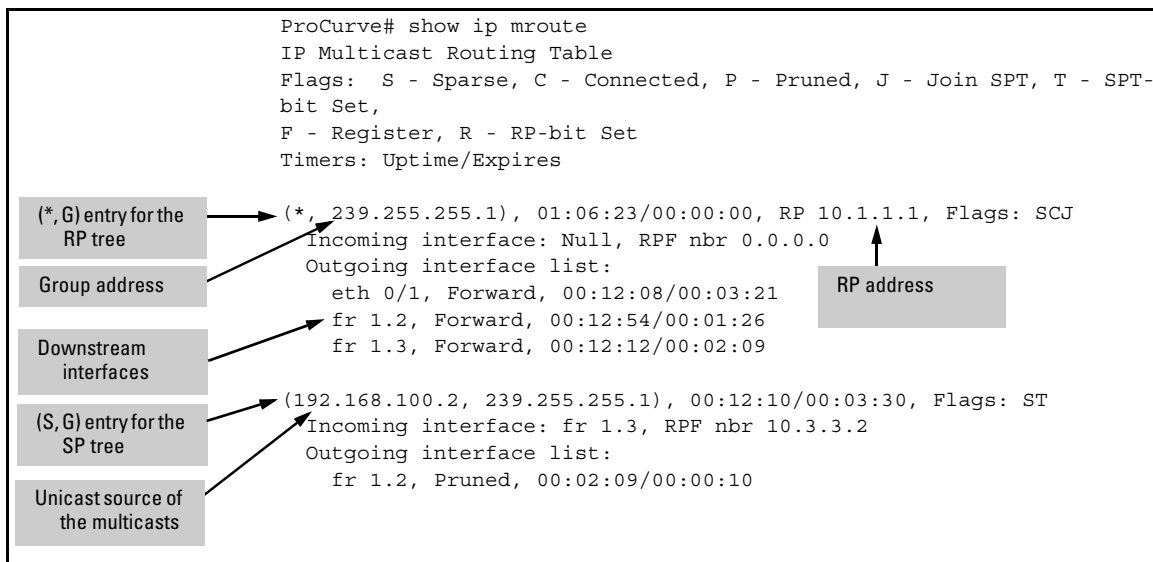


Figure 13-1. Multicast Routing Table

A multicast routing table entry, as shown in Figure 13-1, is identified by the address for the multicast group. It can, optionally, be made specific to a source.

Each entry includes a list of outgoing interfaces. Unlike a unicast routing table entry, a multicast table entry can include multiple forwarding, or outgoing, interfaces. Because a multicast address applies to all hosts who have joined the multicast group, and because these hosts may be in different networks, the router may copy packets destined to a single multicast address and route them out multiple interfaces.

For PIM-SM, the multicast routes form unidirectional trees; traffic for a particular group and from a particular source must always arrive on the same interface. Therefore, a multicast routing table entry also specifies the incoming interface on which the router expects to receive the multicasts. If traffic does not arrive on the correct incoming interface, the router drops it. A router forwards join/prune messages through the corresponding group's incoming interface.

Entries can also include the address of the group's RP and various flags which will be discussed in more detail later.

PIM-SM categorizes entries in a multicast routing table into two types:

- (*, G) entries, which are associated with an RP tree and apply to traffic for a specific group, but from any source
- (S, G) entries, which are primarily associated with an SP tree and apply to traffic for a specific group and from a specific source

When this guide refers to entries, a star (*) indicates a wildcard; G, a multicast group address; and S, the unicast address of the multicast source.

(*, G) entries are based on the RP tree and allow routers to receive multicast packets as soon as any source begins to transmit them. For example, a (*, 239.255.1.1) entry instructs the router how to forward any traffic destined for 239.255.1.1. The outgoing interfaces listed in the entry are downstream interfaces in the RP tree. The incoming interface is the interface through which the router connects to the RP.

(S, G) entries are source-specific and based on the SP tree. For example, a (192.168.1.254, 239.255.1.1) entry applies only to multicast traffic for the group 239.255.1.1 from the multicast source currently streaming traffic from 192.168.1.254. When a router adds an (S, G) entry to its table, it determines the incoming interface using RPF. That is, the incoming interface is the forwarding interface listed for the source address in the unicast routing table. The router copies the outgoing interfaces from the equivalent (*, G) group. The router will later receive PIM join/prunes, with which it can tailor the outgoing interface list to the SP tree.

Although (S, G) entries relate to SP trees, routers that are only part of an RP tree can also store special (S, G) entries with the RP-bit set. These entries prune downstream neighbors from the RP tree for multicasts from a specific source, but allow the neighbors to remain in the RP tree for traffic from other sources for the group.

A router always matches packets to the most specific entry first. In effect, it first processes (S, G) entries and then (*, G) entries. Therefore, by default, the router will always prefer an SP tree when such a tree exists for the source. After a router has not received packets from a source for a certain interval, the (S, G) entry times out, and the router reverts to using the (*, G) entry.

Joining a Shared or RP Tree

A router joins the shared RP tree when it receives:

- an IGMP join from a directly connected host
- a (*, G) join from a PIM neighbor

The router creates a (*, G) entry, completing these steps:

1. It selects the RP for the group. (See “RP Selection” on page 13-17 to learn how the router decides on the RP.)
2. The router uses RPF to determine the best connection to the RP. It sets the incoming interface and RPF neighbor for the entry accordingly.
3. The router includes the interface on which it received the join in the entry’s outgoing interface list.
4. The router sends a (*, G) join to the RPF neighbor.

The upstream router then follows the same steps and sends a (*, G) join to its upstream neighbor. (If the upstream router already has an entry for the group, it simply adds the interface on which it received the join to the entry’s outgoing interface list.)

The process continues until the join reaches the RP.

When the RP creates its own (*, G) entry, it sets the incoming interface to null and the RPF neighbor to 0.0.0.0; the RP is the root of the tree and does not need to forward joins to any upstream neighbor.

The section of the tree from the RP to the local router has now been established. Figure 13-2 illustrates this process.

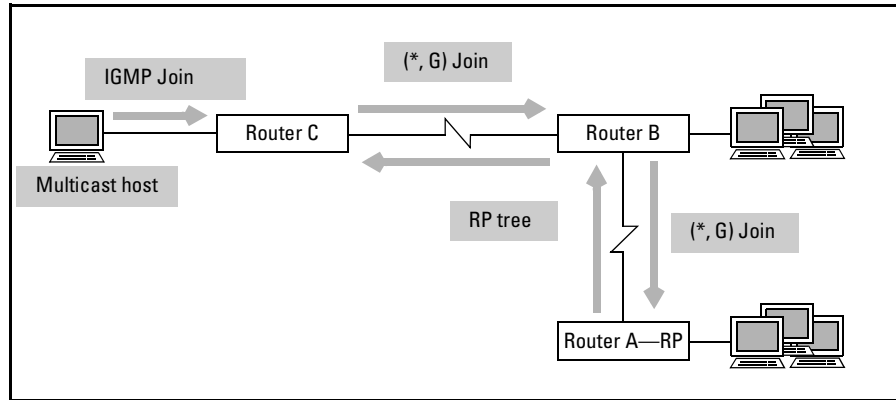


Figure 13-2. Joining a Shared, or RP, Tree

Switching from an RP to an SP Tree

Once a router begins to receive a multicast stream along the RP tree, it can change to an SP tree. In networks with many alternate paths or connections with varying bandwidth, the path through the SP tree may be quicker than that through the RP tree.

PIM-SM allows various methods for determining when a router should switch from the RP to an SP tree. A ProCurve Secure Router switches to the SP tree after it receives a certain number of multicast packets from a source. By default, the router begins to build the SP tree after it receives one packet, or, in effect, as soon as it receives a multicast for a group. (You can alter this threshold.)

RPs

The exception to the rule that only edge routers can initiate the switch to an SP is that the RP itself *must* switch to an SP tree. The J-bit in the (*, G) entry signals the RP to join the SP tree as soon as it receives a register packet from a source.

The RP follows this process to generate an SP tree to the source. (See Figure 13-3):

1. A source registers with the RP and the RP generates an SP tree to draw the multicast traffic towards itself and down the RP tree.

The RP initially receives encapsulated multicast traffic from a new source in unicast register packets. The RP decapsulates these packets and forwards them along the RP tree. It also creates an (S, G) entry and transmits a join towards the source.

2. The RP receives a multicast on the SP tree.

The source's DR receives the SP tree join from the RP. As well as unicasting register packets to the RP, the DR now forwards multicast packets to the RP. When the RP receives the first multicast packet, it sets the SPT-bit for the (S, G) entry. (Note that the RP, as the root of the shared tree, does not have an incoming interface for the (*, G) entry. Therefore, the incoming interface for the RP and the SP tree are always different.)

3. The RP sends a register-stop to the source's DR.

When the RP receives another unicast register packet, it sends a register-stop to source's DR. The DR stops sending the encapsulated multicasts, but continues sending multicasts to the RP because the RP is part of its SP tree.

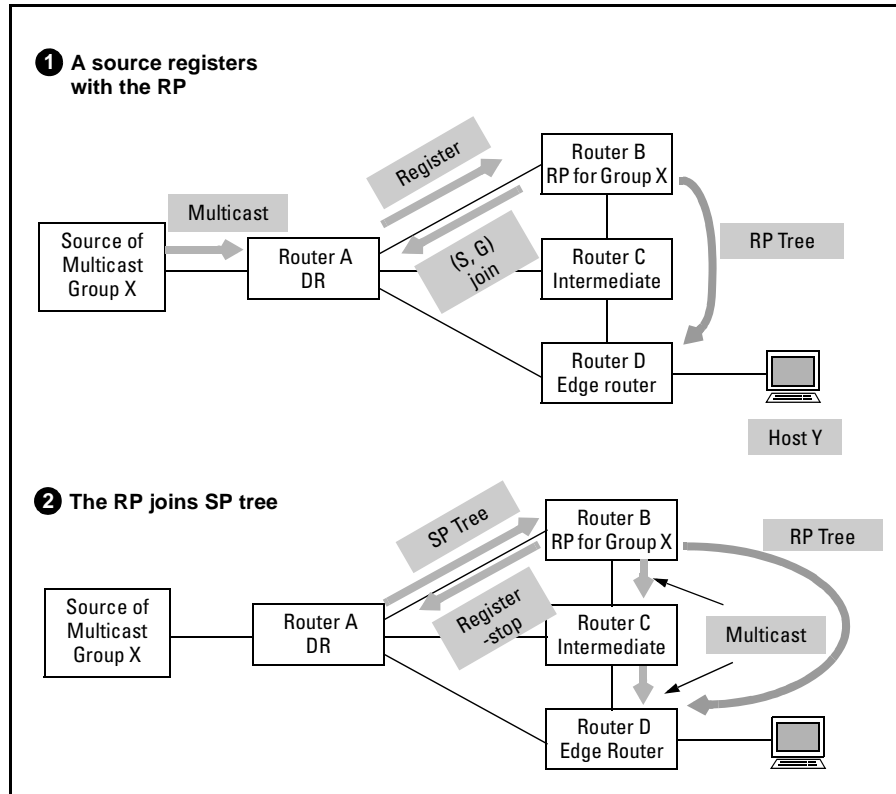


Figure 13-3. Switching to an SP Tree (RP)

4. The RP forwards multicasts over both the RP tree and an SP tree as needed.

When the RP created the (S, G) entry, it copied the outgoing interfaces from the (*, G) entry. This means the RP can continue forwarding traffic to routers that were downstream in the RP tree. As edge routers create their own SP trees, they can, if these trees diverge from the shared tree, send source-specific prunes towards the RP. These prunes indicate that the downstream routers no longer need multicast traffic from the RP.

If all routers directly connected to the RP prune themselves in this way, and if the RP is not itself an edge router for hosts in this multicast group, the RP will prune itself from the SP tree.

Figure 13-4 shows an RP's multicast routing table immediately after all routers in the PIM domain have switched to an SP tree that does not include the RP.

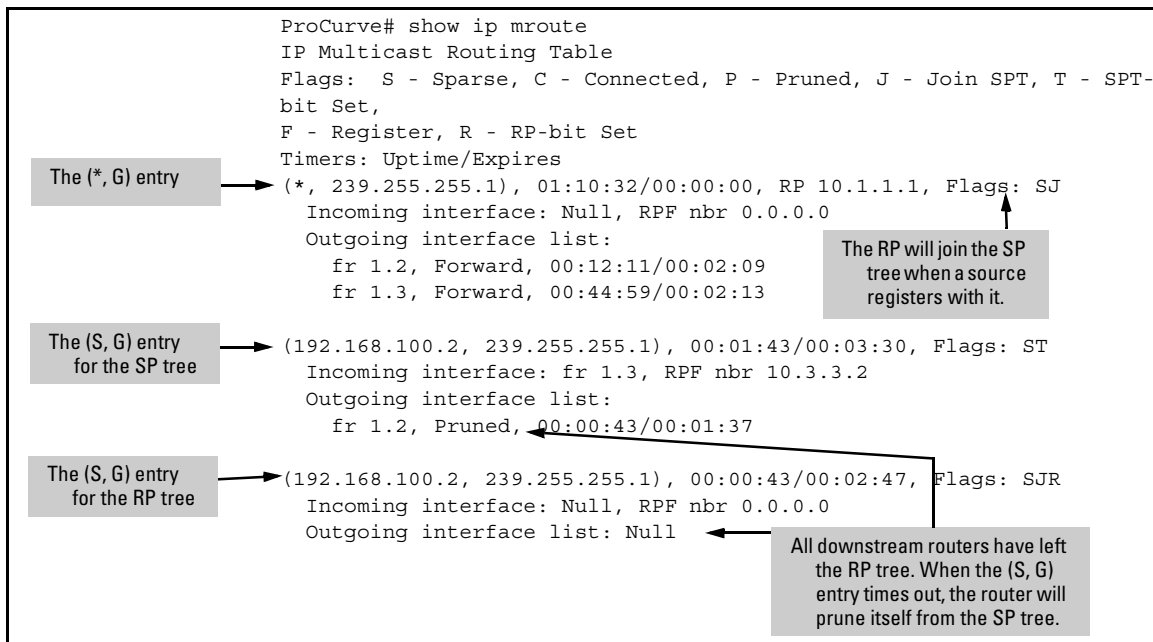


Figure 13-4. Multicast Routing Table for an RP

Edge Routers

Only edge routers that are directly connected to multicast hosts can initiate the transition to an SP tree. Intervening routers can join an SP tree, but only when they receive a join from a downstream neighbor, not when they receive multicast traffic.

When a ProCurve Secure Router receives a multicast packet from a source, it creates an (S, G) entry, taking the source from the packet's source address and the group address from its destination address. An edge router connects directly to a member of the group, so it sets the entry's C bit.

The router then uses RPF to determine which interface connects most directly to the source; this interface becomes the (S, G) entry's incoming interface. If this interface is different from the interface through which the router connects to the RP (and if the entry's C bit is set), then the router switches to the SP tree and prunes itself from the RP tree.

After it receives a multicast, an edge router follows this process to switch to the SP tree. (See Figure 13-5.)

1. The router creates the (S, G) entry, but continues to accept traffic from the RP tree.

An (S, G) entry's SPT-bit signals that the router is using the SP tree exclusively. When the router first creates the (S, G) entry, it clears the SPT-bit so that the multicast stream will not be disrupted while the SP tree is established. The router continues to accept packets from its upstream RP neighbor for as long as the (S, G) entry has a cleared SPT-bit.

2. The router joins an SP tree.

The router sends a join for the specific source and group out the incoming interface for the (S, G) entry. The upstream SP neighbor, in turn, creates an (S, G) entry, determines the forwarding interface for its best connection to the source, and transmits a source-specific join out that interface. This process continues until this branch of the SP tree has reached the root, the source's DR.

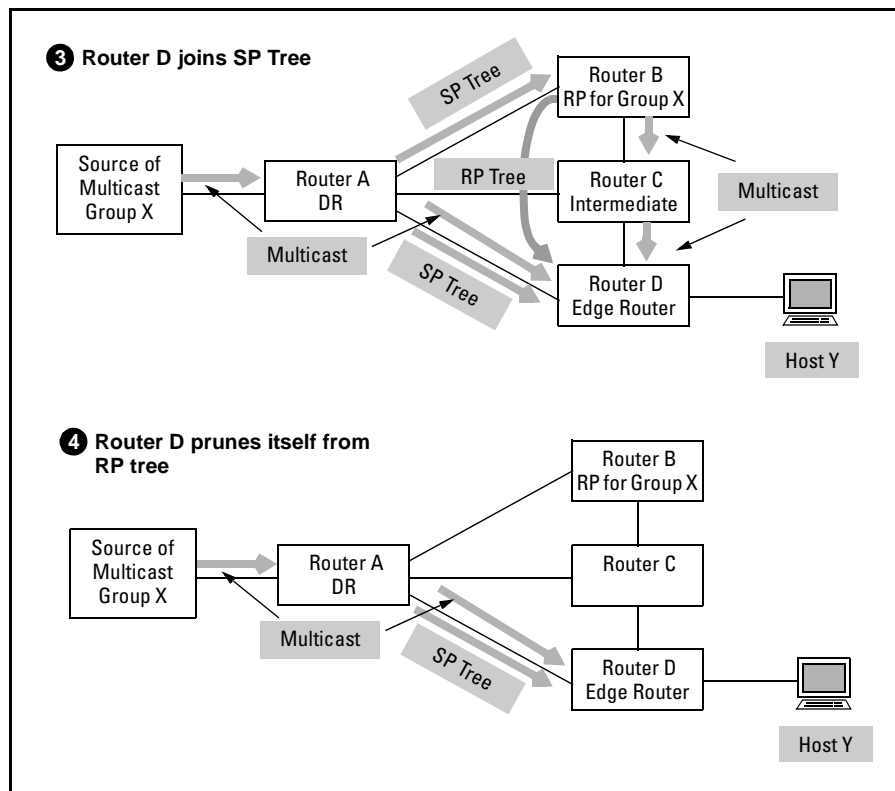


Figure 13-5. Switching to an SP Tree (Edge Router)

3. The router receives multicasts on the SP tree.

As soon as the original router receives a packet on the incoming interface for the (S, G) entry, it sets the entry's SPT-bit, signaling that the SP tree is active. The router now drops any packets except those from the SP neighbor.

4. The router prunes itself from the RP tree.

When it sets the SPT-bit for the entry, the router also sends a source-specific prune to its upstream RP neighbor. This prune removes the router from the RP tree, but only for traffic from that specific source. (The source-specific prune, which has its RP-bit set, is described in "PIM-SM Packets" on page 13-18.)

Note

This process only occurs if the best connection to the source is through a different interface than the best connection to the RP.

A Source's DR

A source's DR follows these steps in the transition from the RP to the SP tree:

1. The DR receives traffic from a multicast source and unicasts them to the RP.

As soon as the DR receives the first multicast packet from a directly connected source, it creates an (S, G) entry. The DR encapsulates the multicast packets in register packets and unicasts them to the RP for distribution over the RP tree.

2. The DR prunes itself from the RP tree.

At the same time, the DR sends a source-specific prune towards the RP to remove itself from the RP tree for traffic from this source.

3. The source, and through it the source's DR, becomes the root of the SP tree.

As edge routers build an SP tree, the DR receives (S, G) joins. When the DR receives an (S, G) join on an interface, it adds that interface to the outgoing interface list for its (S, G) entry. The DR can forward multicasts along an SP tree at the same time that it sends register packets to the RP.

4. The DR stops sending register packets to the RP after it receives a register-stop.

After the RP joins an SP tree, it sends a register-stop to the DR. The DR starts the register-suppression timer in its (S, G) entry and stops sending register packets to the RP.

- The DR continues forwarding multicasts over the SP tree.

Figure 13-6 shows the multicast routing table of a ProCurve Secure Router acting as the DR for a source.

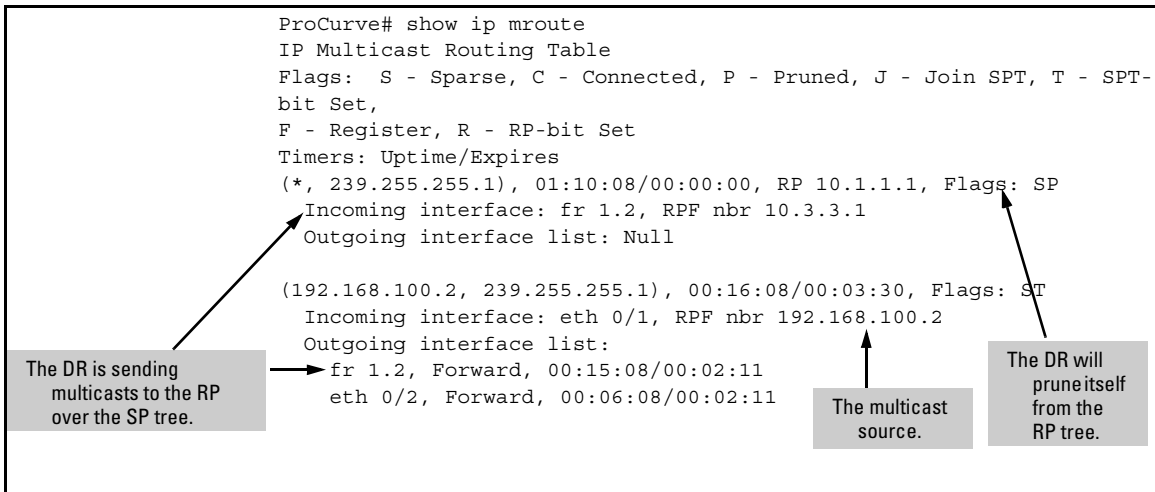


Figure 13-6. Multicast Routing Table of the DR for a Source

Building RP and SP Trees When the Source Begins Multicasting First

The above description of joining a shared tree and switching to an SP tree assumes that all receivers join a group before a source begins multicasting. However, this assumption is often untrue. The first host in a subnet to join a group may do so after a source has begun sending multicasts. The source may even begin sending multicasts before any hosts join the group.

PIM-SM provides for both of these situations. Trees are built and traffic forwarded in much the same way as that described above, but with several modifications that are described below.

A Source Begins Multicasting Before Any Hosts Join Its Group

- When a source registers with an RP that does not yet have an active (*, G) entry for the group, the RP creates both a (*, G) and an (S, G) entry for the traffic.

2. Although the RP creates the (S, G) entry, because the entry's outgoing interface list is null, the RP does *not* send a join for the SP tree. The RP also sends a register-stop to the source's DR.
3. The DR stops sending the encapsulated multicasts. It also drops all multicasts because it has not received any joins for its SP tree.
4. When a host *does* join the multicast group, the RP can activate the SP tree as soon as it receives the join.

Routers between the host and the RP build the shared tree by forwarding joins towards the RP. The RP simply adds the interface on which it receives the IGMP or PIM (*, G) join to *both* the (*, G) and the (S, G) entries. It then sends an SP tree join towards the source. The DR receives the join and begins forwarding multicasts instead of dropping them.

5. The edge router for the host can switch to the SP tree as described in "Edge Routers" on page 13-12.

A Host Joins a Group After Routers Have Already Switched to an SP Tree

A host can, of course, begin receiving multicasts immediately after it joins a group if other hosts in its subnet are also members of the group. However, PIM-SM forwards multicasts conservatively. If a subnet does not have any members in a group, then the subnet will not receive multicast traffic for that group. If a host does join the group, the traffic must be drawn down from the source to the subnet.

Typically, when a router receives an IGMP join for a new multicast group, it joins the corresponding RP tree. This could cause problems when a subnet requests traffic for a group for which the DR is now forwarding traffic on an SP tree.

PIM-SM addresses this issue by always forcing a router to copy the outgoing interface list of a (*, G) entry to an (S, G) entry. When the router receives the IGMP join, it adds the receiving interface to both entries.

In short, as soon as at least one host in subnet joins a group, it can begin receiving multicast traffic on any active SP tree for that group.

RP Selection

When a router adds an entry for a new group to its multicast routing table, it must determine the RP for that group. The router searches its RP set for up to four routers that can support that group. An RP set includes the IP address of every router allowed to become an RP and the multicast groups that each router can support.

The router then hashes each potential RP's address with the multicast group's address. The RP that produces the highest hash value becomes the RP for the group. Because all routers use the same hash function, they will always agree on the same RP for a group as long as they are using the same RP set.

PIM-SM allows these methods for distributing an RP set:

- Automatic—An administrator configures routers to be RP candidates (RP-Cs). A bootstrap router (BSR) collects advertisements from the RP-Cs and, from them, compiles an RP set. The BSR distributes the RP set to each PIM router in the domain.
- Static—The administrator manually configures the same RP set on each PIM router in the domain.

The ProCurve Secure Router supports static RP selection only. You can configure RP sets that contain:

- a single RP, which supports all multicast groups
- multiple RPs, each of which can potentially support any group
- multiple RPs, each of which supports only a specific set of groups
- multiple RPs, some of which can support any group and some of which are limited to certain groups

You must configure the same router to be RP for same set of addresses on each router in the PIM domain. For example, if you want to configure Router A to be the RP for multicast groups 224.0.0.0 through 231.255.255.255 in the network shown in Figure 13-7, then you must specify this configuration on Routers A, B, and C. Similarly if you want Router B to support all groups, you must specify this on each router, not only on Router B.

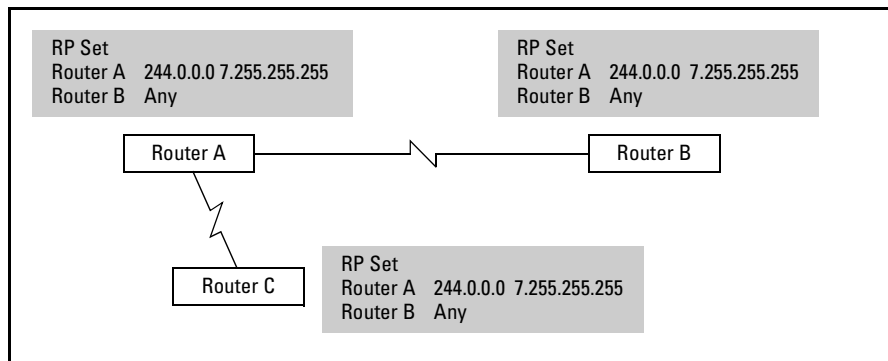


Figure 13-7. Static RP Selection

Note

Because you must configure exactly the same RP set on each router in the domain, attempts to assign specific routers to specific groups can lead to misconfigurations.

PIM-SM Packets

ProCurve Secure Routers exchange PIM-SM packets to:

- build entries in their multicast routing tables
- transition from an RP tree to an SP tree
- eliminate redundant multicasts in multi-access networks

If you so desire, you can learn more about the mechanics of these processes by reading about the various PIM-SM packets. Understanding basic information about these packets will help you to troubleshoot PIM-SM.

Join/Prune Packets

All PIM routers send join/prune packets. Routers send join/prune packets out the incoming interface for the group they want to join or leave. For RP tree join/prunes, this is the interface that connects to the RP or the upstream RP neighbor. For SP tree joins and prunes, this is the interface that connects to the source or the upstream SP neighbor.

The join/prune packet includes one or more group entries. Each group entry includes the group's multicast address, a join list, and a prune list. The join list and the prune list consist of addresses, either for RPs or for multicast sources.

If the router is sending the packet to its RP to either join or withdraw from the group's RP tree, the join or prune list contains a wildcard entry with the RP's address. An exception to this rule occurs when a router withdraws from an RP tree in order to join an SP tree. The router only withdraws from the RP tree as far as traffic from the specific source is concerned, so the prune list contains the source's address.

If the router is sending the packet in order to join or prune an SP tree, the group's join or prune list contains the address for the source specific to that tree.

Receiving (*, G) Join/Prunes. When a router receives a join/prune packet dealing with the RP tree, it can take these actions:

1. It ignores any group entries in which the wildcard source address does not match the RP for that group. This is one reason that it is very important to configure the same RP set on every router in your network.
2. If the *join* list for a group includes the wildcard source, the router searches for a (*, G) entry in its multicast routing table that matches that group. (It creates the entry if it does not exist.) The router then adds the interface on which it received the message to the entry's outgoing interface list.

In some circumstances, the router can also add the interface to a corresponding (S, G) entry. This allows new hosts to join a multicast group even after routers have switched to an SP tree for that group.

3. If the group's *prune* list includes the wildcard source, the router deletes the interface on which it received the packet from the (*, G) entry's outgoing interface.

Note

If the interface connects to a multi-access network, the router schedules the interface for deletion instead of immediately deleting it. This delay gives other routers in the network the opportunity to override the prune. (See the discussion of the override delay and propagation delay timers in "Changing PIM-SM Timers" on page 13-37 for more information.)

Receiving (S, G) Prunes. When a PIM router receives a join/prune message dealing with the SP tree, it either adds or prunes a connection as follows:

- If a group's *join* list includes the specific source, the router adds the interface on which it received the packet to the outgoing interface list for the corresponding (S, G) entry. If necessary, the router creates the entry, copying the outgoing interface list from the (*, G) entry.

If the upstream neighbor is itself part of the SP tree, it prunes the downstream router from its branch of the SP tree. If the upstream neighbor is not part of the SP tree, it creates an (S, G) RP-bit entry to prune the downstream router from its RP tree. (See Router D in Figure 13-8.)

A router follows these steps to prune the connection when it receives an (S, G) RP-bit prune:

1. If the upstream router does not already have an (S, G) entry corresponding to the prune, it creates one, completing these steps:
 - a. It copies the outgoing interface list from the corresponding (*, G) entry. In effect, the upstream router is creating a new entry to control the flow of traffic from the specific source over the RP tree.
 - b. The router calculates the incoming interface for the entry. Generally, the incoming interface for an (S, G) entry is the interface through which the router reaches the source. However, when an (S, G) entry has its RP-bit set, its incoming interface is the interface that connects to the upstream *RP* neighbor. This allows the router to use the (S,G)RP-bit entry to forward (or to suppress) traffic received on the RP tree.
 - c. The router deletes (or marks for deletion) the interface on which it received the prune from the (S, G) entry's outgoing interface list. That is, for as far as traffic from this source is concerned, the upstream router prunes the downstream router from the RP tree.

When this source sends multicast packets, the router will match them to the more specific (S, G) entry instead of to the (*, G) entry. Because the (S,G) entry does not include the pruned interface, the router will not forward the multicasts through it and the redundant traffic will be suppressed.

2. Similarly, if the upstream router already has an (S, G) entry *with* the RP-bit set, it deletes the interface on which it received the packet from this entry's outgoing interface list. In this case, the upstream router had already created an (S, G) RP-bit entry to remove other downstream routers using the SP tree from the RP tree. The upstream router now simply removes the router that sent the prune as well.

3. The upstream router may already have an (S, G) entry *without* the RP-bit set. For example, an RP generally creates an SP tree immediately after a source registers with it. Because the RP copies the outgoing interfaces in the (*, G) entry to the newly created (S, G) entry, the RP continues sending traffic over the connections in its RP tree. However, these connections are now as part of its SP tree.

When such a router receives an (S, G) RP-bit prune, it deletes the interface on which it received the packet from its (S, G) entry's outgoing interface list. This action removes the downstream router from the upstream router's branch of the SP tree. The downstream router remains part of the upstream router's RP tree. However, because the SP tree always overrides an RP tree, the upstream router will not forward multicasts to it. (In Figure 13-8 Router A is part of the SP tree, but it has also pruned Router D from this tree.)

In effect, an (S, G) entry acts as a place-holder, preventing the upstream router from sending redundant multicasts to a router using a different connection in SP tree than it used for the RP tree. (The upstream router will still forward multicast traffic from a different source along the RP tree and out the interface pruned in the (S, G) RP-bit entry.)

Note

The RP-bit prevents the (S, G) entry from being deleted even when the entry does not include any outgoing interfaces. This is necessary to prevent routers from forwarding redundant multicasts. You may therefore see an (S, G) entry in your router's multicast routing table that does not include any outgoing interfaces.

Sending Triggered Join/Prune Packets. A router sends a triggered join/prune packet when the state in one of its multicast routing table entries changes. For example, the router sends a triggered join when it receives an IGMP report for a new group and so creates a new (*, G) entry. Simply adding or deleting an outgoing interface from an existing entry is not enough to trigger a change: the router may be forwarding multicast traffic differently, but it still needs the traffic.

Table 13-1 displays some of the events that trigger join/prune messages and the type of packet the router sends as a result of the event.

Table 13-1. Triggered Join/Prune Packets

Event	Action	Packet Includes	Sent to
<ul style="list-style-type: none"> The router receives an IGMP join for a new or inactive group. The router receives a PIM (*, G) join for a new group. 	The router joins the RP tree.	join for the group with a wildcard source	upstream RP neighbor
The router receives an IGMP join for a group. The router has an (S, G) entry for the group for which the RP-bit is set and the outgoing interface list is empty.	Currently, the router is not forwarding multicasts because its neighbors are using an SP tree that does not include this router. When the router receives the IGMP join, it joins the SP tree so that it can deliver the multicasts to the new host.	join for the group with a specific source address	upstream SP neighbor
The router receives an PIM (*, G) join for a group. The router has an (S, G) entry for the group for which the RP-bit is set and the outgoing interface list is empty.	Currently, the router is not forwarding multicasts because its neighbors are using an SP tree that does not include this router. When the router receives the join, it forwards it towards the RP so that the downstream router can reach the RP.	join for the group with a wildcard source	upstream RP neighbor
The router receives a PIM join for traffic from a new source or from a source for which the (S, G) outgoing interface list is empty.	The router joins an SP tree.	join for the group with a specific source address	upstream SP neighbor
The last outgoing interface is deleted from a (*, G) entry.	The router leaves the RP tree.	prune for the group with a wildcard source	upstream RP neighbor
The last outgoing interface is deleted from an (S, G) entry (RP-bit not set).	The router leaves the SP tree.	prune for the group with a specific source address	upstream SP neighbor
A DR receives multicast traffic from a source in a directly connected subnet.	The router leaves the RP tree for traffic from this source. It will be the root of the SP tree.	prune for the group with a specific source address	upstream RP neighbor
The router receives multicast traffic from a new remote source and creates a new (S, G) entry.	The router determines whether it has a better connection to the source than through the RP. If it does, it begins to establish an SP tree.	join for the group with a specific source address	upstream SP neighbor

Configuring Multicast Support with PIM-SM

Overview

Event	Action	Packet Includes	Sent to
The router receives multicast traffic on its SP tree.	If the SP incoming interface is different from the RP incoming interface, the router sets the STP-bit for the (S, G) entry. In this case, the router starts to use the SP tree exclusively, and so it prunes itself from the RP tree.	prune for the group with a specific source address (RP-bit set)	upstream RP neighbor
The last outgoing interface is deleted from a (S, G) entry (RP-bit set).	The router remains in the RP tree, but refuses traffic from that specific source because all of its neighbors are using an SP tree for that source.	prune for the group with a specific source address (RP-bit set)	upstream RP neighbor
A change in the unicast routing table or in the subnet's DR changes the best connection to the RP.	The router's RP tree changes. After a randomized interval, the router sends a join to the new neighbor.	join for the group with a wildcard address	new upstream RP neighbor
A change in the unicast routing table or in the subnet's DR changes the best connection to a multicast source.	The router's SP tree changes. After a randomized interval, the router sends a join to the new neighbor.	join for the group with a specific source address	new upstream SP neighbor

Sending Periodic Join/Prune Packets. Routers also send join/prune packets at the close of every join/prune interval. (See “Changing PIM-SM Timers” on page 13-37 to set this interval.)

The router transmits these packets out the interfaces that connect to any upstream RP neighbor or upstream SP neighbor. In each packet, the router includes only the joins or prunes relevant to the neighbor receiving the packet.

To the upstream RP neighbor for a group, the router sends a periodic join/prune packet that can include:

- a join for a (*, G) entry with at least one outgoing interface
- a prune for a (*, G) entry with no outgoing interfaces
- a prune for an (S, G) entry with its RP-bit set and no outgoing interfaces
- a prune for an (S, G) entry with the STP-bit set, if the RP neighbor is different from the source-specific neighbor

To the upstream neighbor for a specific source, the router sends a periodic join/prune packet that can include:

- a join for an (S, G) entry with at least one outgoing interface
- a prune for an (S, G) entry with no outgoing interfaces

For example, Router A has an entry for (*, 239.255.1.1) with incoming interface PPP 1, outgoing interface Ethernet 0/2, and RP 192.168.1.1. Router A periodically sends a join/prune packet on PPP 1 which contains an entry for multicast group 239.255.1.1. This entry's join list contains 192.168.1.1, marked with a wildcard bit, for the source address.

Register Packets

Register packets encapsulate multicast traffic. The DR sends the register packet to the RP for the traffic's group to register with the RP and inform the RP of the presence and location of the source.

When a DR receives traffic from a multicast source, it looks up the RP configured to support that group. The DR then encapsulates the multicast packets in register packets and unicasts them to the RP.

The RP decapsulates the packets and forwards them as multicasts along the RP tree.

As soon as a source registers with the RP, the RP sends a join for an SP tree towards the source in order to draw the multicast traffic towards itself and the RP tree. Figure 13-9 illustrates the flow of multicast and register packets when a multicast source initially begins transmitting.

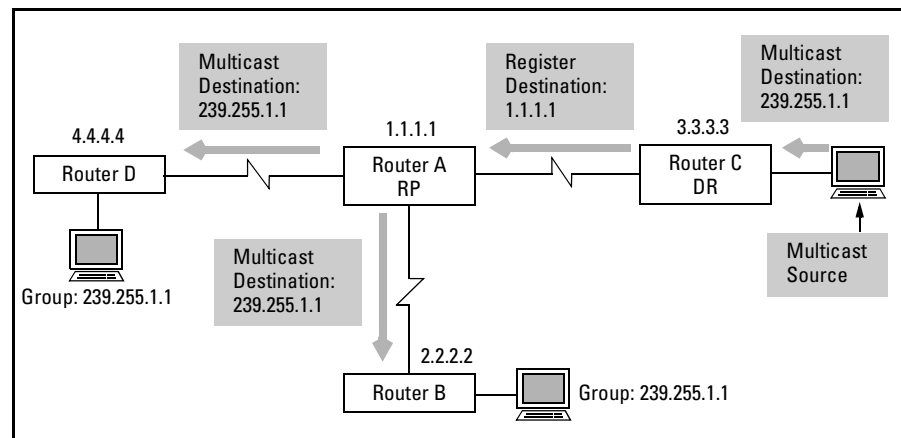


Figure 13-9. Sending Register Packets

Register-Stop Packets

After an RP begins receiving multicasts on the SP tree, it no longer needs the register packets. The RP sends register-stops to the DR for the source, instructing the DR to stop sending the encapsulated traffic. Register-stops are triggered when the RP has an (S, G) with the STP-bit set and receives a register packet.

A router also sends register-stops when it receives encapsulated traffic for a group for which it does not have an entry or for which the entry includes no outgoing interfaces.

Bootstrap Packets

Some PIM routers receive an RP-set (the list of all candidate RPs and the multicast groups that each can support) from bootstrap packets. Because you must statically configure an RP-set on a ProCurve Secure Router, it does not send or receive bootstrap packets.

Hellos

Routers periodically send PIM hellos out all interfaces. Neighbors set a timeout for the router based on a value in the packet. If the neighbor does not receive a hello from the router within this interval, it times out the connection. If the timed-out router was the DR for the subnet, another router on the subnet becomes DR.

Asserts

In a multi-access subnet, routers may receive redundant multicasts from each other. When a router receives a multicast packet on one of the outgoing interfaces for the group, the router knows that this packet has been sent, not by its upstream neighbor, but by another router on the subnet. The router drops the packet and multicasts an assert to all PIM routers in the subnet.

The assert identifies the multicast stream in question with the multicast source and the group addresses. The assert also includes the metric for the router's connection to the source. All PIM routers in the subnet listen for the asserts. They choose the router with the best (lowest) metric to become the subnet's sole forwarder for this multicast stream. All other routers delete the interface on that subnet from the outgoing interface list for the multicast route.

Figure 13-10 illustrates this process.

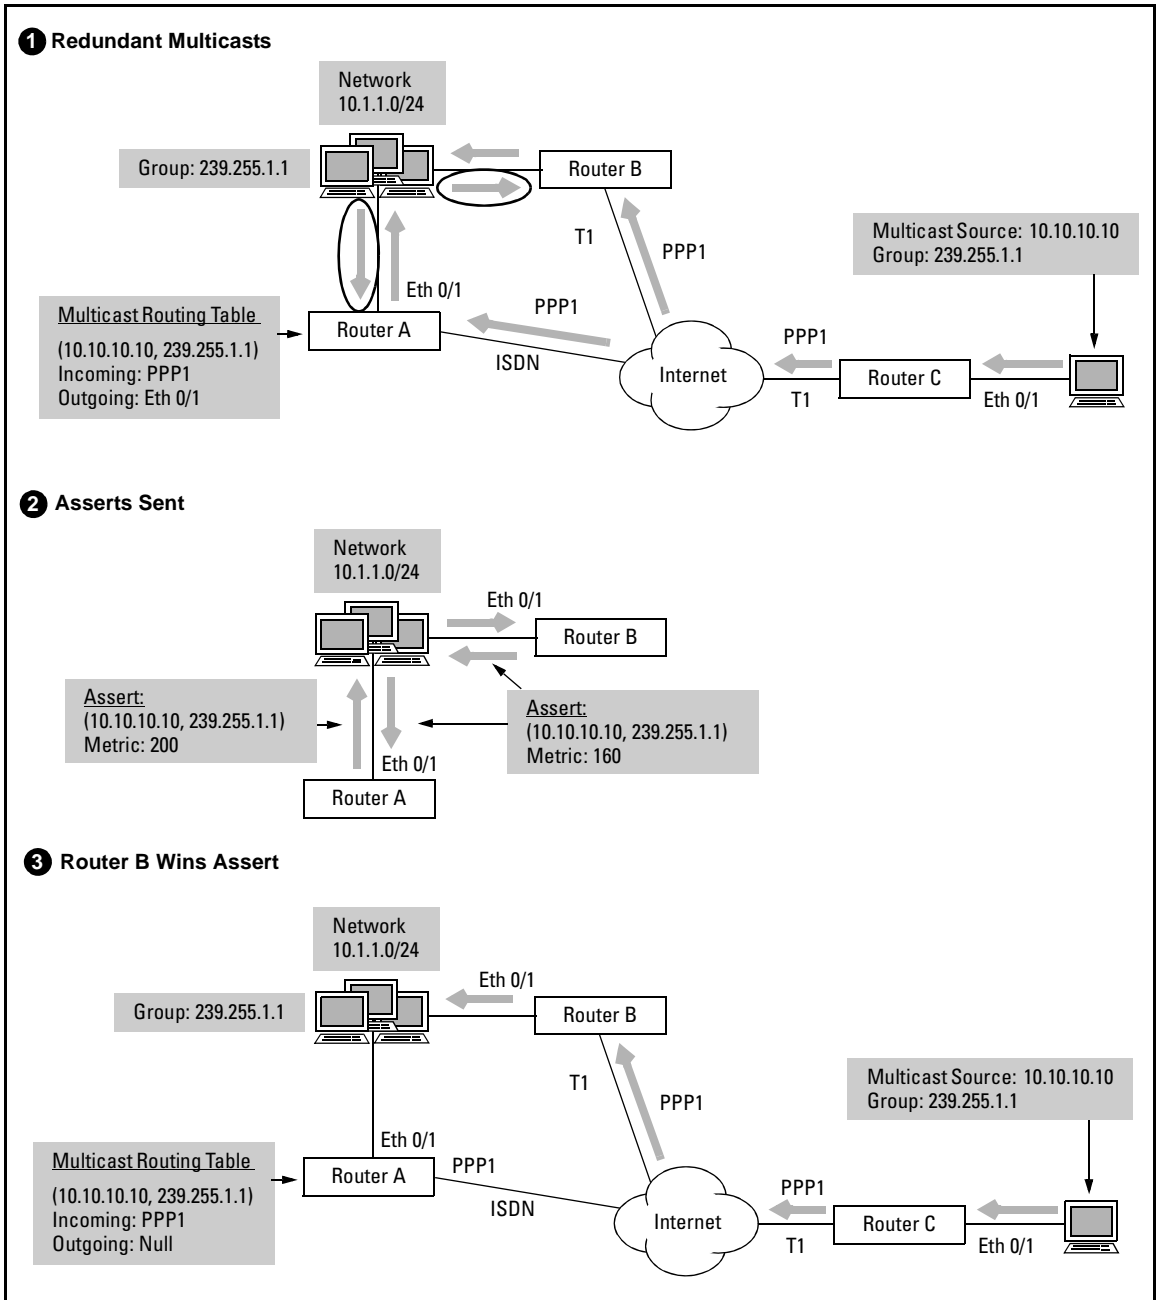


Figure 13-10. Exchanging Asserts

Configuring PIM-SM

To configure PIM-SM on a router, you must:

- enable PIM-SM on router interfaces
- specify the RP

PIM-SM relies on RPF to determine upstream neighbors. The protocol works with whatever routing methods the router uses, including:

- static routing
- Routing Internet Protocol (RIP)
- Open Shortest Path First (OSPF)
- Border Gateway Protocol (BGP)

You must implement one or more of these routing methods in order for PIM-SM to function. See *Chapter 15: IP Routing—Configuring RIP, OSPF, BGP, and PBR*.

Note

The routing table must include an explicit route to each potential RP and multicast source. Default routes are not sufficient. For example, if your network uses OSPF and your router is in a total stub area, the router will only receive a default route for inter-area traffic. It will not be able to use RPF to determine its upstream neighbors for RPs and sources outside of its area. You should remove the **no-summary** option from the **area <area ID> stub** OSPF command.

When configuring PIM-SM on the ProCurve Secure Router, you can optionally:

- configure multiple static RPs, each of which is responsible for a specific multicast group or groups
- change the threshold for switching to an SP tree
- force the router to use the RP tree permanently
- change an interface's DR priority
- change PIM-SM timers

Enabling PIM-SM

You must enable PIM-SM on every interface that connects to a network in the PIM domain. These networks include:

- LAN networks with hosts that may join the multicast groups
- LAN networks through which multicast traffic must transit
- WAN networks through which multicast traffic will travel between remote sites

The Layer 2 interfaces on the ProCurve Secure Router that support PIM-SM are:

- Ethernet interfaces
- Ethernet subinterfaces (VLAN interfaces)
- PPP interfaces
- Frame Relay subinterfaces
- ATM subinterfaces
- HDLC interfaces

To enable PIM-SM on an individual interface, move to that interface's configuration mode context. Then enter this command:

Syntax: ip pim sparse-mode

For example:

```
ProCurve(config)# interface frame-relay 1.101
ProCurve(config-fr 1.101)# ip pim sparse-mode
```

The interface will now listen for and send PIM join/prune messages.

PIM-SM works with IGMP. An edge router adds an interface to the outgoing interface list for a group only when it receives an IGMP report for that group on that interface. The **ip pim sparse-mode** command automatically enables the interface to act as an IGMP server. (See *Chapter 12: Configuring Multicast Support for a Stub Network* for more information on setting IGMP settings.)

After enabling PIM-SM on all interfaces, you should access the PIM sparse configuration mode context. Move to the global configuration mode context and enter:

```
ProCurve(config)# router pim-sparse
ProCurve(config-pim-sparse)#
```

From the PIM sparse configuration mode context, you can:

- specify static RPs
- change the threshold for switching to an SP tree
- force the router to use the RP tree permanently
- change the interval at which the router sends periodic join/prune messages

Configuring a Static RP Set

An RP for a multicast group forms the root of that group's RP tree. The RP tree is a unidirectional tree over which routers initially forward multicast traffic to hosts in the associated group.

An RP must:

- listen for join/prunes for groups
- receive register packets, which encapsulate multicast traffic, from the DR for a source
- decapsulate register packets and forward the multicast packets along the RP tree
- send register-stops to a source's DR after the RP has joined an SP tree
- keep track of routers that have switched to an SP tree so that it will not forward multicasts to routers using an SP tree that differs from the RP tree

You must configure the same RP set on each router in the network. The RP set consists of every router that you will allow to act as an RP and the addresses of groups that you allow that router to support.

A PIM router does not actually match a potential RP to a group address until it receives either an IGMP or a PIM join for that group. The router then performs a hash function on multicast address and the IP address of each RP that can support the group. The statically configured RP that produces the highest hash value becomes the actual RP for the active group.

You do not need to take excessive care when deciding which router or routers to add to the RP set. Because an RP tree is typically only briefly active, an RP is not unduly burdened. However, you *should* be very careful to specify that the same RPs support the same groups on each router in the PIM domain. You must enter the configuration on each PIM router. Otherwise, the router will not know how to forward join/prunes, will not join the RP tree, and will not forward multicasts.

For the simplest configuration, and the configuration least prone to errors, you should allow all RPs to support any group. There is no reason to configure different RPs for various groups unless you expect these conditions to be true:

- only certain areas of the network will use certain groups
- having a router act as RP for groups expected in its area will significantly decrease bandwidth usage
- routers will not immediately switch to the SP tree

Specifying Static RPs that Support All Groups

To specify the IP address of a router that you have selected as a potential RP, move to the PIM configuration mode context and enter:

Syntax: rp-address <A.B.C.D>

The command enables the router specified to act as RP for all multicast groups. Enter this command on all routers in the PIM domain.

You can specify up to four router IP addresses in the RP set. Each router can potentially act as RP for any multicast group. However, the router will only actually do so if selected by the hash function PIM routers perform when they add a new group to their multicast routing tables.

Notes

If you configure only one RP on a router and that RP fails, you must manually reconfigure a new RP on each router.

Every router must have *exactly* the same RPs in their RP sets. For example, if you allow Router A to use itself or Router B as RP, it is not enough to allow Router B to use itself. You must also add Router A to Router B's RP set. Otherwise, Router B would be left out of multicasts for which Router A selects itself as the RP.

The IP address specified for an RP must be exactly the same address specified for that RP on every other router in the WAN. It is not enough to specify any IP address on one of the RP's interfaces. For this reason, it can be a good idea to use the IP address of a loopback interface to identify an RP.

Specifying a Static RP for a Specific Group

Instead of configuring the same routers to support all multicast groups, you can associate specific RPs with specific groups.

You should only use this option if your organization has a particular reason for doing so. Usually, since routers immediately switch to an SP tree, the location of the RP is not as important as it may seem.

However, if you force routers to use the RP tree permanently, the location of the RP in the network topology becomes more important. In this case, configuring a hub router in the area in which you expect certain groups to be active to be RP for those groups may significantly reduce the bandwidth consumed by multicasts.

You can configure multiple RPs for the domain. Each RP can support a single multicast group or a range of groups.

You can overlap address ranges, but doing so can lead to unexpected results. For example, you can configure Router A to support 232.0.0.0 /5 and Router B to support 239.255.255.1 /32. Both routers would then support the multicast group, 239.255.255.1. Router B would not necessarily become the RP for the single group that it supports. You should disable Router A from supporting 239.255.255.1 with a deny statement in the ACL associated with it.

When configuring a static RP set, you should keep these tips in mind:

- At least one RP should be able to support every multicast group. Otherwise, routers may not be able to route necessary multicasts. For example, routers running various networking protocols multicast messages to addresses between 224.0.0.0 and 224.0.0.255.

Note

You should particularly avoid leaving the 224.0.0.0 to 224.0.0.255 range without RP support: doing so poses the risk of disabling one or more protocols needed in the network.

- When you attempt to configure static RPs to support ranges of addresses of varying lengths or to configure some RPs to support ranges and others to support a single group, you open room for misconfigurations. You may unintentionally leave a necessary multicast group without an RP to support it.
- If you expect hosts throughout the PIM domain to join roughly the same groups, you should consider allowing all RPs to support all groups. You could alternatively split the range of multicast groups evenly among the routers you want to act as RPs. However, this opens the possibility for a group being left without support if a router fails.

- If you know precisely which groups your network must support and you know which areas expect traffic for specific groups, you can configure a router to support a single group. For example, the multicast video streamer in Figure 13-11 is the only source that sends traffic to 239.255.255.1. You could associate Router A with this single group when you configure the static RP set on each router in the WAN.

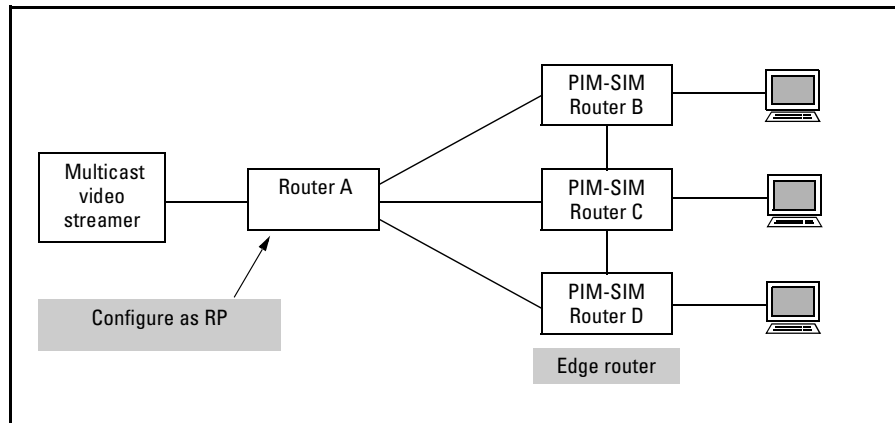


Figure 13-11. Configuring a Static RP

Remember to always configure at least one RP to support all groups so that necessary network protocols are not blocked. You can deny the addresses for the groups specifically supported by another router from the list of groups supported by this default RP.

Specifying the Groups Supported by a Static RP. You use an access control list (ACL) to associate an RP with a particular group or set of groups. When you configure the ACL, you enter the IP address of the group or groups that the RP supports as a permitted *source*.

Complete these steps to configure the ACL:

1. From the global configuration mode context, create the ACL:

Syntax: ip access-list [extended | standard] <listname>

Because you enter the multicast address as a source address, you should use a standard ACL. For example:

```
ProCurve(config)# ip access-list standard RP1
```

You will then enter the ACL configuration mode context.

2. If necessary, you can remove a group from the range of groups for an RP with a deny statement. Use this command:

Syntax: deny [host <A.B.C.D> | <A.B.C.D> <wildcard bits>]

For example, you want Router 1 to be RP for all multicast groups except for group 239.255.255.1, which will be used in only one section of the network. You would configure the ACL for Router 1 to exclude the specific group:

```
ProCurve(config-std-nacl)# deny host 239.255.255.1
```

You would then enter a permit statement to allow all multicast addresses as follows:

```
ProCurve(config-std-nacl)# permit 224.0.0.0 15.255.255.255
```

You must enter the deny statement before the statement that permits the range of groups because the ProCurve Secure Router processes ACL entries in order and stops as soon as it finds a match. The more general permit statement allows the address that should be denied.

3. Enter a permit statement for the group address or range or addresses that the RP will support:

Syntax: permit [host <A.B.C.D> | <A.B.C.D> <wildcard bits>]

Enter an address with wildcard bits to select a range of groups. For example, you can configure an RP to support half of all possible multicast groups, those between 232.0.0.0 and 239.255.255.255. A one in the wildcard bits instructs the router to ignore the associated bit in the address. Another way of thinking of wildcard bits is that adding the bits to the specified address gives you the last address in the range. For this example, you would enter:

```
ProCurve(config-std-nacl)# permit 232.0.0.0 7.255.255.255
```

If you are configuring a static RP support a single group, use the **host** keyword to specify the address for that one group. For example:

```
ProCurve(config-std-nacl)# permit host 239.255.255.1
```

Associating the Static RP with the ACL. Next, you must specify the address of the RP and associate it with the ACL you have configured for it. Move to the PIM sparse configuration mode context and enter this command:

Syntax: rp-address <A.B.C.D> access-group <listname>

Note

You may want to limit an RP that currently supports all groups to only supporting some groups. In this case, you must first enter **no rp-address <A.B.C.D>**. You can then re-enter the command with the specification for the ACL that lists the groups the RP should support.

Specifying When the Router Switches to the SP Tree

After a source registers with the RP, the RP builds an SP tree to the source. The RP can then distribute multicast traffic from this source to routers in the RP tree.

Edge routers join an RP tree so that they can receive and forward multicast packets as soon as a source begins transmitting them. After an edge router receives a multicast packet for a group in which it is active, it too can join the SP tree. The router looks up the source's address in the packet's header and determines whether it has a better route to that source than the one through the RP. If the router does find a better route in its unicast routing table, it joins the SP tree by sending a source-specific join out the forwarding interface for that route.

By default, an edge router begins to switch to an SP tree after it receives one multicast packet from a source. Quickly switching to an SP tree is particularly important for high-rate multicast applications. The SP tree reduces delay and minimizes network congestion.

If your network primarily experiences lower-rate multicasts or brief multicast spurts rather than sustained, high-rate multicasting, you may want to raise the threshold so that the router does not toggle between the RP and SP tree.

To change how many packets the router must receive from a multicast source in order to trigger the switch to an SP tree, move to the PIM sparse configuration mode context and enter:

Syntax: spt-threshold <packets>

You can enter a value between 1 (the default) and 4294967295 for the number of packets.

Note

The SPT threshold only applies to edge routers switching to the SP tree. The RP must generate an SP tree to a registered source, and any SPT threshold that you set does not apply to this tree.

Note

The PIM-SM protocol automatically manages the transition to the SP tree, keeping the RP tree active until convergence is complete. For more information on this process, see “Switching from an RP to an SP Tree” on page 13-9.

Forcing the Router to Use the RP Tree Permanently

A router’s SP tree is tailored to be the best connection between the router and a specific source, and you should almost always allow your ProCurve Secure Router to use this tree as soon as it can. However, if you need to tightly control the flow of multicast traffic in your network, you can force routers to use the RP tree rather than finding their own shortest paths.

Move to the PIM sparse configuration mode context and enter:

Syntax: spt-threshold infinity

Note

The RP for a group always builds an SP tree to a registered source.

For more information on an SP tree versus an RP tree, see “Multicast Trees” on page 13-4 and “Switching from an RP to an SP Tree” on page 13-9.

Note

You should not force a router to use the RP tree permanently out of fear that new hosts will not be able to join a group after routers switch to the SP tree. PIM-SM automatically provides for this situation.

Changing an Interface’s DR Priority

A multi-access network may include several PIM routers. However, only one router should forward join/leave packets for the network. When a network includes a source of multicast traffic, only one router should encapsulate this traffic and send it to the RP. Routers in each subnet elect a DR to perform these functions.

The router in the subnet with the highest DR priority becomes the DR. By default, a ProCurve Secure Router’s DR priority is 1. When multiple routers have the same priority, the router with the highest IP address becomes DR.

If it is important that a specific router become DR, you can raise its priority. If the router must *not* become DR, you should leave the priority at one and raise another router’s priority.

Because a router can have interfaces on several different networks, you set the DR priority for each specific interface. You can assign different interfaces different priorities. For an example, your ProCurve Secure Router connects to VLAN 10 on Ethernet subinterface 0/1.10 and to VLAN 20 on Ethernet subinterface 0/1.20. If you only want the router to be DR on exactly one subnet, you could assign Ethernet subinterface 0/1.10 a high priority, but Ethernet subinterface 0/1.20 a priority of one.

You only need to set a priority on interfaces on multi-access networks. On the ProCurve Secure Router, these are Ethernet interfaces and subinterfaces. To specify the priority, move to the interface configuration mode context and enter:

Syntax: ip pim-sparse dr-priority <value>

You can enter a number from 1 to 4294967295 for the priority value.

Changing PIM-SM Timers

You can alter the length of these PIM-SM timers on the ProCurve Secure Router:

- join/prune period
- hello timer
- neighbor timeout
- override timer
- propagation delay

Enter the commands shown in Table 13-2 to configure PIM timers. The following sections discuss the timers in more detail.

Caution

You should generally leave the timers at their defaults unless you have experience working with PIM-SM. If you do not take great care when setting the timers, you can disrupt PIM-SM's functions. For example, if you lower the neighbor timeout on your router without also lowering the hello timer on neighboring routers, your router may time out neighbors on good connections.

Table 13-2. PIM-SM Timers

Timer	Meaning	Command Syntax	Configured From	Range	Default
join/prune period	time between sending period join/prunes	join-prune-msg-interval <seconds>	PIM configuration mode context	10 to 65535 seconds	60 seconds
hello timer	time between sending hellos	ip pim-sparse hello-timer <seconds>	Ethernet or WAN interface configuration mode context	10 to 3600 seconds (3.5 hours)	30 seconds
neighbor timeout	maximum time the router will wait for a hello before considering a neighbor down	ip pim-sparse nbr-timeout <seconds>	Ethernet or WAN interface configuration mode context	30 to 10800 seconds (3 hours)	105 seconds
override delay	maximum time the router will wait before sending a join to override a prune sent by another router on the multi-access network	ip pim-sparse override-interval <milliseconds>	Ethernet or WAN interface configuration mode context	0 to 65534 milliseconds	2500 milliseconds
propagation delay	time the router expects it will take its join to reach an upstream neighbor	ip pim-sparse propagation-delay <milliseconds>	Ethernet or WAN interface configuration mode context	0 to 32767 milliseconds	500 milliseconds

Join/Prune Period

The timer for the join/prune period determines how often the router sends out periodic join/prunes. These messages renew the router's membership in groups for which it must forward multicasts. They also prune the router from trees for which it no longer needs traffic. The router constructs the join/prune packet and determines through which interfaces to transmit it as described in "Sending Periodic Join/Prune Packets" on page 13-24.

You configure the join/prune period from the PIM sparse configuration mode context. See Table 13-2 for the exact command syntax.

You set the other timers for individual interfaces.

Hello Timer

Routers transmit periodic hellos through PIM interfaces to signal that the connection is still active. The **hello-timer** option determines how often an interface sends a hello. The router also uses this setting to compute the hello holdtime, which it includes in hello packets to instruct neighbors how long to wait for the next hello before removing the connection from any outgoing interface lists. (The hello holdtime should be three and a half times the hello timer to allow for the occasional lost packet.) Shortening the hello timer indirectly reduces the hello holdtime, which means that upstream routers will more quickly stop sending unnecessary multicasts.

The neighbor timeout (**nbr-timeout** option) is the maximum amount of time the router itself will wait for a hello message from a neighbor before considering that neighbor unreachable. You should generally configure the neighbor timeout to three and a half times the neighbor's hello timer. The neighbor timeout configured on the interface takes precedence over the hello holdtime included in the neighbor's hello.

Override and Propagation Delay Timers

The override and propagation delay timers deal with the join that a router sends to override a prune sent by another router in its multi-access network.

For example, a network may have multiple routers serving VLAN 10. One of the routers on VLAN 10 receives multicast traffic that it does not need, so it sends a prune to the upstream neighbor. However, other routers in VLAN 10 may still need the traffic. When a router receives a prune for a group on that group's incoming interface, it knows that another router in the subnet is attempting to prune the connection. If the router still needs the traffic, it schedules a triggered join to the upstream neighbor to prevent the transmission from being disrupted.

The override delay interval (**override-interval** option) is the maximum amount of time that the router will wait before sending its join. Because other routers in the multi-access network may also send joins, the router selects a random time within the interval so that the upstream neighbor is not flooded with all the joins at the same time. The **propagation-delay** option sets the amount of time that the router expects the join will take to reach the neighbor.

When a router receives a prune for a group on an interface that connects to a multi-access network, it schedules that interface for deletion from the appropriate outgoing interface list. The router does not immediately prune the interface so that other routers in the multi-access network can have a chance to override the prune. The period that the router waits for overrides before

pruning the interface is determined by the sum of the override timer and the propagation delay. Take care in altering these timers; they should match on all neighboring routers so that one router does not delete an entry too soon.

Configuration Examples

This section guides you through the process of configuring PIM-SM in several simplified scenarios.

Example 1: Configuring PIM-SM in a Network with a Headquarters and Two Small Remote Sites

The network in Figure 13-12 includes a headquarters with many routers (only three are shown for simplicity), including a WAN router that connects to two remote sites in a Frame Relay network. The remote sites each include only one router which serves both the LANs and provides the Frame Relay connection to the headquarters.

The multicast sources are located at the headquarters site. All hosts in at the headquarters and the remote sites may potentially need to receive multicasts.

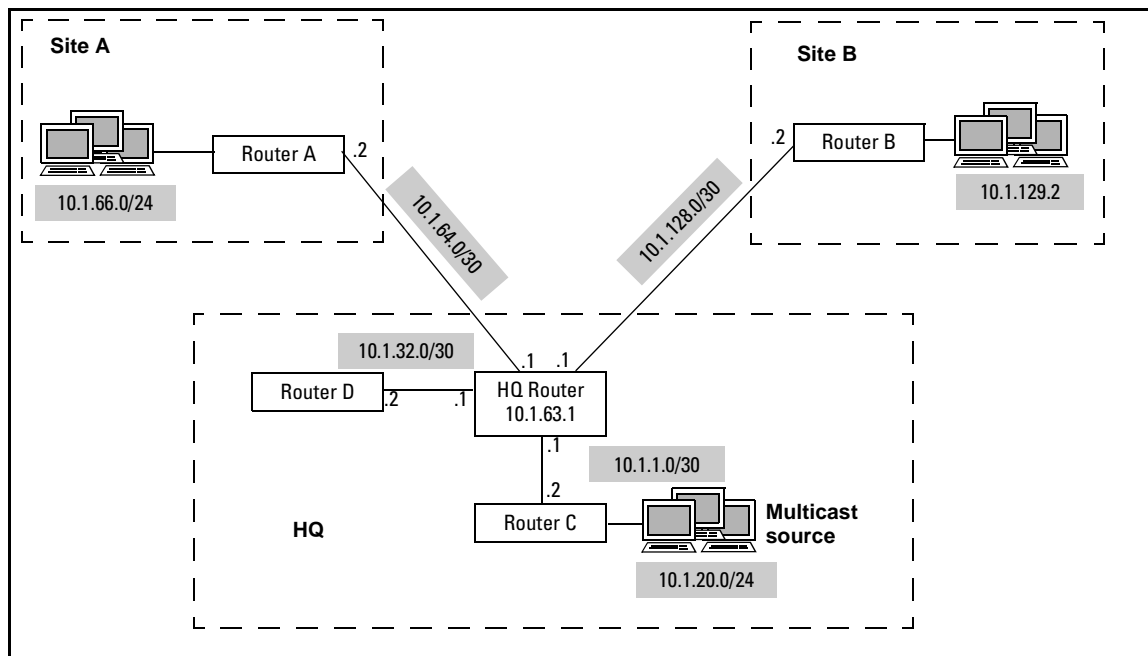


Figure 13-12.Example 1 Network

You should configure PIM-SM on each router interface in the network. Because all sources are at the headquarters, you decide to configure the HQ WAN router as the single RP.

Figure 13-13 shows the running-config for the HQ WAN router (showing only the sections of the configuration necessary for PIM-SM). You would follow these steps to produce this configuration:

1. Enable multicast routing:

```
HQRouter(config)# ip multicast-routing
```

2. Configure the LAN interfaces:

```
HQRouter(config)# interface eth 0/1
HQRouter(config-eth 0/1)# ip address 10.1.1.1 /24
HQRouter(config-eth 0/1)# ip pim sparse-mode
HQRouter(config-eth 0/1)# interface eth 0/2
HQRouter(config-eth 0/2)# ip address 10.1.2.1 /24
HQRouter(config-eth 0/2)# ip pim sparse-mode
```

3. Configure the WAN carrier-lines and the Frame Relay connection. The HQ WAN router uses two carrier-lines to connect to the Frame Relay network:

```
HQRouter(config)# interface t1 1/1
HQRouter(config-t1 1/1)# tdm-group 1 timeslots 1-24
HQRouter(config-t1 1/1)# interface t1 1/2
HQRouter(config-t1 1/2)# tdm-group 1 timeslots 1-24
HQRouter(config-t1 1/2)# interface frame-relay 1
HQRouter(config-fr 1)# frame-relay multilink
HQRouter(config-fr 1)# bind 1 t1 1/1 1 frame-relay 1
HQRouter(config-fr 1)# bind 2 t1 1/2 1 frame-relay 1
```

4. Add a PVC to Site A and enable PIM-SM on this connection:

```
HQRouter(config-fr 1)# interface frame-relay 1.101
HQRouter(config-fr 1.101)# frame-relay interface-dlci 101
HQRouter(config-fr 1.101)# ip address 10.1.64.1 /30
HQRouter(config-fr 1.101)# ip pim sparse-mode
```

5. Add a PVC to Site B and enable PIM-SM on this connection:

```
HQRouter(config-fr 1.101)# interface frame-relay 1.102
HQRouter(config-fr 1.102)# frame-relay interface-dlci 102
HQRouter(config-fr 1.102)# ip address 10.1.128.1 /30
HQRouter(config-fr 1.102)# ip pim sparse-mode
```

6. Configure a routing protocol. In this example, the network uses OSPF. The headquarters is the network backbone (area 0), Site A is stub area 1, and Site B is stub area 2. Note that routers in these areas receive summaries for inter-area traffic, not a default route. This is necessary so that these routers can use RPF to determine upstream neighbors for their multicast trees. Since the HQ WAN router is the ABR for these areas, you should configure it to advertise the area summaries:

```
HQRouter(config)# interface loopback 1
HQRouter(config-loop 1)# ip address 10.1.63.1 /24
HQRouter(config-loop 1)# router ospf
HQRouter(config-ospf)# network 10.1.1.0 0.0.0.255 area 0
HQRouter(config-ospf)# network 10.1.2.0 0.0.0.255 area 0
HQRouter(config-ospf)# network 10.1.64.0 0.0.0.3 area 1
HQRouter(config-ospf)# network 10.1.128.0 0.0.0.3 area 2
HQRouter(config-ospf)# area 0 range 10.1.0.0 255.255.192.0 advertise
HQRouter(config-ospf)# area 1 stub
HQRouter(config-ospf)# area 1 range 10.1.64.0 255.255.192.0 advertise
HQRouter(config-ospf)# area 2 stub
HQRouter(config-ospf)# area 2 range 10.1.128.0 255.255.192.0 advertise
```

7. Configure this router as the RP:

```
HQRouter(config)# router pim-sparse
HQRouter(config-pim-sparse)# rp-address 10.1.63.1
```

8. Activate all interfaces and connect the router to the network:

```
HQRouter(config-eth 0/1)# no shutdown
HQRouter(config-eth 0/2)# no shutdown
HQRouter(config-t1 1/1)# no shutdown
HQRouter(config-t1 1/2)# no shutdown
HQRouter(config-fr 1)# no shutdown
```

```
hostname "HQRouter"  
ip multicast-routing  
interface loop 1  
  ip address 10.1.63.1 255.255.255.0  
  no shutdown  
interface eth 0/1  
  ip address 10.1.1.1 255.255.255.0  
  ip pim sparse-mode  
  no shutdown  
interface eth 0/2  
  ip address 10.1.32.1 255.255.255.0  
  ip pim sparse-mode  
  no shutdown  
interface t1 1/1  
  tdm-group 1 timeslots 1-24 speed 64  
  no shutdown  
interface t1 1/2  
  tdm-group 1 timeslots 1-24 speed 64  
  no shutdown  
interface fr 1 point-to-point  
  frame-relay lmi-type ansi  
  frame-relay multilink  
  no shutdown  
  bind 1 t1 1/1 1 frame-relay 1  
  bind 2 t1 1/2 1 frame-relay 1  
interface fr 1.101 point-to-point  
  frame-relay interface-dlci 101  
  ip address 10.1.64.1 255.255.255.252  
  ip pim sparse-mode  
interface fr 1.102 point-to-point  
  frame-relay interface-dlci 102  
  ip address 10.1.128.1 255.255.255.252  
  ip pim sparse-mode  
router ospf  
  network 10.1.1.0 0.0.0.255 area 0  
  network 10.1.2.0 0.0.0.255 area 0  
  network 10.1.63.0 0.0.0.255 area 0  
  network 10.1.64.0 0.0.0.3 area 1  
  network 10.1.128.0 0.0.0.3 area 2  
  area 0 range 10.1.0.0 255.255.192.0 advertise  
  area 1 stub  
  area 1 range 10.1.64.0 255.255.192.0 advertise  
  area 2 stub  
  area 2 range 10.1.128.0 255.255.192.0 advertise  
!  
router pim-sparse  
  rp-address 10.1.63.1  
!
```

Figure 13-13. Example 1: HQ WAN Router running-config

You would need to make the same configurations on the WAN routers at Site A and Site B. Figure 13-14 shows the running-config for the Router at Site A.

```
hostname "RouterA"
ip multicast-routing
interface loop 1
  ip address 10.1.66.10 255.255.255.0
  no shutdown
interface eth 0/1
  ip address 10.1.65.1 255.255.255.0
  ip pim sparse-mode
  no shutdown
interface t1 1/1
  tdm-group 1 timeslots 1-24 speed 64
  no shutdown
interface fr 1 point-to-point
  frame-relay lmi-type ansi
  no shutdown
  bind 1 t1 1/1 1 frame-relay 1
interface fr 1.100 point-to-point
  frame-relay interface-dlci 100
  ip address 10.1.64.2 255.255.255.252
  ip pim sparse-mode
router ospf
  network 10.1.64.0 0.0.0.3 area 1
  network 10.1.65.0 0.0.0.255 area 1
  network 10.1.66.0 0.0.0.255 area 1
  area 1 stub
!
router pim-sparse
  rp-address 10.1.63.1
!
```

Figure 13-14. Example 1: Router A running-config

If you want to configure a default route at the stub site instead of a routing protocol, then you should enable IGMP proxy instead of PIM-SM. The running-config would be as shown in Figure 13-15.

```
hostname "RouterA"
ip multicast-routing
ip mcast-stub helper-address 10.1.64.1
interface eth 0/1
 ip address 10.1.65.1 255.255.255.0
 ip mcast-stub downstream
 ip mcast-stub helper-enable
 no shutdown
interface t1 1/1
 tdm-group 1 timeslots 1-24 speed 64
 no shutdown
interface fr 1 point-to-point
 frame-relay lmi-type ansi
 no shutdown
 bind 1 t1 1/1 1 frame-relay 1
interface fr 1.100 point-to-point
 frame-relay interface-dlci 100
 ip address 10.1.64.2 255.255.255.252
 ip mcast-stub upstream
ip route 0.0.0.0 0.0.0.0 fr 1.100
end
```

Figure 13-15. Example 1: Router A running-config (IGMP Proxy)

Example 2: Configuring Specific RPs to Support Specific Groups

The network in Figure 13-16 is similar to that in Example 1. However, Site A in this network includes a source that multicasts to 239.255.1.1. Because most hosts for this group will be in the local area, administrators decide to have Router A act as the RP for this group.

Administrators also decide that the headquarters should have two routers that can act as RP. Now, if one of the routers fails, hosts at the headquarters can still receive multicasts. Router D will only support traffic in the 224.0.0.0 /24 range, which enables backup RP support for necessary multicast protocols, but not for privately defined multicast groups.

Administrators will configure an RP set on each router that includes all three of the routers selected to act as RPs.

Note

Router D is not a backup RP in the sense that it will only become an RP if the HQ Router fails. Router D will be the RP for any groups for which the PIM-SM algorithm selects it.

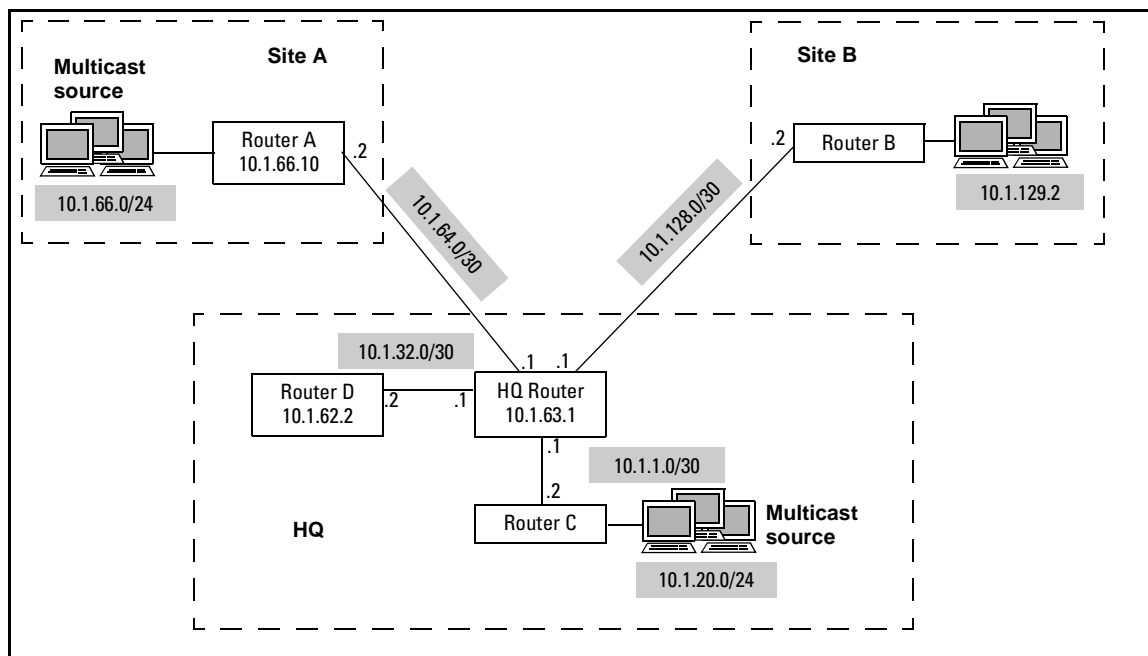


Figure 13-16. Example 2 Network

To configure the HQ WAN router, you would follow these steps:

1. Follow steps 1 through 6 described in Example 1 to configure all router interfaces, to enable Layer 2 interfaces to run PIM-SM, and to configure the routing protocol. (See the running-config in Figure 13-13 on page 13-43.)

2. Configure the ACLs that list the groups supported by each RP:

- a. Configure an ACL that permits RP 10.1.63.1 (the HQ Router) to support any multicast group except 239.255.255.1:

```
HQRouter(config)# ip access-list standard rp1
HQRouter(config-std-nacl)# deny host 239.255.255.1
HQRouter(config-std-nacl)# permit 224.0.0.0 15.255.255.255
```

- b. Configure an ACL that permits Router D (10.1.62.4) to support the multicast addresses used by routing protocols:

```
HQRouter(config)# ip access-list standard rp2
HQRouter(config-std-nacl)# permit 224.0.0.0 0.0.0.255
```

- c. The LAN at Site A supports a multicast server transmitting to 239.255.255.1. Configure an ACL that permits Router A (10.1.66.10) to support only this multicast group:

```
HQRouter(config)# ip access-list standard rp3
HQRouter(config-std-nacl)# permit host 239.255.255.1
```

3. Configure the RP set:

```
HQRouter(config)# router pim-sparse
HQRouter(config-pim-sparse)# rp-address 10.1.63.1 access-group rp1
HQRouter(config-pim-sparse)# rp-address 10.1.62.4 access-group rp2
HQRouter(config-pim-sparse)# rp-address 10.1.66.10 access-group rp3
```

4. Activate all interfaces and connect the router to the network.

```
hostname "HQRouter"
!
router pim-sparse
  rp-address 10.1.63.1 access-group rp1
  rp-address 10.1.62.4 access-group rp2
  rp-address 10.1.66.10 access-group rp3
!
ip access-list standard rp1
  deny host 239.255.255.1
  permit 224.0.0.0 15.255.255.255
ip access-list standard rp2
  permit 224.0.0.0 0.0.0.255
ip access-list standard rp3
  permit host 239.255.255.1
```

Figure 13-17. Example 2: running-config for the RP Set on HQ Router

Troubleshooting PIM-SM

When hosts are not receiving multicasts, you must determine where the traffic is going astray. Because PIM-SM relies on unidirectional trees, you should first troubleshoot the router that directly connects to the hosts, then proceed to the next hop upstream router until you find the point at which the traffic is disrupted.

Monitoring the Multicast Routing Table

Having a solid understanding of the information that a multicast routing table can provide will help you to troubleshoot PIM-SM functions. When you understand how to read the table, you can view it to discover such things as:

- for which multicast groups the router has received PIM or IGMP joins
- through which interfaces the router receives and forwards specific types of multicast traffic
- the IP addresses of upstream RP and SP neighbors for specific groups
- whether the router has joined an SP tree for a group
- whether the router is acting as the DR for a source
- how long specific entries have been active and when they will expire

If you need to review any of these concepts, see the “Overview” on page 13-3.

Use this enable mode command to view the multicast routing table:

```
ProCurve# show ip mroute
```



```

ProCurve# show ip mroute
IP Multicast Routing Table
Legend for entry flags:  S - Sparse, C - Connected, P - Pruned, J - Join SPT, T - SPT-bit Set,
                        F - Register, R - RP-bit Set
                        Timers: Uptime/Expires

(*, G) entry for the RP tree:  (*, 239.255.255.1), 01:06:23/00:00:00, RP 10.1.1.1, Flags: SCJ
                                Incoming interface: Null, RPF nbr 0.0.0.0
                                Outgoing interface list:
                                    eth 0/1, Forward, 00:12:08/00:03:21
                                    fr 1.2, Forward, 00:12:54/00:01:26
                                    fr 1.3, Forward, 00:12:12/00:02:09

(S, G) entry for the SP tree:  (192.168.100.2, 239.255.255.1), 00:12:10/00:03:30, Flags: ST
                                Incoming interface: fr 1.3, RPF nbr 10.3.3.2
                                Outgoing interface list:
                                    eth 0/1, Forward, 00:2:20/00:03:21
                                    fr 1.2, Forward, 00:02:09/00:00:10
    
```

Figure 13-18. Monitoring the Multicast Routing Table

Flags

The first section of the multicast routing table provides a legend for various flags that can be set on entries. (See Figure 13-18.) Table 13-3 provides a more in-depth explanation of the flags that the simple legend does.

Table 13-3. PIM-SM Entry Flags

Flag	Name	Meaning	Valid for Entry Type
S	Sparse	The entry is a sparse mode entry.	<ul style="list-style-type: none"> (*, G) (S, G)
C	Connected	The router connects directly to a host that is a member of this group.	<ul style="list-style-type: none"> (*, G) (S, G)
P	Pruned	The outgoing interface list is null, so the router will send a prune to the upstream RP neighbor, for a (*, G) entry, or SP neighbor, for an (S, G) entry.	<ul style="list-style-type: none"> (*, G) (S, G)

Flag	Name	Meaning	Valid for Entry Type
J	Join SPT	<ul style="list-style-type: none"> For a (*, G) entry on an RP, the RP will generate an SP tree for group traffic immediately after a source for that group registers with it. For a (*, G) entry on a non-RP PIM router, the STP threshold has been exceeded. If this router is an edge router for the group (as indicated by a C flag in the entry), it will switch to the SP tree when it receives the next multicast over the RP tree. That is, the router will create an (S, G) entry and send a join to its upstream SP neighbor. For an (S, G) entry, you may see this flag in an (S, G) RP-bit entry with no outgoing interfaces. The J-bit indicates that, if the router receives a join for this group, it should join the SP tree. 	<ul style="list-style-type: none"> (*, G) (S, G)
T	SPT-bit set	<p>The router is receiving and forwarding multicast traffic using this source-specific entry.</p> <p>After a router receives the first multicast packet over the SP tree, it sets this bit to indicate that the SP tree is active. If the upstream RP neighbor for this group is different from the upstream SP neighbor, the router also sends an (S, G) RP-bit prune to the upstream RP neighbor.</p>	(S, G)
F	Register	The router must send register packets to the RP for this group. The DR sets this bit in its (S, G) entry when it receives multicasts from a directly connected source.	(S, G)
R	RP-bit set	<p>The RP bit indicates that an (S, G) entry relates to the RP tree. When a router sets an entry's RP-bit, it recalculates the entry's incoming interface and RPF neighbor. The RPF neighbor is now the upstream RP neighbor.</p> <p>An (S, G) entry with the R flag is used to prune devices that are using a divergent SP tree from the RP tree.</p>	(S, G)

First Line of a Multicast Routing Table Entry

Figure 13-19 displays the information that you can view in the first line of a multicast routing table entry.

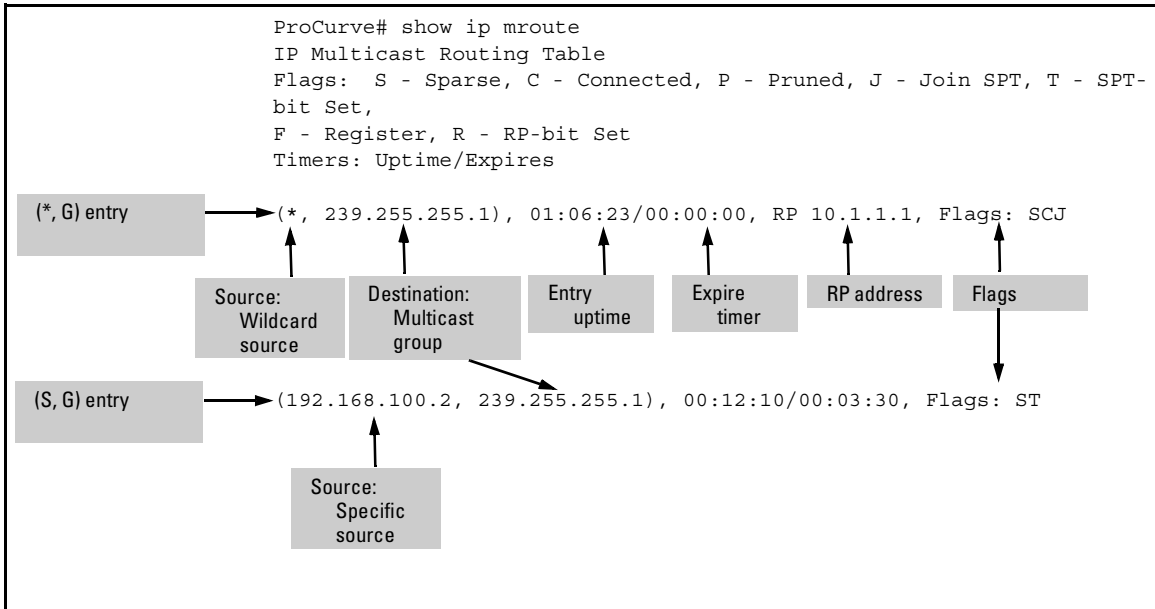


Figure 13-19. First Line of a Multicast Routing Table Entry

The router matches the source and destination of an incoming multicast packet’s IP header to the source and destination listed in an entry. The destination is the multicast group address. A (*, G) entry includes a wildcard source (*) and matches any source address. The router looks for the most specific match, so it always matches a multicast to the (S, G) entry if such an entry exists for its source.

The entry uptime indicates the time since the entry was created. The (*, G) entry in Figure 13-19 has been active for more than an hour.

The expire timer indicates the number of hours, minutes, and seconds left before the entry will be deleted. The router resets this timer every time an interface receives a new join for the group. Routers should send periodic join/prune messages at intervals of about one third the expire timer so that the entry is not constantly deleted and recreated.

In a (*, G) entry, the RP address indicates the device toward which the router should forward its join.

An entry can have more than one flag set at a time. Table 13-4 shows flag settings for some typical multicast routing table entries.

Table 13-4. Flags in Typical Multicast Routing Table Entries

Flags	Meaning
(*, G) entry SC	The router is an edge router for this group.
(*, G) entry SJ	Typically, the router is RP for this group.
(*, G) entry SCJ	Typically, the router is RP for this group, and it also connects directly to hosts that are members of this group.
(S, G) entry ST	The router is receiving and forwarding multicast traffic on an SP tree. If the router is the DR for a directly connected source, the RP has already joined the SP tree and the DR is no longer sends the register packets.
(S, G) entry SJR	The router is receiving multicast traffic on the SP tree, but this tree diverges from the RP tree that the router was originally using. One or more or all of the outgoing interfaces in the corresponding (*, G) entry may be absent from this entry because neighbors are using different connections in the SP tree.
(S, G) entry SFT	The router is DR for a directly connected source. It is currently registering with the RP.

Incoming Interface

The router determines a multicast route's incoming interface and upstream neighbor using RPF. For a (*, G) entry, the router looks up the best route to the RP in its unicast routing table. For an (S, G) entry, the router looks up the best route to the source. The router lists the forwarding interface indicated in the best unicast route as the multicast route's incoming interface. The router enters the next hop address of the unicast route for the address of the multicast route's RPF neighbor. (See Figure 13-20 to learn how to find the incoming interface and RPF neighbor for an entry.)

The incoming interface is very important for multicast routing with PIM-SM because multicast traffic *must* follow a unidirectional flow. If a multicast traffic does not arrive on the interface indicated in the matching entry, the router discards it. Therefore, if an edge router has a null incoming interface for an entry, it will not be able to receive multicast traffic for that entry.

A router should only have a (*, G) entry with a null incoming interface if it is the RP for that group. The RP will receive a register packet from the DR as a unicast, not as a multicast, so it will not discard it. The RP then will create a (S, G) entry for which the incoming interface connects to the source address for the encapsulated multicast.

A router should never have an (S, G) entry without an incoming interface.

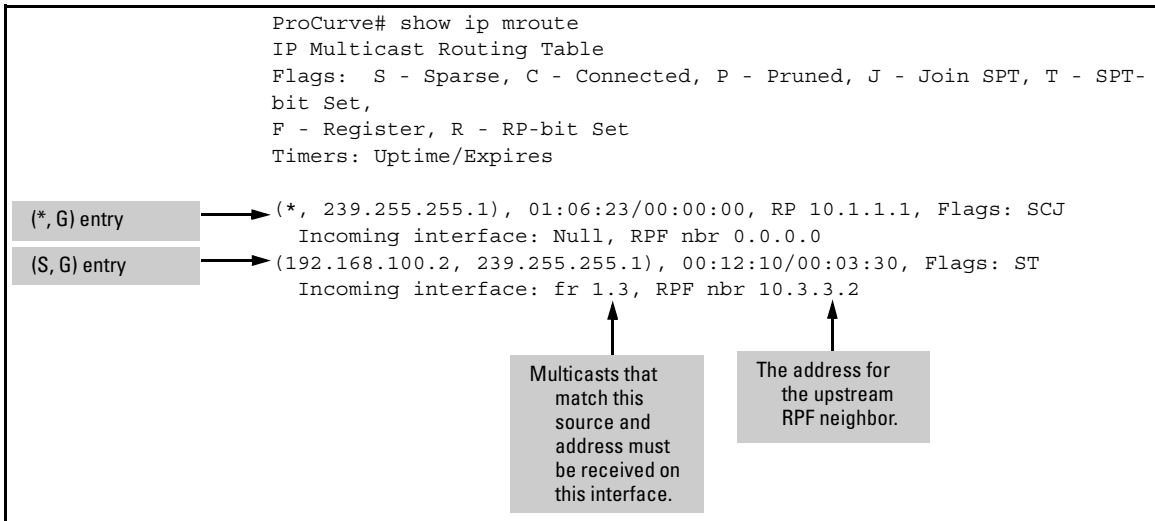


Figure 13-20. Incoming Interface for a Multicast Route

Outgoing Interface List

You can view the outgoing interface list for an entry to determine to which neighbors the router is forwarding multicast traffic. (See Figure 13-21.)

The outgoing interface list for a (*, G) entry deals with the RP tree. It displays the interfaces through which the router forwards the multicasts that it receives from its upstream RP neighbor. An RP forwards decapsulated register packets to the outgoing interfaces.

The outgoing interface list for an (S, G) entry deals with the SP tree. It displays the interfaces through which the router must forward multicast packets sent by a specific source. When a router creates a (S, G) entry, it copies the outgoing interface list from the corresponding (*, G) entry. In this way, downstream neighbors in the process of switching to the SP tree themselves continue to receive traffic.

Note that the router will discard any multicast packets from this source that do not arrive on the incoming interface listed for the entry—for example, packets received from an RP neighbor. This means that once a router has created an (S, G) entry for a multicast group and set its STP-bit, it will only accept and send traffic for this group over the SP tree.

The outgoing interface list for an (S, G) RP-bit entry includes the interfaces that connect to routers who have not joined an SP tree and still need multicasts from the shared RP tree. (See Figure 13-21.)

```
ProCurve# show ip mroute
IP Multicast Routing Table
Flags: S - Sparse, C - Connected, P - Pruned, J - Join SPT, T - SPT-
bit Set,
F - Register, R - RP-bit Set
Timers: Uptime/Expires

(*, 239.255.255.1), 01:06:23/00:00:00, RP 10.1.1.1, Flags: SCJ
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
eth 0/1, Forward, 00:12:08/00:03:21
fr 1.2, Forward, 00:12:54/00:01:26
fr 1.3, Forward, 00:12:12/00:02:09

(192.168.100.2, 239.255.255.1), 00:12:10/00:03:30, Flags: ST
Incoming interface: fr 1.3, RPF nbr 10.3.3.2
Outgoing interface list:
eth 0/1, Forward, 00:2:08/00:02:24
fr 1.2, Forward, 00:02:09/00:00:10
```

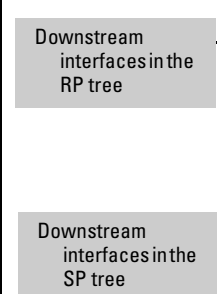


Figure 13-21. Outgoing Interface List

Viewing PIM-SM Information

You can use the commands displayed in Table 13-5 to view information about PIM-SM functions on the router, including:

- which local interfaces implement PIM-SM
- the address of the RP for a group
- the addresses of PIM neighbors
- values for PIM-SM timers
- the number of PIM packets sent and received, including join/prunes and register packets

Table 13-5. PIM-SM show Commands

View	Command Syntax
<ul style="list-style-type: none"> intervals for sending join/prune packets SPT threshold 	show ip pim-sparse
interfaces running PIM: <ul style="list-style-type: none"> interface status DR for the interface's network interface IP address timer values 	show ip pim-sparse interface
a specific interface that is running PIM	show ip pim-sparse interface <interface ID>
routers that can potentially become RPs: <ul style="list-style-type: none"> router's IP address the ACL that selects which group addresses are supported 	show ip pim-sparse rp-set
the router that is currently RP for each active group	show ip pim-sparse rp-map
PIM routers that connect to the local router	show ip pim-sparse neighbor
PIM packets sent and received: <ul style="list-style-type: none"> hello join/prune register register-stop assert 	show ip pim-sparse traffic

You can also debug PIM-SM to monitor PIM-SM events as they occur. For example, if an upstream router is not forwarding multicasts to your router, you can view debug messages to discover whether the local router is sending the proper joins. Use the commands shown in Table 13-6.

Note

PIM debug messages can be draining on the router's processor.

Table 13-6. PIM-SM debug Commands

View	Command Syntax
all messages	debug ip pim-sparse
assert messages	debug ip pim-sparse assert
hellos	debug ip pim-sparse hello
PIM join and prunes	debug ip pim-sparse joinprune
detailed information in PIM messages registers and register-stops	debug ip pim-sparse packets debug ip pim-sparse register

PIM-SM Troubleshooting Process

When local hosts are not receiving multicasts, you should start troubleshooting beginning with the edge router and proceeding up the tree.

Troubleshooting an Edge Router

You should verify the following for the edge router:

- the local router has an IGMP group membership for the multicasts in question (**show ip igmp group**)
- the local router has a multicast route for that group (**show ip mroute**)
- the outgoing interface list for the multicast entry includes the interface that connects to the network in question (**show ip mroute**)
- the incoming interface for the multicast entry is not null and connects to the correct upstream neighbor (**show ip mroute**)
- the router is either implementing a routing protocol or has all necessary static routes (**show ip route**)
- the router is sending PIM joins for the group to the upstream neighbor (**show ip pim-sparse neighbor**, **show ip pim-sparse traffic**, **debug ip pim-sparse joinprune**)

To troubleshoot the edge router, follow these steps:

1. Determine whether the local router has received IGMP joins from the hosts. Enter the following command to view the groups for which the router has received an IGMP report:

```
ProCurve# show ip igmp group
```


2. If you see the group that you are troubleshooting in the list of group memberships, move to step 3.

If the list of group memberships does not include necessary groups, then you must troubleshoot IGMP. Remember that you should enable PIM on LAN interfaces in order for those interfaces to run IGMP. It is also possible that the problem lies with the host. See *Chapter 12: Configuring Multicast Support for a Stub Network* for more information on troubleshooting IGMP.

3. View the multicast routing table by entering this command:

```
ProCurve# show ip mroute
```

The table should include an entry for each multicast group listed in the router's IGMP group memberships. The entry may be a (*, G) entry, or, if the router has already received multicast traffic for this group, the entry may be an (S, G) entry.

Figure 13-22 displays an example of a multicast routing table entry.

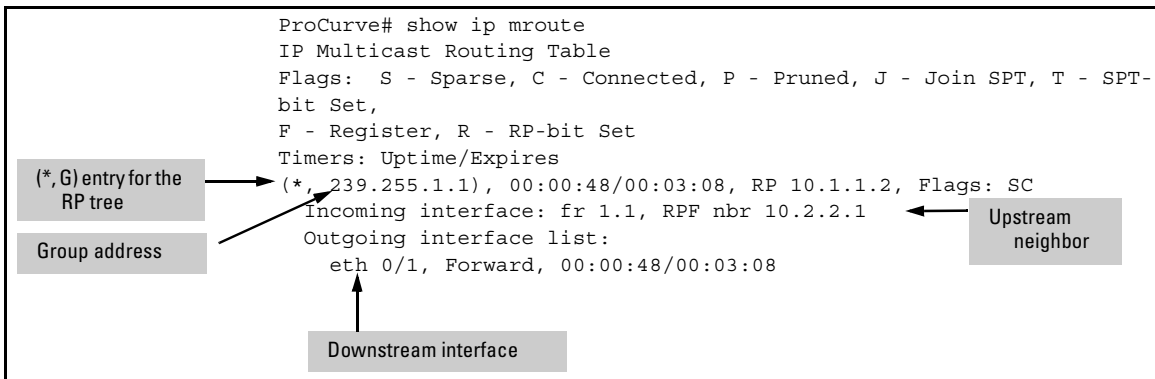


Figure 13-22. Multicast Routing Table Entry

4. If the multicast routing table does not include any entries, then the router may not be implementing PIM-SM. Enter the following command to verify that all the local interfaces that should be running PIM-SM are actually doing so:

```
ProCurve# show ip pim-sparse interfaces
```

View the DR for the network that is experiencing the problems. If the local router is not the DR, it will not forward multicasts for network. You should troubleshoot the DR. If an interface is down, then you must, of course, bring it up before it can rejoin the PIM domain.

5. If the multicast routing table does have an entry for the group in question, view the list of outgoing interfaces in this entry. If the local interface that connects to the network experiencing the problems is not in this list, then the router will not forward multicasts to it.

Note

The multicast routing table may include both a (*, G) and an (S, G) entry for the group. Search the outgoing interface list of either entry.

Typically, an interface listed for an IGMP group should be in the outgoing interface list for the multicast routing table entry for that group. One reason that interface would not be in the list is that another router in a multi-access network is responsible for forwarding multicasts for that group. The local router has lost the assert.

Figure 13-23 shows the multicast routing table for a router in this situation.

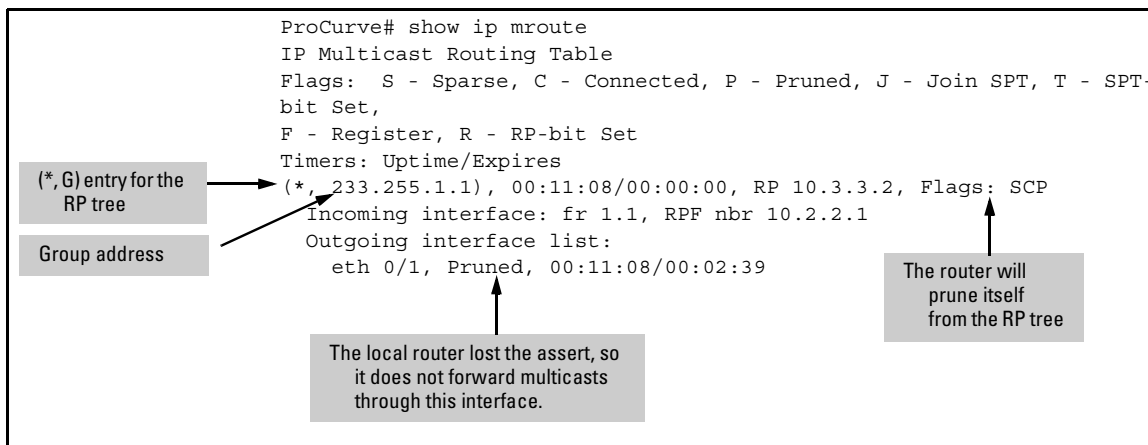


Figure 13-23. Interface Pruned From a Multicast Routing Table Entry

6. If the incoming interface for the multicast routing table entry is null, the router will drop any multicasts it receives for that group. The incoming interface is legitimately null when the local router is the RP for the group. In this case, however, as soon as the group's source begins to transmit multicasts, the router should add an (S, G) entry with the incoming interface set to the interface connected to the source. A (*, G) entry in an RP's multicast routing table should be marked with the J bit to indicate that the RP is allowed to join an SP tree. (See Figure 13-24).

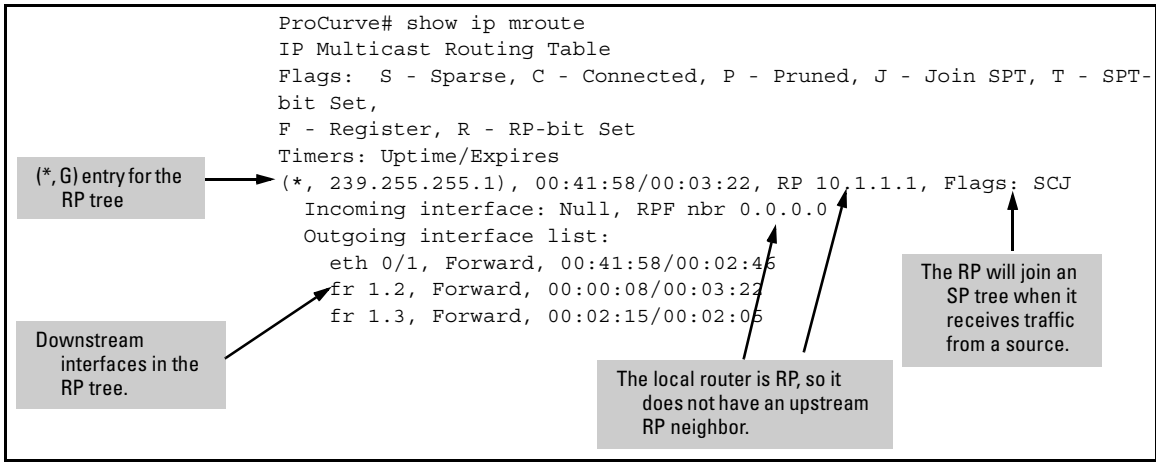


Figure 13-24. Multicast Route for a Group for Which the Local Router is RP

The incoming interface is also null when the router cannot determine the upstream neighbor for the RP or source. Figure 13-25 illustrates an example of this problem.

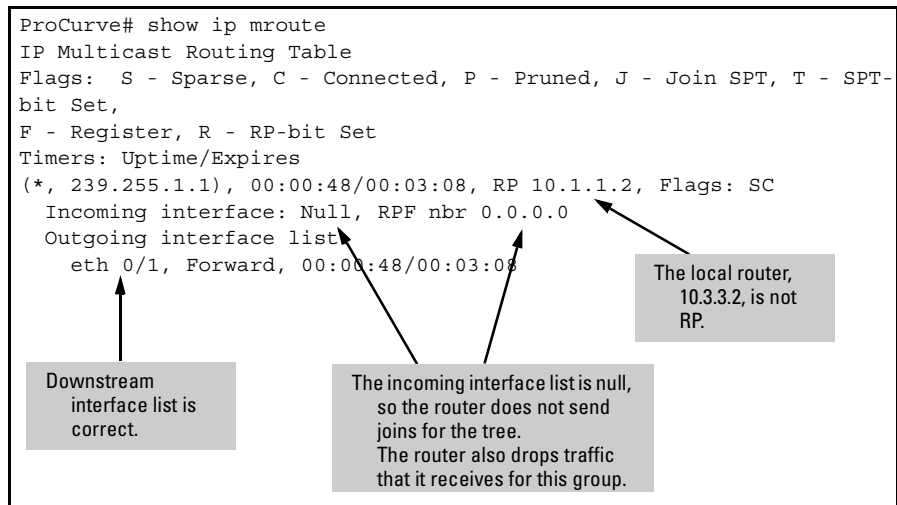


Figure 13-25. Null Incoming Interface List

View the router’s unicast routing table as shown in Figure 13-26:

```
ProCurve# show ip route
```

This table must include an explicit route to the RP or source (depending on the type of entry) in order for the router to determine the incoming interface for an multicast entry. You must either enable a routing protocol on the router or configure a static route to each RP and network that may include a multicast source. If your router runs OSPF, it must receive summary routes, not simply a default route. See *Chapter 15: IP Routing—Configuring RIP, OSPF, BGP, and PBR* for more information.

```
ProCurve# show ip route
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
       IA - OSPF inter area, N1 - OSPF NSSA external type 1
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1
       E2 - OSPF external type 2

Gateway of last resort is 10.2.2.2 to network 0.0.0.0

Default route is not sufficient. → O IA 0.0.0.0/0 [110/50] via 10.2.2.2, ppp 1
C   10.2.2.0/30 is directly connected, ppp 1
C   10.2.2.2/32 is directly connected, ppp 1
C   192.168.64.0/24 is directly connected, eth 0/1
C   192.168.65.0/24 is directly connected, eth 0/2

↑
The RP is 10.1.1.2. (See Figure 13-24.) The table does not include an explicit route to this address.
```

Figure 13-26. Looking for a Route to the RP in the Unicast Routing Table

7. If the incoming and outgoing interfaces for the multicast routing table entry seem correct, then you should verify that the router is sending joins to its upstream neighbors as it should. You can determine whether the router has established contact with a neighbor using this command:

```
ProCurve# show ip pim-sparse neighbor
```

If a neighbor is not present in the list, the connection to that neighbor may be down. Troubleshoot a WAN connection as described in *Chapter 6: Configuring the Data Link Layer Protocol for E1, T1, and Serial Interfaces* in the *Basic Management and Configuration Guide*. If the connection is good, you can try raising the neighbor timeout or returning it to its default.

You can also enter **show ip pim-sparse traffic** to verify that the router is sending join/prune messages. If you want to see the actual messages being sent then you must use the **debug ip pim-sparse joinprune** command as shown in Figure 13-27.

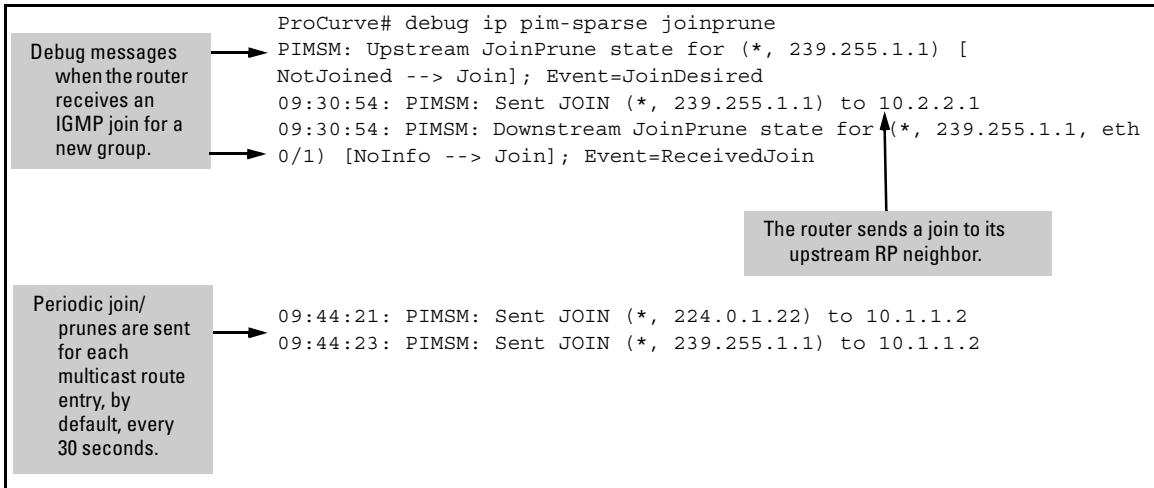


Figure 13-27. Viewing Real-Time Join/Prune Messages

If the router does not seem to be sending messages, try returning the join/prune interval to its default setting. Move to the PIM sparse configuration mode context and enter this command:

```
ProCurve(config-pim-sparse)# no join-prune-msg-interval
```

If you complete these steps and cannot find the problem, you must troubleshoot your router in conjunction with other routers in the domain.

Troubleshooting A Router in Conjunction with Its PIM Neighbors

Note

These instructions assume that you have access to all routers in the PIM domain. If you do not, then you may need to request information from a network administrator at a remote site.

You should first determine whether the local router and its neighbors have selected the same RP for the multicast group.

Troubleshooting RP Sets. When a router does not receive multicast traffic from its upstream neighbors, one of the most likely problems is that the local router and its upstream neighbors have incompatible RP sets. If neighbors select different RPs for a group, the upstream router ignores joins for that group from the downstream router.

Follow these steps to troubleshoot a router's RP set:

1. Compare the local router's RP map, in which the router stores the IP address of the RP for each active group, to its neighbors' maps. Each router should map each group to the same RP. To view the map, enter this enable mode command:

```
ProCurve# show ip pim-sparse rp-map
```

Figure 13-28 shows how to compare the output from this command on two routers.

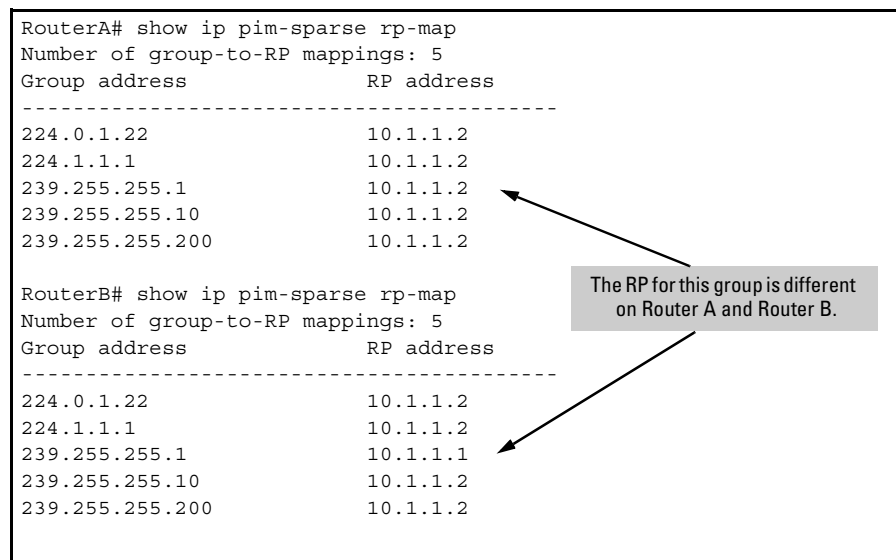


Figure 13-28. Comparing RP Maps

2. If, from router to router, you find discrepancies in the RPs selected, you should view routers' RP sets to find the source of the discrepancies. On a ProCurve Secure Router, the RP set is the manually configured pool of routers that can act as RP for various groups. Every router in the domain should have exactly the same RP set.

Enter this command from the CLI of the router that is using the wrong RP to view its RP set:

```
ProCurve# show ip pim-sparse rp-set
```

Compare this RP set to that configured on a neighboring router that has selected the correct RP.

Figure 13-29 shows RP sets on two neighboring routers in a PIM domain. Administrators in this domain have decided that all RPs should potentially support all multicast groups. In this example, Router A was unable to join the RP tree for group 239.255.255.1 because its RP set did not include the address for one of the routers capable of acting as RP.

```
RouterA# show ip pim-sparse rp-set
Group address      Static-RP-address
-----
224.0.0.0 /4      10.1.1.2

RouterB# show ip pim-sparse rp-set
Group address      Static-RP-address
-----
224.0.0.0 /4      10.1.1.1 ← Entry missing on Router A
224.0.0.0 /4      10.1.1.2
```

Figure 13-29. Viewing an RP Set

3. Comparing RP sets becomes more complicated when RPs have been configured to support specific groups. If possible, you should view the RP set and the ACLs indicated in that set at the same time (as shown in Figure 13-30). You can then determine which groups each RP supports.

```
RouterA# show ip pim-sparse rp-set
Group address      Static-RP-address
-----
rp1                10.1.1.1
rp2                10.1.1.2
rp3                10.3.3.2
RouterA# show access-lists
Standard IP access list rp1
  permit host 239.255.255.1 (1 matches)
Standard IP access list rp2
  deny  host 239.255.255.1 (1 matches)
  permit 224.0.0.0 15.255.255.255 (3 matches)
  permit 224.0.0.0 7.255.255.255 (0 matches)
Standard IP access list rp3
  deny  host 239.255.255.1 (0 matches)
  permit 232.0.0.0 7.255.255.255 (1 matches)
```

Figure 13-30. Viewing an RP Set That Associates Different RPs with Different Groups

In this example, Router 10.1.1.1 supports a single multicast group with the IP address 239.255.255.1. Router 10.1.1.2 supports all multicast groups except 239.255.255.1. Router 10.3.3.2 supports the second half of all multicast groups (232.0.0.0 through 239.255.255.255) except group 239.255.255.1.

The RP set for the upstream router is shown in Figure 13-31.

```
RouterB# show ip pim-sparse rp-set
Group address      Static-RP-address
-----
rp1                10.1.1.1
rp2                10.1.1.2
rp3                10.3.3.2
ProCurveSR7203dl# show access-lists
Standard IP access list rp1
  permit host 239.255.255.1 (1 matches)
Standard IP access list rp2
  deny  host 239.255.255.1 (1 matches)
  permit 224.0.0.0 7.255.255.255 (2 matches)
Standard IP access list rp3
  deny  host 239.255.255.1 (0 matches)
  permit 232.0.0.0 7.255.255.255 (2 matches)
```

This ACL is different than the ACL configured for Router 10.1.1.2 on Router A.

Figure 13-31. Comparing RP Sets

Note the difference in Router B's ACL for the RP at 10.1.1.2. On Router B, this RP only supports the half of all possible multicast groups (224.0.0.0 through 231.255.255.255) rather than all of the groups. Figure 13-32 shows which RPs Router A and B have actually selected for each active group.

```
RouterA# show ip pim-sparse rp-map
Number of group-to-RP mappings: 5
Group address          RP address
-----
224.0.1.22            10.1.1.2
224.1.1.1             10.1.1.2
239.255.255.1        10.1.1.1
239.255.255.10       10.3.3.2
239.255.255.200      10.1.1.2

RouterB# show ip pim-sparse rp-map
Number of group-to-RP mappings: 5
Group address          RP address
-----
224.0.1.22            10.1.1.2
224.1.1.1             10.1.1.2
239.255.255.1        10.1.1.1
239.255.255.10       10.3.3.2
239.255.255.200      10.3.3.2
```

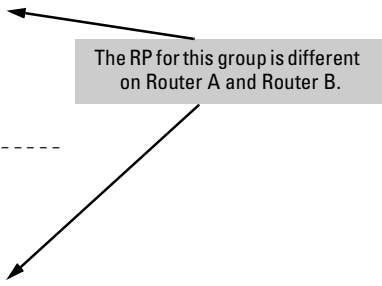


Figure 13-32. Viewing Active RPs in an RP Map

Because Router A allows RP 10.1.1.2 to support all groups, Router A has selected this RP for multicast group 239.255.255.200. Router B, which has an RP set that prohibits RP 10.1.1.2 from supporting this group, selected RP 10.3.3.2. You can view the access lists configured on Router A once again to find the error. (See Figure 13-33.)

In this example, the administrator who configured the ACL for RP 10.1.1.2 accidentally made the RP support all groups. The administrator then attempted to correct the error, but forgot to remove the faulty permit statement. (See Figure 13-33.)

```
RouterA# show ip pim-sparse rp-set
Group address      Static-RP-address
-----
rp1                10.1.1.1
rp2                10.1.1.2
rp3                10.3.3.2
RouterA# show access-lists
Standard IP access list rp1
  permit host 239.255.255.1 (1 matches)
Standard IP access list rp2
  deny host 239.255.255.1 (1 matches)
  permit 224.0.0.0 15.255.255.255 (3 matches)
  permit 224.0.0.0 7.255.255.255 (0 matches)
Standard IP access list rp3
  deny host 239.255.255.1 (0 matches)
  permit 232.0.0.0 7.255.255.255 (1 matches)
```

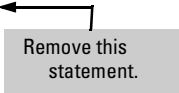


Figure 13-33. Fixing a Misconfigured ACL

To fix the ACL so that Router A selects the correct RP for multicast group 239.255.255.200, you would enter these commands:

```
ProCurve(config)# ip access-list standard rp2
ProCurve(config-std-nacl)# no permit 224.0.0.0 15.255.255.255
```

Figure 13-34 shows another misconfigured RP set. The PIM domain in this example uses the RP at 10.1.1.2 to support half of the multicast groups and the RP at 10.3.3.2 to support the others. The administrator for Router A used extended ACLs instead of standard ACLs and entered the group address as a destination rather than a source. However, routers use the source address in an RP's ACL to determine which groups the RP supports.

```

RouterA# show ip pim-sparse rp-set
Group address      Static-RP-address
-----
rp1                10.1.1.2
rp2                10.3.3.2
RouterA# show access-lists
Extended IP access list rp1
  permit ip any 224.0.0.0 7.255.255.255 (0 matches)
Extended IP access list rp2
  permit ip any 232.0.0.0 7.255.255.255 (1 matches)
  
```

The IP address for the multicast host address should be in the *source* position.

Figure 13-34. Misconfigured Extended ACL for an RP Set

Router A processes the **any** keyword in the source position and allows both routers to support all groups. Router A will sometimes select the same RP as its neighbors, but it will sometimes select the wrong one.

You should use *standard* ACLs to configure your router’s RP set to avoid this problem.

Other Tips for Troubleshooting an Upstream Router. If incompatible RP sets are not the problem, then you should troubleshoot the upstream router in the much same way that you would troubleshoot an edge router. (See “Troubleshooting an Edge Router” on page 13-56.) Verify that:

- the router has a multicast route for that group (**show ip mroute**) and that the interface that connects to the downstream router is up and running PIM (**show ip pim-sparse interfaces**)
- the outgoing interface list for the multicast entry includes the interface that connects to the downstream router (**show ip mroute**)
- the incoming interface for the multicast entry is not null and connects to the correct upstream neighbor (**show ip mroute**)
- the router is either implementing a routing protocol or has all necessary static routes (**show ip route**)
- the router is receiving PIM joins from the downstream router and sending joins to its own upstream neighbor (**show ip pim-sparse neighbor, show ip pim-sparse traffic, debug ip pim-sparse joinprune**)

Quick Start

This section provides the commands you must enter to quickly configure PIM-SM for multicast routing. Only a minimal explanation is provided.

If you need additional information about any of these options, see “Contents” on page 13-1 to locate the section and page number that contains the explanation you need.

1. Enable PIM-SM on each interface that connects to a network that must support or transit multicasts. Move to the interface configuration mode context and enter this command:

Syntax: ip pim sparse-mode

For example, enable PIM-SM on a LAN and a WAN interface:

```
ProCurve(config)# interface eth 0/1
ProCurve(config-eth 0/1)# ip pim sparse-mode
ProCurve(config-eth 1)# interface frame-relay 1.1
ProCurve(config-fr 1.1)# ip pim sparse-mode
```

2. Move to the PIM sparse configuration mode context:

Syntax: router pim-sparse

3. Specify the IP address of the RP:

Syntax: rp-address <A.B.C.D>

For example:

```
ProCurve(config-pim-sparse)# rp-address 10.1.1.1
```

Note

You must configure the same RP address on every router that runs PIM in your network.

The steps above are sufficient for routing multicasts between remote sites in many WAN environments. You can use the following commands to set up more complex features:

4. You can alter when the router switches to an SP tree. Specify the number of packets that the router must receive from a source before it joins the SP tree with this command, entered from the PIM sparse configuration mode context:

Syntax: spt-threshold <packets>

5. You can also prohibit the router from using SP trees at all. Enter this command from the PIM sparse configuration mode context:

Syntax: spt-threshold infinity

6. You can configure different RPs to support different multicast groups. Configure the address or range of addresses for groups that the RP should support in a standard ACL. Then associate the ACL with the RP. Follow these steps:

- a. Create the ACL from the global configuration mode context:

Syntax: ip access-list standard <listname>

- b. If you want to *ensure* that a certain router acts as RP for a certain group, you must disable all other RPs from supporting that group. For example, you want Router A to support group 1. Router B supports a range of groups that includes group 1. In this case, you must enter a deny statement for the group 1 in the ACL for Router B. Enter this command:

Syntax: deny [host <A.B.C.D> | <A.B.C.D> <wildcard bits>]

Use the **host** keyword to specify a single address. Use wildcard bits, which operate on reverse logic of subnet masks, to select a range of addresses. (Remember to enter the deny statement before permitting the range of addresses that includes the host or hosts denied.)

- c. Specify the address or range of addresses for the multicast groups that the RP can support:

Syntax: permit [host <A.B.C.D> | <A.B.C.D> <wildcard bits>]

Note

Simply because an RP can support a group does not mean that it will actually do so. If a second RP supports the group as well, that second RP may be selected instead.

- d. Move to the PIM sparse configuration mode context:

Syntax: router pim-sparse

- e. Specify the IP address of the RP and associate the ACL with the RP:

Syntax: rp-address <A.B.C.D> access-group <listname>

Caution

Routing protocols send messages to multicast addresses 224.0.0.0 through 224.0.0.255. You should be very careful to configure at least one RP to support these address so that routers can transit routing updates between remote sites.

Configuring Multicast Support with PIM-SM
Quick Start