Configuring IP Multicast

SYSTEM ADMINISTRATOR GUIDE

Copyright

© Ericsson AB 2009–2011. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

SmartEdge	is a registered trademark of Telefonaktiebolaget LM Ericsson.
NetOp	is a trademark of Telefonaktiebolaget LM Ericsson.

Contents

1 Overview

This document provides an overview of IP multicast, and describes the tasks and commands used to configure, monitor, troubleshoot, and administer IP multicast features through the SmartEdge router.

This document applies to both the Ericsson SmartEdge® and SM family routers. However, the software that applies to the SM family of systems is a subset of the SmartEdge OS; some of the functionality described in this document may not apply to SM family routers.

For information specific to the SM family chassis, including line cards, refer to the SM family chassis documentation.

For specific information about the differences between the SmartEdge and SM family routers, refer to the Technical Product Description *SM Family of Systems* (part number 5/221 02-CRA 119 1170/1) in the `Product Overview` folder of this Customer Product Information library.

Note: In the following descriptions, the term controller card refers to any version of the Cross-Connect Route Processor (XCRP4) Controller card, including the controller carrier card, unless otherwise noted.

There are three basic types of IP communication: unicast, broadcast, and multicast. Unicast communication occurs between a source host and a single, unique destination host; it is one-to-one communication. Unicast packet headers specify a single IP address of a destination host. Broadcast communication occurs between a source host and all other hosts on the network; it is one-to-all communication. Broadcast packet headers specify an IP broadcast address that includes all destination hosts on the subnet. Multicast communication, by contrast, falls somewhere between unicast and broadcast communication.

Multicast communication enables a source host to send IP packets to any number of hosts, anywhere within an IP network; it is one-to-any communication. That is, multicast communication is not limited to sending packets to a single destination host, or sending packets to every host on the network. Instead, multicast enables a source host to send IP packets to as many destination hosts as necessary, but no more than that. The advantages of multicast communication, unlike broadcast communication, which floods the network with unnecessary traffic, is that a source host can communicate with more than one destination host without sending traffic to every host on the network. This results in an economic use of bandwidth.

The main challenge for multicast communication is developing a method for determining which hosts will receive multicast traffic, and which hosts will not receive the traffic. Several different multicast protocols have been developed, each with its own unique approach to addressing the multicast challenge. The SmartEdge router supports the following multicast protocols:

- Internet Group Management Protocol (IGMP)
- Protocol Independent Multicast (PIM)
- Source-Specific Multicast (SSM)
- Multicast Source Discovery Protocol (MSDP)
- Anycast RP
- Multicast VPNs
- Remote Multicast Replication (RMR)

These multicast protocols are described in the following sections.

1.1 Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is the method by which local hosts join a multicast group. A host that wants to join a multicast group should immediately transmit an unsolicited Membership Report for that group to the multicast-enabled router for that network. The router maintains a list of multicast group memberships for each attached network, and a timer for each membership. The designated router (DR), which is the multicast-enabled router with the highest IP address on the network, periodically sends a general query to learn which groups have members on an attached network, and a group-specific query to learn if a particular group has any members on an attached network.

Note:

> An IGMP group can have more than one source, indicated with the following notation:
>
> - (*, G)—The group (G) uses every possible source IP address
> - (S, G)—The group (G) uses a particular source IP address

The sections that follow describe additional IGMP-related features.

1.1.1 IGMP Bandwidth Limitation

The IGMP bandwidth limitation feature is targeted at applications where many potential receivers share the same port. When too many receivers join at the same time, the aggregate bandwidth exceeds that of the physical port, resulting in unacceptable service. The loss of packets is more visible for video and audio types of applications, in the form of interruptions, than for unicast Transmission Control Protocol (TCP) applications, where the sender backs off and retransmits. With this feature, you can decide when to reject new IGMP joins, and you can set priorities among receivers.

1.1.2 IGMP Membership Tracking

The IGMP membership tracking feature allows explicit tracking of group membership for all multicast hosts in a multiaccess network. Because it allows the instant-leave feature to work on a multiaccess network, membership tracking enables much lower leave latency and faster channel surfing.

Membership tracking, which is enabled by default, works with IGMP Version 2 (IGMPv2) and IGMP Version 3 (IGMPv3). The sections that follow describe how membership tracking works within each IGMP version.

1.1.2.1 Membership Tracking with IGMPv2

When a host running IGMPv2 joins a group, it sends a membership report to the router. The router adds the host’s IP address to the group membership list, which enables the router to track which hosts are members of a particular group on the same multiaccess network. When a host sends an IGMPv2 Leave message, it is removed from the group membership list.

1.1.2.2 Membership Tracking with IGMPv3

When a host running IGMPv3 joins a group from a source list, it sends a membership report for a group and source as the include source list. The router adds the host’s IP address to the list of interested members for all the sources in the source list. When a host removes a source from its source list, the router removes the host from the group’s source record, and if the host was the last interested host for that source, and the circuit is configured with instant-leave, the router performs an instant-leave operation for the source record.

1.1.3 IGMP Snooping

IGMP snooping is an Ethernet bridge feature which ensures that multicast traffic is forwarded only to those nodes that are interested in receiving that particular type of multicast traffic. Without IGMP snooping, multicast traffic is flooded to all the ports in a broadcast domain. When IGMP snooping is enabled, multicast traffic is not flooded to all circuits in the network. Instead, IGMP snooping forwards multicast packets received at the bridge for a particular group only to those hosts that have subscribed to that group.

The advantages of enabling IGMP snooping on a router are as follows:

- IGMP snooping reduces the amount of multicast traffic in a network, thereby reducing packet processing.

- IGMP snooping increases the bandwidth of the host because the host does not have to receive and filter all the multicast traffic generated by the network.

IGMP snooping is supported on Ethernet bridges and VPLS instances.

IGMP snooping is compliant with the following RFCs:

- RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*.
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

The sections that follow describe IGMP snooping-related features.

1.1.3.1 Configure IGMP Snooping on an Ethernet Bridge

Use the **`igmp snooping`** command in IGMP snooping in bridge configuration mode to enable IGMP snooping on a particular Ethernet bridge and enter IGMP snooping configuration mode for that bridge. After you are in IGMP snooping configuration mode, you can modify the default IGMP snooping configuration parameters for the bridge, if desired, or you can exit IGMP snooping configuration mode and run the default IGMP snooping configuration on the bridge. The default configuration for the following attributes can be modified from within IGMP snooping configuration mode:

- Whether IGMP proxy reporting is enabled or disabled on the bridge
- Static (*,G) and (S,G) membership
- Interval at which IGMP group-specific host query messages are sent out
- Number of IGMP packets that can be lost before the bridge goes down
- Whether the bridge uses IGMP snooping Version 2 or Version 3

To further customize the IGMP snooping configuration for a particular Ethernet bridge circuit, create an IGMP snooping profile, as described in IGMP Snooping Profile. After you create an IGMP snooping profile, you must reference it in a bridge profile. All bridge Ethernet circuits using that bridge profile assume the IGMP snooping configuration from the referenced IGMP snooping profile.

1.1.3.2 Configure IGMP Snooping on a VPLS Instance

When IGMP snooping is enabled on a bridge, all VPLS circuits that are bound to that bridge run IGMP snooping.

You can further customize the IGMP snooping configuration for a particular VPLS instance by creating an IGMP snooping profile, as described in IGMP Snooping Profile. After you have created an IGMP snooping profile, you must reference it in a bridge profile that is then referenced in a VPLS profile. All circuits that use that VPLS profile assume the IGMP snooping configuration settings from the referenced IGMP snooping profile.

1.1.3.3 Configure IGMP Snooping on a LAG

When IGMP snooping is enabled on a bridge, any LAG that is bound to that bridge runs IGMP snooping.

When a new Ethernet or dot1q circuit is added to a bundle, preexisting LAGs on that bundle are not rehashed. Only new LAGs include new circuits in their hashing calculation. When a circuit is deleted from a circuit bundle, all (*,G) and (S,G) groups that are using the circuit as an outgoing circuit choose another outgoing circuit through rehashing.

1.1.3.4 IGMP Snooping Profile

An IGMP snooping profile simplifies IGMP snooping configuration by providing values that can be applied to any number of Ethernet bridge or VPLS circuits. Instead of having to define IGMP snooping parameters individually for each circuit, you need to define the parameters only once in the IGMP snooping profile and then apply that profile to the appropriate circuits.

When an IGMP snooping profile is referenced by a bridge profile that is applied to a bridge or VPLS circuit, that circuit assumes the configuration settings from the IGMP snooping profile.

An IGMP snooping profile defines the following parameters:

- Static membership for a multicast group
- Ability of circuits to send and receive multicast data
- Maximum number of (*,G) and (S,G) groups a single circuit is allowed to join
- Access list to filter IGMP control messages and traffic
- Discovery or specification of multicast routers

After an IGMP profile is configured, it can then be referenced in a bridge profile. The bridge profile can be attached to a circuit that is bound to an Ethernet bridge that has IGMP snooping enabled, or it can be referenced by a VPLS profile (which is attached to a VPLS instance). The configuration in the referenced IGMP snooping profile is then automatically applied to the appropriate circuits.

Figure 1 illustrates how IGMP snooping profiles are attached to circuits.

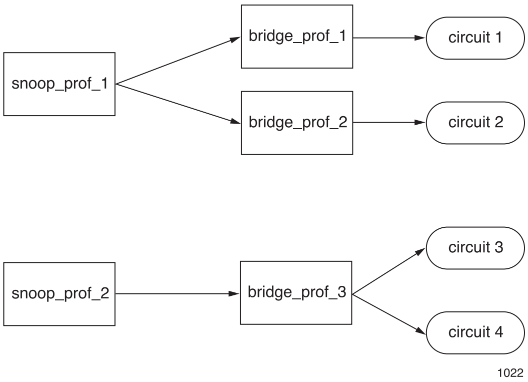

Figure 1 Example: How IGMP Snooping Profiles Are Attached to Circuits

In this figure, there are two IGMP snooping profiles, called `snoop_prof_1` and `snoop_prof_2`. Circuits 1 and 2 use the IGMP snooping configuration defined in the IGMP snooping profile called `snoop_prof_1`. Circuits 3 and 4 use the IGMP snooping configuration defined in the IGMP snooping profile called `snoop_prof_2`. To attach the IGMP snooping profiles to their respective circuits, each profile is referenced in a bridge profile that is attached to the circuit, as follows:

- The bridge profile called `bridge_prof_1`references the IGMP snooping profile called `snoop_prof_1`. Since`bridge_prof_1` is attached to circuit 1, circuit 1 takes on the IGMP snooping configuration defined in `snoop_prof_1`.

- The bridge profile called `bridge_prof_2` also references `snoop_prof_1. Since bridge_prof_2` is attached to circuit 2, circuit 2 takes on the IGMP snooping configuration defined in `snoop_prof_1`.

- To attach the IGMP snooping profile called `snoop_prof_2` to circuits 3 and 4, the user references `snoop_prof_2` in the bridge profile called `bridge_prof_3`. Since `bridge_prof_3` is attached to circuits 3 and 4, those circuits take on the IGMP snooping configuration defined in `snoop_prof_2`.

Note: Both IGMP snooping and bridge profiles are configured in global configuration mode and are applied directly to circuits. IGMP snooping and bridge profiles are not associated with any particular context, bridge, or interface.

1.1.3.5 Join and Leave a Multicast Group

When the SmartEdge router receives a request from a host to join an IP multicast group, the circuit connecting the host to the router is added to the outgoing circuit list of the IGMP state entry for the requested group. The IGMP state entry contains information about how to forward a multicast data packet. The entry is updated every time an IGMP join or leave message is received for each (*,G) or (S,G) group.

In addition, a bridge replies to queries for static groups when no other hosts are already joined to that group.

A static IGMP state entry can be added for a (*,G) or (S,G) group. Data is forwarded to statically joined circuits, even if an IGMP join message is not received by the circuits. A bridge can be configured to statically join (*,G) and (S,G) routes. In such cases, the bridge sends an IGMP join message to the appropriate upstream router if no other circuits are already joined to the specified group.

Use the **`static`** command in IGMP snooping profile configuration mode to configure the static membership setting for a multicast group in an IGMP snooping profile. All circuits attached to that IGMP snooping profile have a static membership with the specified multicast group.

When the IGMP router receives a leave message from a host, the associated circuit is removed from the outgoing circuit list.

1.1.3.6 IGMP Proxy Reporting

The IGMP proxy manages IGMP host tracking, messages, and queries. By default, the IGMP proxy is disabled, and the router is passive and does not suppress IGMP messages. When the IGMP proxy is disabled, the router listens for only IGMP messages.

When enabled on the SmartEdge router, the IGMP proxy runs in passive mode, in which host tracking is performed.

When passive mode proxy reporting is enabled, IGMP join messages for existing IGMP states are suppressed. Only the first IGMP join message is propagated toward the source. In addition, IGMP leave messages are forwarded from the hosts to the IGMP routers only when the last member leaves a group.

When the IGMP router queries the network, the IGMP snooping bridge replies on behalf of the hosts that are connected to it.

Use the **`proxy-reporting`** command in IGMP snooping bridge configuration mode to enable proxy reporting on an Ethernet bridge. To return the Ethernet bridge to the default configuration, in which proxy-reporting is disabled, use the **`no proxy-reporting`** command.

1.1.3.7 IGMP Query Solicitation

IGMP query solicitation facilitates the relearning of routers when topology changes occur. By default, the router periodically sends out general query solicitation messages to learn about any topology changes in the network. In an IGMP snooping configuring, you can configure a router to generate an immediate response to IGMP general query messages that are received from a bridge. IGMP query solicitation is enabled on all bridge circuits by default unless IGMP router monitoring is disabled in the IGMP snooping profile that is used by the bridge circuit. (Use the `no mrouter` command in IGMP snooping profile configuration mode to disable IGMP router monitoring in an IGMP snooping profile). Bridges that have IGMP snooping enabled forward query solicitation messages to all IGMP routers in a system. Use the `no query-solicitation` command in IGMP snooping configuration mode to disable the generation of IGMP query solicitation messages by a bridge.

By default, router interfaces send IGMP query messages only at a specified query interval (Use the `query-interval` command to set the IGMP query interval). Use the `igmp query solicitation` command in interface configuration mode to enable the immediate generation of an IGMP general query messages in response to a query solicitation message received by a router interface. When IGMP general query is enabled on a router interface, an IGMP query response message is sent to the appropriate bridge as soon as the query solicitation message is received, regardless of the configured query interval.

If query solicitation is enabled on a bridge circuit, query solicitation messages are sent when an STP-port-unblock event occurs or the state of a bridge circuit changes from down to up.

When an IGMP router receives a general leave message from a circuit, it responds with an IGMP general query that records topology changes.

1.1.4 Multicast Traffic Distribution

Ethernet and dot1Q link groups support multicast traffic distribution when the SmartEdge router is connected to an IGMP device. Traffic distribution enables the SmartEdge router to use more of the available bandwidth when forwarding traffic over a link group, which ensures faster service without packet loss.

Without traffic distribution, multicast traffic is forwarded on a single link and the entire bandwidth of the link group is not utilized. For example, if a link group contains four links, 75% of the bandwidth of that link group is not used when multicast traffic traverses one link only. With the multicast traffic distribution feature, multicast traffic traversing an Ethernet or a dot1Q link group is distributed across all the active links in a link group. Each multicast flow uses a hash function to select an active link. The hash function uses the following information to distribute the traffic evenly across the link group.

- Source IP address in the packet header
- Group IP address in the packet header
- Number of active links in the link group

Any change in the status of a link group triggers a rehashing function, which redistributes the traffic accordingly. For example, if an active link that is carrying multicast traffic goes down, the rehashing function selects another active link to carry the multicast traffic so that packets are not dropped. If no active links are available, the outgoing circuit is removed.

Note: Only Ethernet and dot1Q link groups support multicast traffic distribution.

Multicast traffic distribution is automatically enabled on an Ethernet or dot1Q link group when the group is bound to an IP or bridge interface with the **`bind interface`** command, and when either of the following features is enabled on the bridge:

- PIM
- IGMP snooping

The following statements apply to multicast traffic distribution:

- Multiple multicast streams may be forwarded on one link, even when other links are available.
- When an active link state changes in the control plane, the rehashing delay for all the multicast streams over a link group depends on the availability of CPU and communication with the forwarding plane.
- An individual link group can contain links from a single line card or from multiple line cards.
- After a seamless XCRP switchover or a restart of PIM or IGMP, multicast traffic continues to be forwarded on the same link it traversed before the switchover if the same number of links are active after the switchover or restart. If a different number of active links is available after an XCRP switchover or a restart of PIM or IGMP, multicast traffic is redistributed.

1.2 Protocol Independent Multicast

Protocol Independent Multicast (PIM) is a multicast routing protocol that runs over an existing unicast infrastructure. As its name implies, PIM is IP routing protocol independent; that is, regardless of the unicast routing protocol used to populate the unicast routing tables, PIM uses those tables to perform multicast forwarding tasks. PIM also relies on IGMP to provide and maintain all multicast group membership information.

PIM is implemented in the following three ways:

- Protocol Independent Multicast Dense Mode
- Protocol Independent Multicast Sparse Mode
- PIM-Dual Join Mode

Note: If you configure PIM (of any mode) for an interface, it is not necessary to explicitly configure IGMP.

1.2.1 Protocol Independent Multicast Dense Mode

Protocol Independent Multicast Dense Mode (PIM-DM) uses source distribution, or shortest path trees (SPTs), to distribute multicast traffic to receivers in the network. PIM-DM uses Hello messages to establish neighbor adjacencies and builds an initial SPT based on the neighbor adjacencies. The initial form of the SPT is also referred to as a broadcast tree because PIM routers use it to distribute multicast traffic in a broadcast-like manner; that is, multicast traffic is sent to all PIM-DM routers, regardless of whether they want to receive the traffic.

After the initial flood of multicast traffic on a PIM-DM network broadcast tree, the tree is trimmed to a minimum spanning tree. PIM-DM routers send prune messages to remove themselves from the SPT if they meet any of the following conditions:

- The PIM router is a leaf router and has no directly connected receivers.
- The PIM router is a nonleaf router on a point-to-point link and receives a prune message from its neighbor.
- The PIM router is a nonleaf router on a LAN segment with no directly connected receivers, has received a prune message from a neighbor on a LAN segment, and has no other neighbor on the LAN segment that can override the prune message.

Prune messages can be overridden by Join messages sent by downstream neighbors that want to continue, or begin receiving multicast traffic on the specified SPT. Pruned branches are restored periodically to see if new multicast group members have joined since the branch was pruned.

Note: The PIM-DM State Refresh feature keeps pruned branches from being automatically restored to the PIM-DM network by periodically forwarding control messages down the broadcast tree. The control messages refresh the prune state on the outgoing interfaces of each router in the broadcast tree. Enabling this feature is useful in situations in which restoring previously pruned branches would consume too much bandwidth by reflooding unwanted traffic over the PIM-DM network.

The PIM-DM flooding and pruning mechanism is optimal for only densely populated groups.

1.2.2 Protocol Independent Multicast Sparse Mode

Protocol Independent Multicast Sparse Mode (PIM-SM) differs from PIM-DM in the following ways:

- Routers with directly attached multicast receivers, or downstream receivers, are required to join a sparse mode distribution tree by transmitting explicit join messages. If a router does not become part of the distribution tree, it does not receive multicast traffic.

- PIM-SM uses a rendezvous point (RP) to serve as a distribution point for multicast traffic from one or more related multicast sources.

PIM-SM sends multicast traffic only to locations on the network that explicitly request membership in a multicast group. The requests are called PIM Join messages, which are sent hop by hop toward the multicast source, creating an SPT. As the PIM Join message is sent up the tree, routers along the path establish the multicast forwarding state so that multicast traffic can be sent back down the path. Likewise, PIM prune messages can be sent hop by hop toward the multicast source to remove locations from the multicast group.

On a PIM-SM network, SPTs are trees created by a collection of joins in which the root of the tree is also the multicast source; however, the root of an SPT does not need to be the multicast source, but can be a location called the rendezvous point. SPTs with an RP as its source are called shared trees. With a shared tree, multiple multicast sources share the same tree structure by forwarding their multicast traffic to the RP, where it is then distributed down the shared tree.

Any router on a network can be specified as the RP, or multiple routers can be specified as candidate RPs (C-RPs). For C-RPs, an RP election process determines which router serves as the RP. The bootstrap router (BSR) eliminates the need to manually configure each router on the network with the RP information by distributing group-to-RP mapping information to all routers on the network. During the RP election process, all C-RPs send their candidacy advertisements to the BSR, and the BSR distributes the group-to-RP mappings.

For redundancy, multiple candidate BSRs (C-BSRs) can be specified. A BSR election process, based on the router priority level, determines which C-BSR serves as the BSR.

PIM-SM is not supported on multibind interfaces. If you are configuring a multibind interface, you must use the **`pim sparse-mode passive`** command in interface configuration mode to prevent PIM messages from being exchanged on the egress interface, while allowing the interface and its circuits to be populated in a multicast forwarding entry by receiving an IGMP report or a data packet.

Consider the following restrictions when configuring multicast for subscribers:

- If the multicast source uses a router-mode CPE, the SmartEdge router runs in active PIM-SM only.
- If the multicast receiver uses a router-mode CPE, the SmartEdge router runs in active PIM-SM or passive PIM-SM.
- If a bridge-mode CPE exists between the multicast source or receiver and the SmartEdge router, the SmartEdge router can run in active PIM-SM or passive PIM-SM.

Note: The PIM-SM explicit join mechanism is optimal only for sparsely populated groups.

1.2.3 PIM-Dual Join Mode

Protocol Independent Multicast-Dual Join Mode (PIM-Dual Join) is an extension of PIM-SM that provides an extra level of protection for multicast join requests transmitted in this mode. The PIM-Dual Join feature enables the SmartEdge router to join the same multicast stream on two individual interfaces, which provides the ability to support fast failover of subscriber multicast streams when network failures occur. Service providers use PIM-Dual Join Mode to protect IP Multicast traffic from link failure when a SmartEdge router at the provider site sends a Multicast join request over a link.

Protocol Independent Multicast-Dual Join Mode (PIM-Dual Join) is an extension of PIM-SM that provides an extra level of protection for multicast join requests transmitted in this mode. The PIM-Dual Join feature enables the SmartEdge router to join the same multicast stream on two individual interfaces, which provides the ability to support fast failover of multicast streams when network failures occur. Service providers use PIM-Dual Join Mode to protect IP Multicast traffic from link failure when a SmartEdge router at the provider site sends a Multicast join request over a link.

Note: The PIM-Dual Join method applies mostly to links or nodes through which traffic passes at 50 milliseconds per second through the SmartEdge router.

PIM-Dual Join Mode sends packets over both primary and secondary join request sessions of specified multicast group addresses on two different reverse path forwarding (RPF) interfaces on upstream routers. The RPF interface enables a router to examine all packets received on that interface. The router checks to make sure that the source address of a packet arriving on that interface appears in the routing table. It also identifies the interface that first received the packet. This process helps identify problems associated with spoofing or malformed IP source addresses.

Note: The PIM-Dual Join method can occur only if secondary RPF interfaces exist and are enabled.

PIM downloads multicast route entries with several attributes and sends both primary and secondary RPFs. The primary RPF has a setting of RPF-P, which indicates it is the active RPF entity. The secondary RPF has a setting of RPF-S, which indicates it is the backup RPF entity that sends Multicast Join requests when the active RPF entity goes down. It also sends addresses of both primary and secondary outgoing interfaces (OIFs), which are the interfaces connected to the device of the subscriber.

PIM downloads multicast route entries with several attributes and sends both primary and secondary RPFs. The primary RPF has a setting of RPF-P, which indicates it is the active RPF entity. The secondary RPF has a setting of RPF-S, which indicates it is the backup RPF entity that sends Multicast Join requests when the active RPF entity goes down. It also sends addresses of both primary and secondary outgoing interfaces (OIFs), which are the interfaces connected to the far-end device.

PIM-Dual Join Mode uses the RPF route from the SmartEdge router as its primary interface. You select the secondary interface by specifying the RPF interface in the **`pim-dual-join`** command. When the primary RPF is active, the SmartEdge router checks RPFs on incoming IP multicast traffic marked for forwarding over the primary RPF. The SmartEdge router drops secondary IP multicast traffic that arrives on secondary RPFs because it recognizes this traffic as invalid since the primary path is working. When the primary path is not working, the SmartEdge router recognizes the secondary IP multicast traffic as valid.

IGP routing advertises the most optimal route to the routing table of the SmartEdge router. The router then advertises that route to PIM. When PIM-Dual Join Mode occurs, no secondary RPF route initially exists, although the SmartEdge router sends a secondary join request on the RPF interfaces specified in the **`pim dual-join`** command.

Before you configure PIM-Dual Join Mode, you must understand the following:

- Supported Topologies
- PIM-Dual Join Rules and Restrictions

1.2.3.1 Supported Topologies

The PIM-Dual Join functionality is based on PIM SSM and can map IGMPv2 and IGMPv3 ASM joins to a PIM SSM group join for the configured multicast server. This functionality means that no RP configuration is needed on the SmartEdge router because the configuration could potentially lead to incorrect behavior of the PIM-Dual Join functionality.

The PIM-Dual Join feature supports both square and star network topologies.

Figure 2 shows an example of a square network topology.

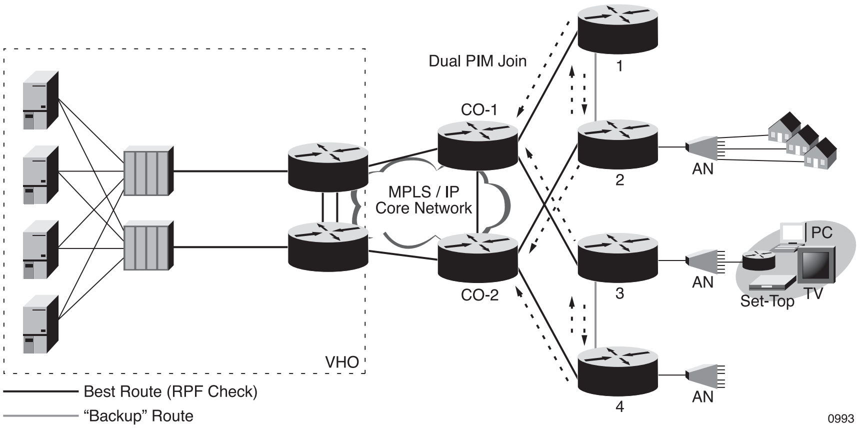

Figure 2 Example of a Square Network Topology

Figure 3 shows an example of a star network topology.

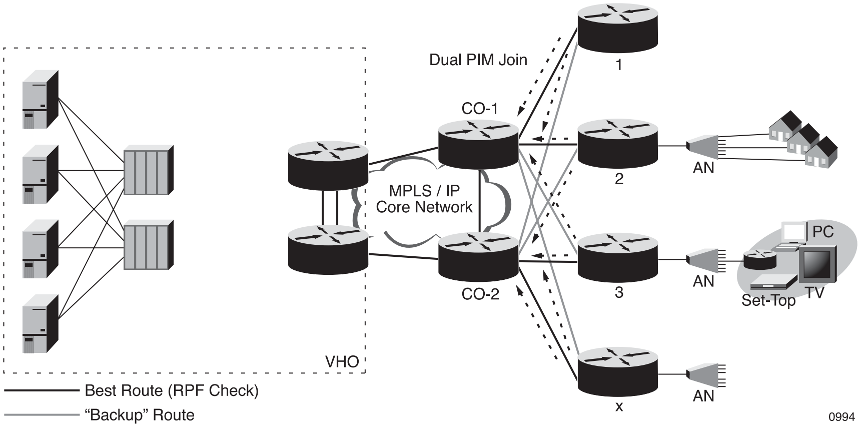

Figure 3 Example of a Star Network Topology

Note: The SmartEdge router uses a maximum of two interfaces: a primary interface and a secondary RPF interface that pulls down multicast traffic.

1.2.3.2 PIM-Dual Join Rules and Restrictions

Keep the following rules and restrictions in mind when configuring PIM-Dual Join on the SmartEdge router:

- The primary RPF interface is always selected by the RIBd process, and the configured RPF interfaces automatically pick the secondary RPF interface. If the first configured RPF interface is the same as the primary RPF interface selected by the RIBd process, then the second configured RPF interface is used as the secondary RPF interface. If the first configured RPF interface is different from the primary RPF interface selected by the RIBd process, then the first configured RPF interface is used as the secondary RPF interface. If the first configured RPF interface is not in a working up state, then the second configured RPF interface is used instead.

- When PIM-Dual Join is configured on the SmartEdge router, only one of the received multicast streams is forwarded to the multicast receivers.

- A static OIF can be configured for a PIM-Dual Join entry.

- Only one PIM-Dual Join entry is permitted for each source or group. Users cannot configure multiple PIM-Dual Join entries for the same group or source.

- BFD is supported as the link failure detection to ensure PIM-Dual Join fails over from the primary RPF interface to the secondary RPF interface.

- Mandatory source specification for PIM-Dual Join entries prevents users from incorrectly configuring Dual-PIM Join entries without specifying a multicast server IP address.

- By default, the PIM-Dual Join behavior is to wait for an IGMP Join packet before it sends PIM-Dual Join messages upstream in the network. Use the **send-join** command to pull down the multicast stream before an IGMP Join is received. The **send-join** command sends join messages upstream on the RPF primary and secondary interface without any OIFs being present. If a join message from IGMP is received by PIM after the send-join feature is enabled for a specific group, the SmartEdge router does not send any new join messages upstream in the network. Instead, the SmartEdge router immediately adds the client to the OIF list and starts forwarding the multicast stream.

1.3 Source-Specific Multicast

The source-specific multicast (SSM) feature is an extension of multicast routing. In SSM, traffic is forwarded to receivers from only those multicast sources to which the receivers have explicitly joined. For multicast groups configured to use SSM, only source-specific multicast distribution trees (MDTs) are created, and not shared trees.

The PIM-SSM routing protocol supports the implementation of SSM and is derived from PIM-SM.

SSM is supported by IGMPv3, but if you enable IGMPv3 on the SmartEdge router interface, you can translate IGMPv1, v2 subscriber reports to IGMPv3 by configuring SSM mapping using the `ssm-map` command.

The address range 232.0.0.0 through 232.255.255.255 is reserved for multicast applications and protocols. Existing IP multicast receivers cannot receive traffic when trying to use addresses in a defined SSM range, unless they are SSM enabled.

For more information on SSM routing, see the Internet Draft, *Source-Specific Multicast for IP, draft-ietf-ssm-arch-00.txt*.

1.4 Multicast Source Discovery Protocol

The Multicast Source Discovery Protocol (MSDP) is the method used to link interdomain RPs so that multicast messages can be forwarded to other domains that have active group membership. RPs in a PIM domain know about all active sources in its own domain, but not other domains; however, if an RP from one domain is peered with another RP in a different domain, it can send source active messages from one domain to the other.

Using MSDP provides the following benefits:

- A multicast distribution tree can be divided into different segments.
- Local members for each segment can join their local segments.

- Each segment depends only on its own RP. Each RP has information about multicast sources of each domain, so members in each segment can stay in their local segment.
- Each domain does not have to globally send its member information.

1.5 Anycast RP

In a basic PIM-SM network, a single RP is used by all multicast sources and receivers. Anycast RP is a mechanism that provides RP redundancy and load-sharing capabilities by allowing the use of multiple RPs within a single multicast domain. Assuming that the sources are evenly spaced around the network, an equal number of sources register with each RP. That is, the process of registering the sources is shared equally by all the RPs in the network.

All routers acting as RPs must be configured with a loopback interface using the same anycast RP address. All downstream routers use that anycast RP address as IP address for their local RP. To facilitate communication between RPs, each router acting as an RP must also be configured with its own unique IP address, which is used only to send and receive messages from the other RPs.

When a source registers with one RP, a message is sent to the other RPs informing them that there is an active source for a particular multicast group. The result is that each RP knows about the active sources in the area of the other RPs. If any of the RPs were to fail, IP routing would converge and one of the RPs would become the active RP in more than one area. New sources would register with the backup RP. Receivers would join toward the new RP and connectivity would be maintained.

Our implementation of anycast RP eliminates the dependency on MSDP by removing MSDP peering between the anycast RPs; however, to advertise internal sources to routers outside of the routing domain, MSDP may still be required.

1.6 Multicast VPNs

Standard Border Gateway Protocol/Multiprotocol Label Switching Virtual Private Networks (BGP/MPLS VPNs) do not provide a way for IP multicast traffic to travel from one VPN site to another. Implementing multicast domain trees (MDTs) provides a scalable solution to support IP multicast over BGP/MPLS VPNs. Currently, MDTs support only IPv4 multicast.

When a network uses many VPNs, where each VPN can have many multicast groups, and each multicast group can have many multicast transmitters, it is not scalable to have one or more distribution trees for each multicast group. A scalable IP multicast solution for MPLS/BGP VPNs requires that the amount of VPN-specific information maintained by the P routers must be proportional only to the number of VPNs that run over the backbone. The amount of VPN-specific information in the P routers is not sensitive to the number of multicast groups or to the number of multicast transmitters within the VPNs. However, there is a trade off to using this scalable solution; nodes that are not on a path to a multicast receiver may still receive multicast packets, and will have to discard them. That is, greater scalability reduces multicast route optimization.

A multicast-enabled VPN has a corresponding multicast domain. A provider edge (PE) router that attaches to a multicast-enabled VPN belongs to the corresponding multicast domain. For each multicast domain, there is a default MDT through the backbone, connecting all of the PE routers that belong to that multicast domain. A PE router may be in as many multicast domains as there are VPNs attached to it. However, each multicast domain has its own MDT. The MDTs are created by running PIM in the backbone, and in general an MDT also includes P routers on the paths between the PE routers. For MDTs to work properly, the following conditions must be met:

- PIM must be the multicast routing protocol used in the VPN.
- PIM must be the multicast routing protocol used in the backbone network.
- The backbone network must support IP multicast forwarding.

Default MDTs are constructed automatically as the PE routers in the domain come up. Construction of a default MDT does not depend on multicast traffic in the domain. That is, it exists before any multicast traffic is detected.

In a multicast-enabled VPN, each customer edge (CE) router has a PIM adjacency to a PE router, but CE routers at different sites do not have PIM adjacencies to each other. Multicast packets from within a VPN are received from a CE router by an ingress PE router. The ingress PE router encapsulates the multicast packets and forwards them across the default MDT to all PE routers connected to the specified VPN. If a PE router receiving the multicast packets is not on the path to any multicast receiver of that multicast group, it discards the multicast packet.

For the SmartEdge router implementation of multicast VPNs, the default MDT group must be configured on an intercontext interface in a VPN context. This interface is similar to a loopback interface in that it is not bound to anything and does not need an IP address. It creates an intercontext circuit between the VPN context and the local context. PIM-SM must also be configured on this intercontext interface. The MDT encapsulation type must be configured on a loopback interface in the local context. The loopback interface is used to source multicast packets on the MDT.

1.7 BFD for PIM

When enabled on a PIM interface, Bidirectional Forwarding Detection (BFD) ensures early detection of network failures, such as router control plane failures or loss of router adjacency over nondirect links. The advantages of BFD-enabled PIM interfaces include:

- PIM responds faster to interface state changes and control plane failures. When BFD is disabled, PIM receives interface state change notifications through ISM, which takes time and can subject the interface to system loads. When BFD is enabled on an interface, BFD immediately notifies PIM of interface state changes so that PIM can respond accordingly. BFD also detects control plane failures in cases when the physical port remains in the Up state.
- Without BFD, PIM takes longer to detect a loss of router adjacency over nondirect links; three hello intervals (30 seconds by default) pass before PIM detects the failure. For example, when a Layer 2 device exists between a PIM-enabled router and a neighbor, PIM does not immediately detect when the link between the Layer 2 device and the PIM router fails; in this case, the PIM router has to wait for the neighbor state machine to indicate that the neighbor is no longer present.

  When BFD is enabled on the PIM interface, PIM immediately detects when a neighbor is not present.

BFD is enabled by default on a PIM interface, but can be disabled by using the *no pim bfd* command in interface configuration mode.

1.8 Remote Multicast Replication

Remote multicast replication (RMR) is used to optimize the delivery of multicast services. RMR requires a multicast controller (MC), such as the SmartEdge router, to maintain complete subscriber awareness and subscriber control for all traffic types, and a multicast replicator (MR), such as a digital subscriber line access multiplexer (DSLAM), to perform IGMP snooping and multicast traffic replication on the subscriber-side link. Figure 4 shows the network topology used for RMR.

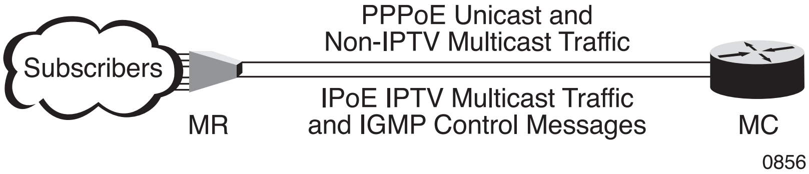
PPPoE Unicast and
Non-IPTV Multicast Traffic
Subscribers
IPoE IPTV Multicast Traffic
and IGMP Control Messages
MR
MC
0856

Figure 4 Remote Multicast Replication Network Topology

The IP over Ethernet (IPoE) circuit is configured to carry multicast traffic and IGMP control messages between the MC and the MR. The MC starts forwarding a multicast stream when it receives the first IGMP join message, and stops forwarding the stream when it receives the last IGMP leave message. The MR replicates the incoming multicast stream to all subscribers that need a copy of that stream, thus reducing bandwidth usage on the IPoE circuit. The MR makes its multicast replication and forwarding decisions by snooping the IGMP join and leave messages from the subscriber.

The MR performs multicast replication, but it does not support any routing functions, user authentication, or billing functions. These functions are supported by the MC (SmartEdge router) through a Point-to-Point Protocol over Ethernet (PPPoE) circuit from the subscriber to the MC.

A single MC can support multiple MRs. When multiple MRs are used, the MC performs per-MR multicast replication, while each MR performs per-subscriber multicast replication.

To configure RMR on the SmartEdge router, you must enable an interface to forward the multicast data and IGMP control messages, and you must enable an IGMP service profile forward multicast data for IGMP messages received on the PPPoE circuit on the IPoE interface.

1.9 Packet Replication for High-Bandwidth Multicast Traffic

For high-bandwidth multicast applications such as IPTV, you can enable PMA traffic replication to destination line cards.

PMA traffic replication is only supported on PPA3-based line cards: 20-port Gigabit Ethernet card (ge4-20-port) and 4-port 10 Gigabit Ethernet card (10ge-4-port).

For more information on the requirements and restrictions, see the *optimize replication* command.

1.10 PIM-SM and PIM-SSM MDT Discovery

The SmartEdge OS supports PIM-SM default MDT and default MDT auto-discovery for PIM-SSM through BGP subaddress family identifier MDT (SAFI-MDT). Such multicast-enabled VPNs are sometimes called draft-rosen multicast VPNs, in reference to the `draft-rosen-vpn-mcast-08` specification (now RFC 4364, `BGP/MPLS IP Virtual Private Networks (VPNs)`). In a draft-rosen network, each multicast-enabled VPN corresponds to a multicast domain (MD). Each MD has a default MDT through the backbone of the network connecting all of the provider edge (PE) routers belonging to the MD. The MDT is constructed when the PE routers are brought up, through auto-discovery of the member PE routers. The mechanism for constructing the MDT varies with the version of PIM being used:

- In a PIM-SM environment, a rendezvous point (RP) provides rendezvous and auto-discovery services to the PE routers belonging to the multicast domain, establishing the PIM adjacencies between the routers. The source and receiver PE routers auto-discover one another through the RP. In this scenario, the local PE acts as the source, sending Register messages to the RP. The RP then builds a Shortest Path First (SPF) tree toward the source PE. The remote PE acts as the receiver for the MDT group, sending (*, G) Join messages toward the RP and joining the distribution tree for the group.

- In a PIM-SSM environment, no RP is required. Instead, the PE routers use SAFI-MDT to auto-discover one another directly. Using PIM-SSM auto-discovery allows the PE to directly join to a source tree rooted at another PE for MDT without an RP. Eliminating the RP reduces management overhead and eliminates a potential point of failure. Also, forwarding delay is reduced.

To support PIM-SSM MDT auto-discovery, PE source addresses and the MDT SSM group addresses must be distributed to the PIM protocol over BGP. To do this, BGP uses a subaddress family called MDT-SAFI, which has a value of 66. This information is maintained within the BGP protocol, wherein it is considered an MDT route.

MDT routes are treated by BGP in a similar way as unicast VPN routes: best-path selection is performed on the PE routers and route reflectors (RRs). However, while VPN routes are downloaded to the Routing Information Base (RIB), MDT routes are downloaded to PIM.

MDT routes can be advertised with route targets (RTs) but RTs are not required, since the PE router can identify the multicast VPN (mVPN) to which the routes belong by the MDT group address. For the SmartEdge OS, the decision to accept an MDT route is performed in two steps.

1. In the first phase of filtering, if BGP receives the routes with RTs attached, the routes are accepted only if at least one VPN context is configured to import at least one of the RTs attached to the route. If the route does not have any RTs attached, BGP passes the routes to PIM for the second phase of filtering.

2. In the second phase of filtering, PIM uses the MDT group address configured in any VPN context.

1.10.1 Advertising the MDT Routes and Setting Up the Default MDT

When the MDT default group address is configured in a VPN context, PIM passes this information to BGP coupled with the local PE router's source address for the multicast streams. This address is the address of a loopback interface in the local context; it is known throughout the AS. The next hop for BGP to use is also provided by PIM. This hop is just the IP address of the loopback interface.

When BGP receives this information, it advertises the MDT route to all of its iBGP PE neighbors and its RR neighbors. When a PE router receives these MDT routes from a BGP peer that has been enabled with the MDT address family, BGP performs the first level of filtering, using the RTs applied to the routes. BGP calculates the best path for the MDT routes and then converts this information into the information required by PIM and downloads the converted route to PIM. PIM caches this information and filters it based on the MDT group address for the context and then sends an (S,G) Join message toward the source address.

1.10.2 Withdrawing MDT Routes and Pruning the Default MDT

When the default MDT group address configuration is removed, the mVPN instance can no longer participate in the mVPN network. In this case, PIM removes the information from BGP, and BGP withdraws the corresponding MDT routes from all PE routers.

When the receiving PE routers recognize that the MDT routes have been withdrawn, they inform PIM. PIM removes the information from its cache and sends (S,G) pruning messages for this MDT group toward the PE router that sent the MDT route withdrawals.

1.10.3 BGP Restart

When a source PE router restarts BGP, PIM detects the BGP restart. When the restart completes, PIM sends all locally configured MDT group information to BGP as a triple consisting of (source address, MDT group, next hop) when the restart has completed. BGP updates its local cache with these entries and readvertises the MDT routes to remote PE routers. In addition, PIM performs stale processing for MDT route entries learned from BGP. If BGP does not refresh a particular entry after restart, the entry is deleted from the local cache and an (S,G) pruning message is sent to the source.

If BGP graceful restart is enabled, the peers of the restarting router do not remove the MDT entries from their databases, and PIM information is not withdrawn. If the restarting BGP router does not readvertise the MDT routes, the destination PE PIM modules remove the information and tear down the MDTs toward the restarting PE router.

1.10.4 PIM Restart

In a PIM restart, the BGP module at the source PE router marks all local MDT routes learned from PIM as stale and performs stale processing, including MDT route withdrawals, if the stale time expires. When PIM is restored, BGP relearns the local routes from PIM and readvertises them, if necessary, to all PE routers. At the destination PE, BGP detects the PIM restart and re-downloads all MDT routes learned from the remote PE routers.

1.10.5 Interaction of Configuration

If BGP on a local PE router is up and correctly configured for mVPN but the PIM process is not up or is not yet configured for multicast, BGP caches MDT route entries learned from remote PE routers. When PIM is enabled and available, BGP sends the MDT routes to PIM.

If a configured MDT group in a particular VPN context on a source PE router is changed, the PIM process on the source PE informs BGP to remove the entry for the previously configured group. PIM then sends the new group with the source address, MDT group, and next-hop information to BGP. BGP withdraws the old MDT route and advertises the new MDT route.

For the default MDT mechanism to work correctly, all PEs participating in a particular VPN context must configure the same MDT group. In the case where an MDT group on a particular PE router in a specific VPN context is misconfigured, PIM still caches MDT-SAFI entries sent by BGP as long as the group is in the address range configured for SSM. However, no PIM (S,G) Join message is sent to the advertising source PE router. When the local MDT group is modified to be correct, PIM sends the (S,G) Join toward the source PE router.

If PIM-SSM default MDT auto-discovery is the only default MDT mechanism used (that is, if all PE routers have MDT groups configured only in the SSM address range), PIM RP configuration is not required in the local context.

1.11 QoS Adjustments for RMR

When functioning as a Broadband Remote-Access Server (BRAS), the SmartEdge router can perform traffic management on a Virtual Local Area Network (VLAN) or a subscriber session that represents the flow of traffic destined for an access line or other specific portion of the access network. The outgoing rate enforced by the SmartEdge router on the traffic of a VLAN or a subscriber session may be based on the downstream bandwidth capacity of that network segment. When this is the case, the rate enforced by the SmartEdge router may need to be adjusted to account for additional traffic inserted into the downstream traffic flow through Remote Multicast Replication (RMR). For more information on RMR, refer to section Remote Multicast Replication or section 6.3.2.3 of the TR-101 standard.

You cannot use QoS adjustments for RMR when the SmartEdge router functions as the Multicast Controller as described in Section 1.8. QoS adjustments for RMR are designed to be used in conjunction with an external Multicast Controller configured for RMR. Configure the network and multicast application so that IGMP Joins and Leaves for applicable multicast streams are forwarded to the SmartEdge router on the session or virtual LAN (VLAN) to which each corresponding QoS adjustment should be applied.

When a subscriber joins a multicast channel, the bandwidth of the subscriber needs to be adjusted downwards to prevent downstream congestion due to multicast replication.

For example, suppose the SmartEdge router is enforcing a downstream rate of 10 Mbps for a DSL subscriber, which matches the physical downstream sync speed of the subscriber access line. If the SmartEdge router allows up to 10 Mbps of traffic for the subscriber, but an additional 2 Mbps of multicast traffic is inserted towards the subscriber downstream by the RMR multicast replicator for an IPTV channel that the subscriber is viewing, the forwarding capacity of the access line is exceeded, causing indiscriminate loss of traffic. Using the QoS adjustment feature, when the SmartEdge router receives a subscriber's IGMP join, it reduces the rate of traffic toward the subscriber to 8 Mbps. The

SmartEdge router reserves 2 Mbps to account for additional traffic added by the multicast replicator. This ensures that the capacity of the access line is never exceeded.

Because the SmartEdge router does not receive direct information about how much downstream bandwidth is consumed by multicast traffic, it infers this information from other sources:

- When a subscriber joins or leaves a multicast channel, IGMP join or leave messages for that group address are generated. Those IGMP messages are forwarded on the subscriber circuit to the SmartEdge router, which uses this information to determine the additional bandwidth required for multicast traffic inserted onto the access line. You use the **`igmp group-bandwidth`** command to determine the criteria for matching joins to multicast groups as well as the corresponding amount of bandwidth to adjust.

- When a join message is received, the SmartEdge router subtracts the bandwidth value from the final rate value enforced for the applicable circuit binding (metering or queuing, or both if applicable).

- When a corresponding IGMP leave message is subsequently received, the value previously subtracted is restored to the enforced rate.

- If L2-aware shaping is enabled (using a **`qos profile overhead`** ***`name`*** CLI command for a circuit with a PWFQ binding), the SmartEdge router calculates the overhead of the packet header encapsulation on the access line and determines the final shaping rate. The average-packet-size value specified in the **`igmp-group-bandwidth`** command is combined with the Layer 1 and Layer 2 overhead specified in the overhead profile to calculate the appropriate rate adjustment. If the QoS overhead profile is configured with a rate factor, the final PWFQ binding rate at the circuit level is adjusted downward accordingly.

The RMR adjustments are applied to appropriate circuit-level rates for a given binding (metering or queuing). These rates can be the circuit-level rate value configured in the metering or PWFQ policy itself, or a customized rate provisioned for the circuit using RADIUS VSA 196 (Dynamic-QoS-Parameter), VSA 157 (Qos-Rate-Outbound), ANCP, or TR-101.

QoS adjustments for RMR supports the following circuit types:

- 802.1 PVCs with static bindings (for example, a nonsubscriber IPoE circuit)

- PPPoE and CLIPS subscriber sessions over ethernet ports, 802.1Q PVCs or 802.1Q tunnels

- PPPoE and CLIPS subscriber sessions over hitless and economical Access-LAG

QoS adjustments for RMR is supported for IGMP versions 1, 2, and 3.

2 Configuration and Operations Tasks

Note: In this section, the command syntax in the task tables displays only the root command.

To configure IP multicast, perform the tasks described in the sections that follow.

2.1 Configuring IGMP

To configure IGMP, perform the tasks described in Table 1. Enter all commands in interface configuration mode, unless otherwise noted.

Table 1 Configuring IGMP

Task	Root Command	Notes
Configure a router to join a multicast group. Enter this command on a subscriber interface leading to a multicast source within one hop.	*igmp join-group*	If you already configured PIM (any mode) for the interface, it is not necessary to explicitly configure IGMP.
Configure IGMP membership on the interface.	*igmp access-group*	Only multicast groups permitted by the access control list (ACL) are accepted on the interface.
Configure the interval at which the router sends IGMP group-specific host query messages.	*igmp last-member-query-interval*	—
Configure the interval at which the router sends IGMP host query messages.	*igmp query-interval*	—
Configure the maximum response time specified in IGMP queries.	*igmp query-max-response-time*	—
Configure the IGMP robustness variable.	*igmp robust*	The group membership interval, other query present interval, startup query count, and last member query count are all determined by the robustness variable.

Table 1 Configuring IGMP

Task	Root Command	Notes
Configure the interface to operate in either IGMP version 1, version 2, or version 3 mode.	*igmp version*	—
Configure the estimated bandwidth required by each of the specified groups.	*igmp group-bandwidth*	Enter this command in context configuration mode. Before configuring the group bandwidth, determine the rate at which senders transmit multicast data for each group. You can use inbound rate limiting to ensure that the groups' recommended bandwidth is not exceeded.
Configure the total maximum bandwidth allowed for multicast data traffic on a port or channel.	*igmp maximum-bandwidth*	Enter this command in ATM, AU-3, port, or STM-1 configuration mode. When the addition of a new group would cause the bandwidth usage on this port to exceed the maximum bandwidth, if an IGMP subscriber with a lower-priority exists on this port, the lower priority subscriber's group subscription is terminated to reclaim the bandwidth; otherwise, the new IGMP join request is rejected.
Ensure that all mtrace queries are received within the administratively scoped domain of the router.	*igmp mtrace-prohibit*	Enter this command in context configuration mode.
Configure an IGMP service profile.	For the complete list of tasks used to configure an IGMP service profile, see Configuring an IGMP Service Profile.	—
Enable the specified IGMP service profile on the interface.	*igmp service-profile*	Enter this command in context configuration mode.

2.2 Configuring an IGMP Service Profile

To configure an IGMP service profile, perform the tasks described in Table 2. Enter all commands in IGMP service profile configuration mode, unless otherwise noted.

Table 2 Configuring a Service Profile

Task	Root Command	Notes
Create a service profile, and access IGMP service profile configuration mode.	*igmp service-profile*	Enter this command in context configuration mode.
Enable Instant Leave on the interface.	*instant-leave*	Instant Leave allows IGMP to perform a 0-delay leave upon receiving an IGMPv2 leave message. If the router is an IGMP querier, it sends an IGMP last member query with a 100 millisecond last member query response time; however, the router does not wait for 100 milliseconds before it prunes off the group. This allows channel surfing applications to function more efficiently.
Configure the maximum number of IGMP-joined groups allowed per interface.	*max-groups*	If the addition of a new group on a circuit causes the total number of joined groups to exceed the maximum number allowed, one of the following actions is taken: • If the `drop-old` keyword is specified for the service profile, the oldest IGMP group on the circuit is dropped and the new IGMP report accepted. • If the `drop-old` keyword is not specified for the service profile, the new IGMP membership report is dropped.

Table 2 Configuring a Service Profile

Task	Root Command	Notes
SSM map source address and group access list. Translates IGMPv1, v2 membership reports to IGMPv3.	*ssm-map*	An example of the use of this command is found on the *ssm-map* page. Use the same parameter values for the `ssm-map` command when it is entered in the IGMP service profile and IGMP snooping.
Allow QoS adjustments for Remote Multicast Replication (RMR) on the traffic of a circuit.	*multicast adjust-qos-rate*	Configuring QoS adjustment allows you to avoid congestion when multicast traffic is inserted downstream of the SmartEdge router. You can apply QoS adjustment to metering and queuing bindings.
Set the delay between the subscriber leaving the multicast channel and the removal of QoS adjustment.	*multicast adjust-qos-rate delay-interval*	—
Enable the forwarding of multicast data for IGMP messages received on the PPPoE subscriber circuits on an out-of-band (separated from the PPPoE circuit) IPoE interface.	*multicast destination*	The IGMP service profile must be bound to a subscriber record through a configuration or a RADIUS attribute. For the multicast destination command to work properly, the out-of-band IPoE interface on which the multicast data is to be forwarded must be multicast-enabled; use the `multicast output` command (in interface configuration mode) to enable the out-of-band IPoE interface to forward multicast data.

Table 2 Configuring a Service Profile

Task	Root Command	Notes
Configure the priority of the interface when the maximum bandwidth in the service profile has been exhausted.	*priority (IGMP)*	When the addition of a new group would cause the maximum bandwidth, as specified by the **igmp maximum-bandwidth** command, to be exceeded on the port, the router searches for subscribers joined on the same port with a lower priority. The router drops the lower priority subscribers and reclaims their bandwidth until it gets enough bandwidth to service the higher priority subscriber. If it cannot reclaim enough bandwidth, the new group join will be rejected.

Table 2 Configuring a Service Profile

Task	Root Command	Notes
Create a static multicast route, (*,G) or (S,G), with a subscriber circuit as the outgoing interface (OIF).	*static-group*	PIM normally creates dynamic multicast routes; the **`static-group`** command allows you to create static multicast routes. An OIF is an outgoing circuit that receives traffic destined for a given multicast group. When you configure the static multicast route in IGMP service profile configuration mode, the OIF is a subscriber circuit. To configure all subscriber circuits on a multibind interface to receive multicast traffic for a specified multicast group, configure the **`static-group`** command in an IGMP service profile that is bound to a subscriber (default) profile.
Enable IGMP groups to be sticky.	*sticky-groups*	Groups defined by the ACL will never be dropped, unless an explicit leave for that group is received.

2.3 Configuring PIM-DM

To configure PIM-DM, perform the tasks described in Table 3. Enter the commands in interface configuration mode.

Table 3 Configuring PIM-DM

Task	Root Command	Notes
Enable PIM-DM on an interface.	*pim dense-mode*	—
Enable the origination of PIM-DM State Refresh control messages.	*pim state-refresh origination-interval*	The PIM-DM State Refresh feature keeps pruned branches from being automatically restored to the PIM-DM network by periodically forwarding control messages down the broadcast tree. The control messages refresh the prune state on the outgoing interfaces of each router in the broadcast tree. Enabling this feature is useful in situations in which restoring previously pruned branches consumes too much bandwidth by reflooding unwanted traffic over the PIM-DM network.

2.4 Configuring PIM-SM

To configure PIM-SM, perform the tasks described in Table 4. Enter all commands in interface configuration mode, unless otherwise noted.

Table 4 Configuring PIM-SM

Task	Root Command	Notes
Enable PIM-SM on an interface.	*pim sparse-mode*	—
Configure an administrati vely scoped boundary for multicast routing.	*ip multicast boundary*	An administratively scoped boundary prevents forwarding of multicast data packet destined for group addresses denied by the ACL.

Table 4 Configuring PIM-SM

Task	Root Command	Notes
Accept or reject an IP address as being a valid RP address for a specific multicast group.	*pim accept-rp*	Enter this command in context configuration mode. To determine if the RP should be accepted, the router checks the group-to-RP mapping cache for a matching entry for the group. If there is a matching entry, and the acl-name argument is specified, the router compares the RP address to the ACL to determine if the filter permits the RP address.
Enable anycast RP functionality on a PIM-SM router.	*pim anycast-rp*	Enter this command in context configuration mode.
Configure the router to neither send nor receive BSR messages.	*fast-hello*	This command should be configured on routers that connect to bordering PIM domains to create a PIM domain boundary that blocks the flow Protocol Independent Multicast Version 2 (PIMv2) BSR messages across the domain border.
Configure a router to begin serving as a C-BSR, and participate in the BSR election process.	*pim bsr-candidate*	Enter this command in context configuration mode. If this router wins the BSR election, all candidate RPs will advertise their candidacy to this router. The BSR caches and advertises the RP sets via the PIM bootstrap messages to the entire PIM domain.
Specify the election priority value for a DR.	*pim dr-priority*	—
Set the PIMv2 Hello interval.	*pim hello-interval*	Range is from 10 to 1800 seconds; the default interval is 30 seconds.
Filter PIM messages from neighbors.	*pim neighbor-filter*	—

Table 4 Configuring PIM-SM

Task	Root Command	Notes
Set the protocol parameters to be compatible with PIM-SM specifications, or to be compatible with legacy implementations, such as traditional Cisco implementations.	*pim operation-mode*	—
Configure a router with the RP address for all multicast group addresses permitted by an ACL.	*pim rp-address*	Enter this command in context configuration mode. The **`pim rp-address`** command is usually used on very simple PIM-SM networks where the RP address is manually configured on each router in the network. More complicated networks should use PIMv2's Bootstrap Router feature which allows routers on a network to dynamically learn the RP address. If an ACL is not specified, this RP address is used for the entire multicast address space.
Configure a C-RP on an interface for group address ranges permitted by an ACL.	*pim rp-candidate*	Enter this command in context configuration mode. If an ACL is not specified, this RP address is used for the entire multicast address space.

Table 4 Configuring PIM-SM

Task	Root Command	Notes
Enable a PIM-SM leaf router to continue using a shared tree, instead of switching to an SPT.	*pim spt-threshold infinity*	Enter this command in context configuration mode.
Create a static multicast route, (*,G) or (S,G), with the specified interface as the outgoing interface (OIF).	*pim static group*	Enter this command in context configuration mode. PIM normally creates dynamic multicast routes; the **`pim static group`** command allows you to create static multicast routes. An OIF is an outgoing circuit that receives traffic destined for a given multicast group. For this command, the OIF is a regular interface. For multibind interface OIFs, configure the **`static-group`** command in an IGMP service profile that is bound to a subscriber (default) profile.

2.5 Configuring PIM-Dual Join Mode

To configure PIM-Dual Join Mode, perform the tasks described in Table 5. Enter all commands in interface configuration mode, unless otherwise noted.

Table 5 Configuring PIM-Dual Join Mode

Task	Root Command	Notes
Enter PIM-Dual Join Mode.	*pim dual-join*	Include the **`source`** ***`ip_addr`*** construct to specify the address of the device generating packets to pass through a link in PIM-Dual Join Mode.
Specify the interface over which the SmartEdge router performs a reverse forwarding check on incoming IP multicast packets.	*Examples*	—
Specify the interfaces connected to the far-end device.	*oif-interface*	—

2.6 Configuring MSDP

To configure MSDP, perform the tasks described in Table 6. Enter all commands in MSDP router configuration mode, unless otherwise noted.

Table 6 Configuring MSDP

Task	Root Command	Notes
Enable MSDP within a context, and access MSDP router configuration mode.	*router msdp*	Enter this command in context configuration mode.
Configure a default peer from which to accept all MSDP source active (SA) messages.	*default-peer*	A default peer is needed in topologies where MSDP peers do not coexist with BGP peers. In such a case the reverse path forwarding (RPF) check on SAs may fail, and no SAs will be accepted. In these cases you can configure the peer as a default peer, and bypass RPF checks. An MSDP peer must already be configured before it can be made a default peer.
Configure an MSDP peer to be a member of a mesh group.	*mesh-group*	The MSDP mesh group is a mechanism to reduce SA flooding. Peers in the same mesh group will not forward SA messages to each other. The originator will send the SAs to all its peers.
Configure an interface as the originating RP address.	*originating-rp*	The IP address of the interface is used as the RP address in all SAs originated by the router.
Configure an ACL to filter incoming SA messages learned from the local RP.	*originating-rp sa-filter*	—
Configure an MSDP peer.	For the complete list of tasks used to configure an MSDP peer, see Configuring an MSDP Peer.	—

2.7 Configuring an MSDP Peer

To configure an MSDP peer, perform the tasks described in Table 7. Enter all commands in MSDP peer configuration mode, unless otherwise noted.

Table 7 Configure an MSDP Peer

Task	Root Command	Notes
Create an MSDP peer and enter MSDP peer configuration mode for peer-specific configurations.	*peer (MSDP)*	Enter this command in MSDP router configuration mode. The `no shutdown` command is enabled by default after you configure an MSDP peer.
Associate a text description with an MSDP peer.	*description (MSDP peers)*	—
Configure a peer's autonomous system (AS) number.	*peer-as*	—
Configure an ACL to filter SA messages coming from another peer.	*sa-filter*	Use the following command syntax: **`sa-filter in`** ***`acl-name`***
Configure an ACL to filter SA messages going to another peer.	*sa-filter*	Use the following command syntax: **`sa-filter out`** ***`acl-name`***
Disable a configured MSDP peer.	*shutdown (MSDP peers)*	—

2.8 Configuring Multicast for Subscribers

To configure multicast for subscribers, perform the tasks described in Table 8. Enter all commands in subscriber configuration mode.

Consider the following restrictions before configuring multicast for subscribers:

- PIM-SM is not supported on multibind interfaces. If you are configuring multicast routing on subscribers, you must first use the **`pim sparse-mode passive`** command in interface configuration mode to prevent PIM messages from being exchanged on the egress interface, while allowing the interface and its circuits to be populated in a multicast forwarding entry by receiving an IGMP report or a data packet.

- If the multicast source uses router-mode CPE, the SmartEdge router runs in active PIM-SM only.

- If the multicast receiver uses router-mode CPE, the SmartEdge router can run in active PIM-SM or passive PIM-SM.

- If bridge-mode CPE exists between the multicast source or receiver and the SmartEdge router, the SmartEdge router can run in active PIM-SM or passive PIM-SM.

Table 8 *Configuring Multicast for Subscribers*

Task	Root Command	Notes
Enable an existing IGMP service profile on a single subscriber record, a named subscriber profile, or a default subscriber profile.	*ip igmp service-profile*	The service profile used is determined in the following order: • Subscriber profile • Default subscriber profile • Service profile configured on the subscriber's parent interface If a service profile is not defined in the subscriber record, it inherits the service profile from the default subscriber profile. If the default subscriber profile is not configured with an service profile, the service profile configured on the interface is used.

Table 8 Configuring Multicast for Subscribers

Task	Root Command	Notes
Configure the multicast receive permissions for a subscriber record or for the default subscriber record.	*ip multicast receive*	Permission attributes are applied in the following order: • Subscriber record • Default subscriber record • System defaults If a permission is not defined in the subscriber record, it inherits the value of the permission from the default subscriber record. If the permission is not defined in the default subscriber record, the system default values are used. For multicast routing to function on subscribers, you must use the **`pim sparse-mode passive`** command in interface configuration mode to enable PIM-SM on the interface.
Configure the multicast send permissions for a subscriber record or for the default subscriber record.	*ip multicast send*	If the **`permit`** keyword is used without the **`unsolicit`** keyword, the subscriber must join a group prior to sending unsolicited multicast data. If used together (**`permit unsolicit`**), a subscriber is allowed to send unsolicited multicast traffic. Permissions are examined in the following order: • Subscriber record • Default subscriber record • System defaults. If a permission is not defined in the subscriber record, it inherits the value of the permission from the default subscriber record. If the permission is undefined in the default subscriber record, the system default values are used. For multicast routing to function on subscribers, you must use the **`pim sparse-mode passive`** command in interface configuration mode to enable PIM-SM on the interface.

2.9 Enabling PIM Graceful Restart

PIM graceful restart allows the SmartEdge router and its neighbors to continue forwarding multicast packets without disrupting network traffic. Because neighboring routers assist, the SmartEdge router can quickly restart the PIM process without having to recalculate algorithms from scratch. To enable PIM graceful restart, perform the task described in Table 9. Enter the command in context configuration mode.

Table 9 Enabling PIM Graceful Restart

Task	Root Command	Notes
Enable PIM graceful restart on the specified context.	*pim graceful-restart*	—

2.10 Enabling PIM SSM

To enable SSM, perform the task described in Table 10. Enter the command in context configuration mode.

Table 10 Enabling SSM

Task	Root Command	Notes
Enable SSM routing on the specified context and specify the address range to be reserved for SSM.	*pim ssm*	Use the **`default`** keyword to use the default address range for SSM or create an access control list to specify the SSM address range.

2.11 Configuring PIM-SSM Auto-Discovery

To enable PIM-SSM auto-discovery, both PIM and BGP must be configured. The BGP configuration is divided across local and VPN contexts.

- In the local context, all participating PE routers must be enabled for the MDT address family.
- In the VPN context, the MDT address family must be enabled. RTs may optionally be used for importing and exporting MDT routes.

To configure PIM-SSM auto-discovery, perform the tasks described in Table 11.

Table 11 Enabling PIM-SSM Auto-Discovery

Task	Root Command	Notes
Enable SSM routing in the local context and specify the address range to be reserved for SSM. For PIM-SSM auto-discovery of MDT, the SSM range must be enabled in the local context.	*pim ssm*	Use the **`default`** keyword to reserve the default address range for SSM, or create an access control list to specify the SSM address range.
Configure the MDT default group for the mVPN context.	*mdt default-group*	Enter this command in VPN-RD context configuration mode.
Set the encapsulation for MDT.	*mdt encapsulation*	Enter this command in interface configuration mode.
Enable the MDT address family for BGP.	*address-family ipv4 mdt*	Specify the address family for BGP router instances and BGP neighbors, in or out of VPN context.
Direct the local router to export MDT routes.	*export route-target (BGP VPN)*	Enter this command in router BGP VPN mode.
Direct the local router to import MDT routes.	*import route-target (BGP VPN)*	Enter this command in router BGP VPN mode.

2.12 Disabling BFD for a PIM Interface

By default, BFD is enabled on PIM interfaces and for each neighbor on the interface. To disable BFD for a PIM interface, perform the Table 12.

Table 12 Disabling BFD for a PIM Interface

Task	Root Command	Notes
Access interface configuration mode for the PIM interface on which you want to disable BFD.	*interface (BFD)*	Enter this command in context configuration mode.
Disable BFD for the PIM interface.	*no pim bfd*	For examples and more information about BFD, see *Configuring BFD*.

2.13 Enabling Multicast VPNs

Multicast VPNs use MDTs on PE routers to support IP multicast over BGP/MPLS VPNs. To enable multicast VPNs, perform the task described in Table 13. Enter both commands in interface configuration mode.

Table 13 Enabling Multicast VPNs

Task	Root Command	Notes
Specify the default MDT group.	*mdt default-group*	Configure this command on an intercontext interface in a VPN context. This interface is similar to a loopback interface in that it is not bound to anything and does not need an IP address. It creates an intercontext circuit between the VPN context and the local context. PIM-SM must also be configured on this intercontext interface.
Specify the multicast MDT encapsulation type.	*mdt encapsulation*	Configure this command on a loopback interface in the local context. The loopback interface is used to source multicast packets on the MDT.

2.14 Enabling RMR

Remote multicast replication (RMR) is used to enable IP multicast services. To enable RMR, perform the task described in Table 14.

Table 14 Enabling Multicast VPNs

Task	Root Command	Notes
Enable an interface to forward multicast data and to send and receive IGMP control messages.	*multicast output*	Enter this command in interface configuration mode. An IP over Ethernet (IPoE) circuit, on a Gigabit Ethernet port or an 802.1Q permanent virtual circuit (PVC) configured on it, must be configured on the interface to carry the multicast services. The MAC addresses received from IGMP control packets on the IPoE circuit are compared to the subscriber's MAC address received on the corresponding PPPoE circuit. By default, if the MAC addresses do not match, the IGMP control packet is dropped. Use the **`accept-unknown-mac`** keyword to accept IGMP control packets that have MAC addresses that do not match the subscriber's MAC address.
Enable the forwarding of multicast data for IGMP messages received on the PPPoE subscriber circuits on an out-of-band (separated from the PPPoE circuit) IPoE interface.	*multicast destination*	Enter this command in IGMP service profile configuration mode. The IGMP service profile must be bound to a subscriber record through a configuration or a RADIUS attribute.

2.15 Configuring QoS Adjustments for RMR

This section describes how to configure QoS adjustments for remote multicast replication (RMR). To enable QoS adjustment for RMR, perform the tasks described in Table 15.

Table 15 Configuring QoS Adjustments for RMR

Step	Task	Root Command	Notes
1.	Create an ACL to specify the matching criteria for one or more multicast groups, and enter access control list configuration mode to add options and rules.	*ip access-list*	Enter this command in context configuration mode.
2.	Create an ACL statement using permit conditions for each multicast group.	*permit*	To explicitly set its order in the list, use the seq permit construct for each statement.
3.	Create an ACL statement using deny conditions for each multicast group.	*deny*	To explicitly set its order in the list, use the seq deny construct for each statement.
4.	Configure the adjustment bandwidth and other adjustment parameters for each multicast group.	*igmp group-bandwidth*	Enter this command in context configuration mode, specifying the name of an ACL created in step 1 for the **`group-list`** ***`acl-name`*** construct and the **`qos-adjust`** keyword to indicate that joins to this IGMP group should result in QoS adjustments.
5.	In the IGMP profile of a subscriber, enable QoS adjustments for the subscriber.	*multicast adjust-qos-rate*	Enter this command in IGMP service profile configuration mode. In an IGMP service profile, up to two instances of this command can be configured: one for metering and one for queuing.
6.	Adjust the default delay interval to use when processing multicast leaves for the subscriber.	*multicast adjust-qos-rate delay-interval*	Enter this command in IGMP service profile configuration mode.

2.16 Configuring IGMP Snooping

To configure IGMP on your router, perform the tasks in the sections that follow.

2.16.1 Enabling IGMP Snooping on an Ethernet Bridge

This section describes how to enable IGMP snooping on an Ethernet bridge with the default IGMP snooping configuration settings, as described in Table 16.

Table 16 Ethernet Bridge IGMP Snooping Default Configuration

Parameter	Command Used to Modify Default Setting	Default
IGMP proxy reporting	*proxy-reporting*	disabled
Groups statically joined to the bridge	*join-group*	No groups are statically joined to this bridge
Interval at which IGMP group-specific host query messages are sent	*query-interval*	125 seconds
Number of IGMP packets that can be lost before the bridge goes down	*robust*	2
Fast IGMP message processing in the PPA	*turbo*	disabled
IGMP snooping version used by the bridge	*version*	IGMP version 2

To modify the default IGMP snooping configuration settings on an Ethernet bridge, see Modifying the Default IGMP Snooping Parameters on a Bridge.

To enable IGMP snooping on a specific Ethernet bridge, perform the tasks described in Table 17.

Table 17 Enabling IGMP Snooping on an Ethernet Bridge

#	Task	Root Command	Notes
1	Access global configuration mode.	*configure*	Enter this command in exec mode.
2	Access context configuration mode.	*context*	Replace the ***`ctx-name`*** argument with the name of the context containing the bridge on which you want to enable IGMP snooping.
3	Access bridge configuration mode for a specified bridge.	*bridge*	Replace the ***`bridge-name`*** argument with the name of the Ethernet bridge on which you want to enable IGMP snooping.
4	Enable IGMP snooping on an Ethernet bridge.	*igmp snooping*	—

Table 17 Enabling IGMP Snooping on an Ethernet Bridge

#	Task	Root Command	Notes
5	Verify that IGMP snooping is enabled on the bridge you configured in Task 1 through Task 4.	*show igmp snooping bridge*	—
If you want to change the default IGMP configuration on the Ethernet bridge, see Modifying the Default IGMP Snooping Parameters on a Bridge. If you want to create an IGMP snooping profile, see Creating and Configuring an IGMP Snooping Profile.			

2.16.2 Modifying the Default IGMP Snooping Parameters on a Bridge

To modify the default IGMP snooping parameters on an Ethernet bridge, perform the tasks described in Table 18.

Table 18 Modifying the Default IGMP Snooping Parameters on an Ethernet Bridge

Task	Root Command	Notes
Access global configuration mode.	*configure*	Enter this command in exec mode.
Access context configuration mode.	*context*	Replace the ***ctx-name*** argument with the name of the context containing the bridge on which you want to enable IGMP snooping.
Access bridge configuration mode.	*bridge*	Replace the ***bridge-name*** argument with the name of the Ethernet bridge whose IGMP snooping parameters you want to modify.
Access IGMP snooping configuration mode.	*igmp snooping*	
Enable the proxy reporting feature on the Ethernet bridge.	*proxy-reporting*	When the proxy reporting feature is enabled, the SmartEdge router suppresses unnecessary IGMP messages, but does not generate new IGMP messages or manipulate forwarded IGMP messages. IGMP proxy reporting is disabled by default.

Table 18 Modifying the Default IGMP Snooping Parameters on an Ethernet Bridge

Task	Root Command	Notes
Statically join (*,G) or (S,G) bridges to a group, regardless of the number of circuits already joined to the group.	*join-group*	Replace the ***ip-address*** argument with the IP address of the group to which (*,G) and (S,G) bridges will be statically bound. Enter the IP address in the format ***A.B.C.D***. To specify a source for the multicast group, Include the optional >**source** ***source-address*** construct in the **join-group** command. Replace the ***source-address*** argument with the IP address of an appropriate source.
Modify the IGMP query interval.	*query-interval*	Replace the ***interval*** argument with the interval at which IGMP group-specific host query messages are sent. Range is from 1 through 65535 seconds.
Disable the generation of IGMP query solicitation messages by a bridge.	*no query-solicitation*	The generation of IGMP query solicitation messages by a bridge is enabled by default. (1)
Modify the number of IGMP packets that can be lost before the multicast group membership of a bridge expires.	*robust*	Replace the ***packet-number*** argument with the expected packet loss for this bridge. Range is from two to seven packets. If a bridge is particularly susceptible to packet losses, we recommend increasing this value. The default expected packet loss is two packets.
Sets the SSM map source address and group access list. Translates IGMPv1, v2 membership reports to IGMPv3.	*ssm-map*	An example of the use of this command is found on the *ssm-map* page. Use the same parameter values for the `ssm-map` command when it is entered in the IGMP service profile and IGMP snooping.
Enable fast IGMP message processing in the PPA.	*turbo*	Fast IGMP message processing is disabled by default.

Table 18 Modifying the Default IGMP Snooping Parameters on an Ethernet Bridge

Task	Root Command	Notes
Set the IGMP snooping version used by this bridge.	*version*	Enter `2` to select IGMP version 2, or `3` to select IGMP version 3. By default, all bridges use IGMP version 2.
Commit all configuration changes.	*commit*	—
If you want to create an IGMP snooping profile, see Creating and Configuring an IGMP Snooping Profile.		

(1) Use the ***query-solicitation*** *command to reenable the generation of IGMP query solicitation messages on a bridge.*

2.16.3 Creating and Configuring an IGMP Snooping Profile

This section describes how to create and configure an IGMP snooping profile.

When a circuit that belongs to a bridge that has IGMP snooping configured that is not bound to an IGMP snooping profile, then IGMP snooping runs in the default mode, as described in Table 19.

Table 19 IGMP Snooping Profile Default Configuration

Parameter	Command Used to Modify Default Setting	Default
Static memberships for a multicast group	*static*	Not configured
IGMP control message filtering and access-list match counting	`access-group` command See *access-group (IGMP snooping profile configuration mode)*.	Not configured
Maximum number of IGMP-joined groups allowed for each interface	*max-groups*	Unlimited
SSM map source address and group access list.	*ssm-map*	Not configured
IGMP router monitoring	*mrouter*	Enabled
Receipt of multicast data	*receive*	Enabled
Sending of multicast data	*send*	Permitted

To create an IGMP snooping profile that customizes the default IGMP snooping configuration data, perform the tasks described in Table 20. When the IGMP snooping profile is applied to a bridge circuit, that circuit assumes the configuration settings from the IGMP snooping profile.

Table 20 Creating an IGMP Snooping Profile

Task	Root Command	Notes
Enter global configuration mode.	*configure*	Enter this command in exec mode.
Create an IGMP profile and enter IGMP snooping profile configuration mode.	*igmp snooping profile*	Replace the ***profile-name*** argument with a name that identifies this IGMP profile.
Add a static membership to a multicast group.	*static*	Replace the ***ip-address*** argument with the IP address of the multicast group you want a circuit to join. Use the optional >**source** ***source-address*** construct to specify a particular router as a source for the multicast group.
Specify the SSM map source address and group access list. Translates IGMPv1, v2 membership reports to IGMPv3.	*ssm-map*	An example of the use of this command is found on the *ssm-map* page. Use the same parameter values for the `ssm-map` command when it is entered in the IGMP service profile and IGMP snooping.
Specify an access group whose messages you want to filter on the associated circuit.	**access-group** command See *access-group (IGMP snooping profile configuration mode)*.	Replace the ***group-name*** argument with the name of the access group whose messages you want to filter.
Limit the number of (*,G) and (S,G) groups an associated circuit is allowed to join.	*max-groups*	Replace the ***number*** argument with the number of (*,G) and (S,G) groups that a circuit is allowed to join. Range is 1 to 100000. By default, a circuit is allowed to join an unlimited number of (*,G) and (S,G) groups.

Table 20 Creating an IGMP Snooping Profile

Task	Root Command	Notes
Enables multicast router monitoring for circuits attached to the specified IGMP snooping profile.	*mrouter*	Enter the **mrouter** command without any optional keywords to enable IGMP router monitoring on a circuit. When IGMP queries are received by an associated circuit, that circuit is declared to be a multicast router circuit. Include the **static** keyword to disable IGMP router monitoring on an associated circuit. No monitoring or active discovery is performed on this circuit. A circuit is always assumed to be a multicast router circuit. By default, every circuit has IGMP router monitoring enabled, unless it is overwritten by the **mrouter** setting in the bridge profile.
Disable the receipt of multicast data by an associated circuit.	*receive deny*	To return the circuit to the default setting, in which the receipt of multicast data is permitted, use the **receive permit** command.
Disable the sending of multicast data by an associated circuit.	*send deny*	To return the circuit to the default setting, in which the sending of multicast data is permitted, use the **send permit** command.
Save your changes to the IGMP snooping profile.	*commit*	—
Add the IGMP snooping profile you created to a particular Ethernet bridge, as described in Adding an IGMP Snooping Profile to a Bridge Profile.		

2.16.4 Adding an IGMP Snooping Profile to a Bridge Profile

To add an IGMP snooping profile to a bridge profile, perform the tasks described in Table 21.

Table 21 Adding an IGMP Snooping Profile to a Bridge Profile

Task	Root Command	Notes
Access global configuration mode.	*configure*	Enter this command in exec mode.

Table 21 Adding an IGMP Snooping Profile to a Bridge Profile

Task	Root Command	Notes
Access a bridge profile.	*bridge profile*	Replace the ***prof-name*** argument with a name that identifies this bridge profile to which you want to add the IGMP snooping profile.
Add the IGMP snooping profile to the bridge profile, and applies those settings to all circuits that belong to that bridge profile.	*igmp snooping profile*	Replace the ***profile-name*** argument with the name of the IGMP snooping profile you want to add to the bridge profile.
To add the bridge profile to a VPLS profile, see Adding a Bridge Profile to a VPLS Profile. To enable IGMP snooping directly on an Ethernet circuit, enable IGMP snooping on the desired Ethernet bridge, as described in Enabling IGMP Snooping on an Ethernet Bridge.		

2.16.5 Adding a Bridge Profile to a VPLS Profile

Note: Perform this task only if you want to enable IGMP snooping on a VPLS instance. If you want to enable IGMP snooping on an Ethernet bridge, perform the tasks described in Adding an IGMP Snooping Profile to a Bridge Profile.

To add a bridge profile to an existing VPLS profile, perform the tasks described in Table 22.

Table 22 Adding a Bridge Profile to a VPLS Profile

Task	Root Command	Notes
Access global configuration mode.	*configure*	Enter this command in exec mode.
Enter VPLS profile configuration mode for a VPLS profile.	*vpls profile*	Replace the ***prof-name*** argument with the name of the VPLS profile to which you want to add the bridge profile.
Enter VPLS neighbor configuration mode for a VPLS neighbor.	*neighbor (VPLS)*	Replace the ***ip-addr*** argument with the IP address of the VPLS neighbor to which you want to add the bridge profile.
Add a bridge profile to a VPLS profile for a particular VPLS neighbor. Because the bridge profile references the IGMP snooping profile, all settings in the IGMP snooping profile are applied to all circuits that belong to that bridge profile.	*bridge profile*	Replace the ***prof-name*** argument with the name of the bridge profile you want to add to the VPLS profile.

2.17 Enabling IGMP General Query Messages in Response to IGMP Query Solicitation Messages on an Interface

To enable the generation of IGMP query response messages on an interface, perform the tasks described in Table 23.

Table 23 Enabling the Generation of IGMP Query Response Messages on an Interface

Task	Root Command	Notes
Access global configuration mode.	*configure*	Enter this command in exec mode.
Access context configuration mode.	*context*	—
Enter interface configuration mode.	*interface*	—
Enable the generation of IGMP general query response messages on an interface.	*igmp query solicitation*	For IGMP query solicitation to work, IGMP query solicitation must be enabled on the interface and on the bridge where query messages are exchanged. IGMP query solicitation is enabled on all bridges by default. For more information on enabling and disabling IGMP query solicitation on a bridge, see Modifying the Default IGMP Snooping Parameters on a Bridge.

2.18 Configuring Packet Replication for High-Bandwidth Multicast Traffic

You can configure line cards to enable packet mesh ASIC (PMA) traffic replication for increased multicast bandwidth.

Table 24 Configuring Optimization of Packet Replication for Multicast Traffic

Task	Root Command	Notes
Enable PMA traffic replication on the current line card to the destination line cards.	*optimize replication*	Enter this command in card configuration mode. Use the **`optimize replication`** command for the efficient support of high-bandwidth multicast services such as IPTV. When not enabled, packet replication is performed in the ingress packet processing ASICs (iPPAs).

2.19 IP Multicast Operations

To monitor, troubleshoot, and administer IP multicast, perform the tasks described in the sections that follow.

2.19.1 IGMP

To manage IGMP functions, perform the appropriate tasks described in Table 25. Enter the **`show`** commands in any mode; enter the **`clear`** and **`debug`** commands in exec mode.

Table 25 IGMP Operations Tasks

Task	Root Command
Clear all dynamically learned IGMP groups from the IGMP cache table.	*clear igmp group*
Clear IGMP specified rows in the snooping routing table.	*clear igmp snooping*
Clear all traffic statistics maintained by IGMP.	*clear igmp traffic*
Enable the generation of IGMP debug messages.	*debug igmp*
Display the current IGMP configuration information for the current context.	*show configuration igmp*
Display the configured IGMP bandwidth profiles for ports.	*show igmp bandwidth-profile*
Display circuit-specific information for the IGMP.	*show igmp circuit*
Display IGMP-connected group membership information.	*show igmp group*
Display bandwidth recommendations for multicast groups. Display includes whether a group requires QoS rate adjustment or not.	*show igmp group-bandwidth*

Table 25   IGMP Operations Tasks

Task	Root Command
Display information about IGMP interfaces.	*show igmp interface*
Display service profile information, or bandwidth usage for all circuits.	*show igmp profile*
Displays information about a specified access list that is associated with an IGMP snooping instance.	*show igmp snooping access-group name*
Displays IGMP snooping information for a specific bridge interface or all bridge interfaces that are currently configured on the router.	*show igmp snooping bridge*
Displays IGMP snooping-related information about circuits that are bound to bridge interfaces that have IGMP snooping enabled.	*show igmp snooping circuit*
Displays a per-bridge list of IGMP groups and their associated circuits.	*show igmp snooping group*
Displays per-bridge list of circuits that are facing multicast routers.	*show igmp snooping mrouter*
Display information about IGMP traffic statistics.	*show igmp traffic*

2.19.2 MSDP

To manage MSDP functions, perform the appropriate tasks described in Table 26. Enter the **`show`** commands in any mode; enter the **`clear`** and **`debug`** commands in exec mode.

Table 26   MSDP Operations Tasks

Task	Root Command
Clear the connection to MSDP peers.	*clear msdp peer*
Clear MSDP SA cache entries from the router.	*clear msdp sa-cache*
Clear all MSDP peer statistics (counters).	*clear msdp statistics*
Enable the generation of MSDP debug messages.	*debug msdp*
Display the current MSDP configuration information for the current context.	*show configuration msdp*
Display information about configured MSDP peers.	*show msdp peer*
Display information about MSDP SA messages cached on the router.	*show msdp sa-cache*
Display summary information about configured MSDP peers.	*show msdp summary*

2.19.3 PIM

To manage PIM functions, perform the appropriate tasks described in Table 27. Enter the **show** commands in any mode; enter the **clear**, **debug**, **mrinfo**, and **mtrace** commands in exec mode.

Table 27 PIM Operations Tasks

Task	Root Command
Delete source and group entries from the PIM routing table.	*clear ip mroute*
Clear the connection to MSDP peers.	*clear msdp peer*
Clear all dynamically learned PIM RP mappings in the local database.	*clear pim rp*
Clear all traffic statistics maintained by PIM.	*clear pim traffic*
Enable the generation of debug messages for multicast routing table entries.	*debug ip mrouting*
Enable the generation of PIM debug messages.	*debug pim*
Enable the generation of PIM BFD debug messages.	*debug pim bfd*
Enable the generation of PIM multicast data packet debug messages.	*debug pim packet*
Enable the generation of debug messages for PIM routes downloaded to Packet Processing ASICs (PPAs).	*debug pim ppa*
Enable the generation of bootstrap router (BSR) and rendezvous point (RP) debug messages.	*debug pim rp-mapping*
Query a neighboring multicast router to determine which routers are peers of the local router.	*mrinfo*
Trace the path from a source to a destination branch on a multicast distribution tree.	*mtrace*
Display the current PIM configuration information for the current context.	*show configuration pim*
Display the PIM routing table.	*show ip mroute*
Display BSR and candidate RP (C-RP) information.	*show pim bsr-router*
Display circuit-specific information for the PIM.	*show pim circuit*
Display information about PIM-enabled interfaces.	*show pim interface*
Display multicast domain tree (MDT) information.	*show pim mdt*
Display information about PIM neighbors.	*show pim neighbor*
Display information about the PPA state from the PIM perspective.	*show pim ppa*
Display information about the RP to which the specified group hashes (maps).	*show pim rp mapping*

Table 27 PIM Operations Tasks

Task	Root Command
Display information about the group-to-RP mapping cache.	*show pim rpf*
Display reverse path forwarding (RPF) information for a specified multicast source.	*show pim rp-hash*
Display the Source-Specific Multicast (SSM) setting in PIM.	*show pim ssm*
Display PIM traffic statistics.	*show pim traffic*

2.19.4 IP Multicast Manager

The IP multicast manager handles all communications between the control plane and forwarding plane for multicast protocols. The multicast manager receives *, G and S,G states from multicast protocols, maintains a copy of the information it its local database, and sends any updates to the forwarding plane. When the forwarding plane restarts, multicast protocols are not required to provide any updates. Instead, the multicast manager reloads the multicast entries. To manage IP multicast manager functions, perform the appropriate tasks described in Table 28. Enter the **`show`** command in any mode; enter the **`debug`** commands in exec mode.

Table 28 IP Multicast Manager Operations Tasks

Task	Root Command
Enable the generation of debug messages by the IGMP snooper when it is communicating with the IP multicast manager.	*debug ipmul*
Enable the generation of debug messages for IP multicast manager events.	*debug mcast*
Enable Protocol Independent Multicast (PIM) multicast manager debug messages.	*debug pim mcastmgr*
Display routes from the IP multicast manager database.	*show ip mfib*

3 Configuration Examples

The sections that follow provide IP multicast configuration examples.

3.1 PIM-SM

The following example demonstrates how three routers (Router A, Router B, and Router C) are configured to correctly operate on a PIM-SM local network. Figure 5 shows the simple PIM-SM network topology used for the configuration example.

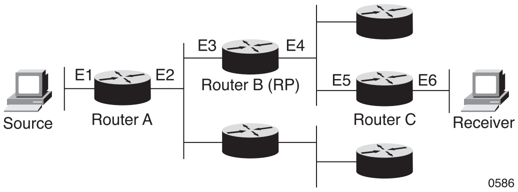

Figure 5 Simple PIM-SM Network Topology

In this example, Router A is directly connected to the source, and Router C is directly connected to the receiver. Because Router A is the only router directly connected to the source, it serves as a PIM DR for the network. If multiple routers were connected to the source, the router with the highest IP address would be selected as the PIM DR.

The **`pim sparse-mode`** interface configuration mode command enables PIM-SM on the interface. The **`pim rp-address`** global configuration mode command enables all routers in the PIM-SM network to statically configure Router B as the rendezvous point (RP). An ACL can be specified with the *rp-addr* argument to permit multicast traffic for a particular group with this RP.

Enabling PIM-SM on an interface also enables IGMP on the same interface. For each local network, an IGMP querier is selected; for example, Router C is the IGMP querier for the local network connected to the receiver. If multiple routers were connected directly to the receiver, the router with the lowest IP address serves as the IGMP querier. The IGMP querier is responsible for sending IGMP host-query messages to all hosts on the local network.

Router A, which is directly connected to the source and the DR for its local network, sends PIM register messages on behalf of the source to the RP. Router C, on behalf of the receiver, sends PIM join and prune messages to the RP to advertise the group membership.

The configuration for **RouterA** is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface E1
[local]Redback(config-if)#ip address 10.2.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E2
[local]Redback(config-if)#ip address 11.1.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#ip access-list 1
[local]Redback(config-access-list)#seq 10 permit 224.0.0.0 15.255.255.255
[local]Redback(config-access-list)#exit
[local]Redback(config-ctx)#pim rp-address 10.2.1.2 1
```

The configuration for **RouterB** (RP) is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface E3
[local]Redback(config-if)#ip address 10.2.1.2/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E4
[local]Redback(config-if)#ip address 10.4.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#ip access-list 1
[local]Redback(config-access-list)#seq 10 permit 224.0.0.0 15.255.255.255
[local]Redback(config-access-list)#exit
[local]Redback(config-ctx)#pim rp-address 10.2.1.2 1
```

The configuration for **RouterC** (IGMP querier) is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface E5
[local]Redback(config-if)#ip address 10.4.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E6
[local]Redback(config-if)#ip address 44.1.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#ip access-list 1
[local]Redback(config-access-list)#seq 10 permit 224.0.0.0 15.255.255.255
[local]Redback(config-access-list)#exit
[local]Redback(config-ctx)#pim rp-address 10.2.1.2 1
```

3.2 PIM-Dual Join

The following example demonstrates how the implementation generates two join requests to the same SmartEdge router to correctly operate on a PIM-Dual Join implementation. To perform this task, you must configure a PIM-Dual Join group.

Figure 6 shows the PIM-Dual Join network topology used for the configuration example:

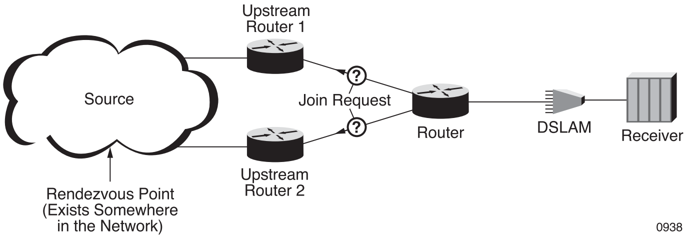

Figure 6   PIM-Dual Join Implementation

The configuration for the PIM-Dual Join implementation is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#pim dual-join group 225.100.1.1 source 192.110.30.6
[local]Redback(config-pim-dual)#rpf-interface int1 int2
[local]Redback(config-pim-dual)#oif-interface int3
[local]Redback(config-pim-dual)#exit
[local]Redback(config-ctx)#exit
```

3.3 MSDP for Two PIM-SM Domains

The following example demonstrates how to configure MSDP to link two PIM-SM domains, using MSDP, so that multicast messages can be forwarded from one domain to the other. Figure 7 shows the PIM-SM interdomain network topology used for the configuration example.

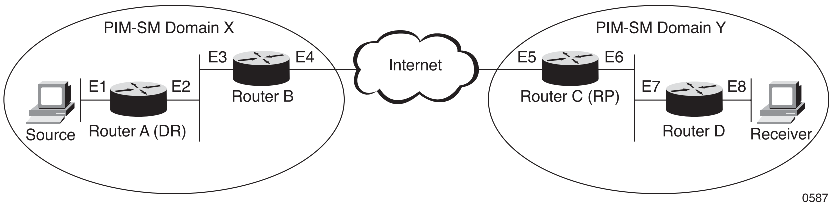

Figure 7   Interdomain PIM-SM Network Topology

This example can be expanded to several PIM-SM domains. Each domain can use BGP for interdomain routing. MSDP is used for interdomain source discovery.

Each PIM-SM domain has one or more RPs that belong to the domain. MSDP allows RPs in different domains to share information about active sources. RPs know about the receivers in their local domain. Because RPs share information about the active sources in each domain, each RP can forward data accordingly if there is an active receiver in their local domain for a particular source.

For RPs to share information with each other, RPs are configured as MSDP peers. There can be multiple peers in between two RP MSDP peers. Each RP establishes an MSDP peering session with another RP in another domain.

To keep this configuration example simple, the following assumptions are made:

- The two domains, Domain X and Domain Y, are externally peered using Multiprotocol BGP (MBGP), thus, Router B and Router C are external MBGP peers and MSDP peers.
- The two domains are different LAN segments.
- Static routing is being used instead of other Internet gateway protocols like Open Shortest Path First (OSPF), internal Border Gateway Protocol (iBGP), Intermediate System-to-Intermediate System (IS-IS), and so on.

The configuration for **RouterA** (DR) is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface lo1 loopback
[local]Redback(config-if)#ip address 10.200.1.1/32
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E2
[local]Redback(config-if)#ip address 102.1.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E4
[local]Redback(config-if)#ip address 11.1.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
```

Static RP for Domain X configuration:

```
[local]Redback(config-ctx)#pim rp-address 10.200.1.2
```

Static route configuration:

```
[local]Redback(config-ctx)#ip route 10.200.1.2/32 102.1.1.2
```

The configuration for **RouterB** is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface lo1 loopback
[local]Redback(config-if)#ip address 10.200.1.2/32
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E1
[local]Redback(config-if)#ip address 102.1.1.2/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E2
[local]Redback(config-if)#ip address 104.1.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
```

Static RP for Domain X configuration:

```
[local]Redback(config-ctx)#ip pim rp-address 10.200.1.2
```

eBGP configuration:

```
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#router-id 10.200.1.2
[local]Redback(config-bgp)#address-family ipv4 multicast
[local]Redback(config-addrfamily)#network 11.1.1.0/24
[local]Redback(config-addrfamily)#exit
[local]Redback(config-bgp)#peer-group eMBGP external
[local]Redback(config-peergroup)#ebgp-multihop 5
[local]Redback(config-peergroup)#update-source lo1
[local]Redback(config-peergroup)#address-family ipv4 unicast
[local]Redback(config-addrfamily)#exit
[local]Redback(config-peergroup)#address-family ipv4 multicast
[local]Redback(config-peergroup)#neighbor 10.200.1.3 external
[local]Redback(config-neighbor)#remote-as 200
[local]Redback(config-neighbor)#peer-group eMBGP
[local]Redback(config-neighbor)#exit
[local]Redback(config-peergroup)#exit
[local]Redback(config-bgp)#exit
```

MSDP configuration—peering between two RPs:

```
[local]Redback(config-ctx)#router msdp
[local]Redback(config-msdp)#peer 10.200.1.3 local-tcp-source lo1
[local]Redback(config-msdp-peer)#no shutdown
[local]Redback(config-msdp-peer)#exit
[local]Redback(config-msdp)#exit
```

Static route configuration:

```
[local]Redback(config-ctx)#ip route 10.200.1.1/32 102.1.1.1
[local]Redback(config-ctx)#ip route 10.200.1.3/32 104.1.1.2
[local]Redback(config-ctx)#ip route 11.1.1.0/24 102.1.1.1
```

The configuration for **RouterC** (RP) is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface lo1 loopback
[local]Redback(config-if)#ip address 10.200.1.3/32
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E2
[local]Redback(config-if)#ip address 104.1.1.2/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E4
[local]Redback(config-if)#ip address 105.1.1.1/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
```

eBGP configuration:

```
[local]Redback(config-ctx)#router bgp 200
[local]Redback(config-bgp)#router-id 10.200.1.3
[local]Redback(config-bgp)#address-family ipv4 multicast
[local]Redback(config-addrfamily)#network 44.1.1.0/24
[local]Redback(config-addrfamily)#exit
[local]Redback(config-bgp)#peer-group eMBGP external
[local]Redback(config-peergroup)#ebgp-multihop 5
[local]Redback(config-peergroup)#update-source lo1
[local]Redback(config-peergroup)#address-family ipv4 multicast
[local]Redback(config-addrfamily)#exit
[local]Redback(config-peergroup)#neighbor 10.200.1.2 external
[local]Redback(config-neighbor)#remote-as 100
[local]Redback(config-neighbor)#peer-group eMBGP
[local]Redback(config-neighbor)#exit
[local]Redback(config-peergroup)#exit
[local]Redback(config-bgp)#exit
```

Static RP for Domain Y configuration:

```
[local]Redback(config-ctx)#ip pim rp-address 10.200.1.3
```

BGP configuration:

```
[local]Redback(config-ctx)#router bgp 200
[local]Redback(config-bgp)#router-id 10.200.1.3
[local]Redback(config-bgp)#address-family ipv4 multicast
[local]Redback(config-addrfamily)#network 44.1.1.0/24
[local]Redback(config-addrfamily)#exit
[local]Redback(config-bgp)#peer-group eMBGP external
[local]Redback(config-peergroup)#ebgp-multihop 5
[local]Redback(config-peergroup)#update-source lo1
[local]Redback(config-peergroup)#address-family ipv4 unicast
[local]Redback(config-addrfamily)#exit
[local]Redback(config-peergroup)#address-family ipv4 multicast
[local]Redback(config-addrfamily)#exit
[local]Redback(config-peergroup)#neighbor 10.200.1.2 external
[local]Redback(config-neighbor)#remote-as 100
[local]Redback(config-neighbor)#peer-group eMBGP
[local]Redback(config-neighbor)#exit
[local]Redback(config-peergroup)#exit
[local]Redback(config-bgp)#exit
```

Static route configuration:

```
[local]Redback(config-ctx)#ip route 10.200.1.2/32 104.1.1.1
[local]Redback(config-ctx)#ip route 10.200.1.4/32 105.1.1.2
[local]Redback(config-ctx)#ip route 44.1.1.0/24 105.1.1.2
```

MSDP configuration—configure MSDP peering between two RPs:

```
[local]Redback(config-ctx)#router msdp
[local]Redback(config-msdp)#peer 10.200.1.2 local-tcp-source lo1
[local]Redback(config-msdp-peer)#no shutdown
```

The configuration for **RouterD** is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface lo1 loopback
[local]Redback(config-if)#ip address 10.200.1.4/32
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E1
[local]Redback(config-if)#ip address 105.1.1.2/24
[local]Redback(config-if)#pim sparse-mode
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#interface E2
[local]Redback(config-if)#ip address 44.1.1.1/24
[local]Redback(config-if)#pim sparse-mode
```

Static RP for Domain Y configuration:

```
[local]Redback(config-if)#ip pim rp-address 10.200.1.3
[local]Redback(config-if)#exit
```

Static route configuration:

```
[local]Redback(config-ctx)#ip route 10.200.1.3/32 105.1.1.1
```

3.4 Multicast VPNs

Multicast-enabled VPNs use MDTs to support IP multicast over BGP/MPLS VPNs. Figure 8 shows the multicast VPN network topology used for the configuration example.

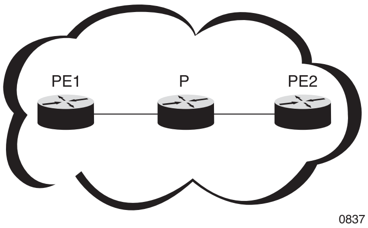

Figure 8 Multicast VPN Network Topology

Multicast-enabled VPNs are configured on both PE routers, PE1 and PE2. In the **local** context, the MDT encapsulation type is configured on loopback interface, **lo1**, which must be the same interface used for BGP peering. The loopback interface is used to source multicast packets on the MDT. An intercontext P2P interface is also configured in the local context, and is used to pass traffic between the VPN and the local context. (This interface does not need an IP address.)

A generic intercontext interface, **ic-local**, is configured in the VPN context, **VPN1**. This interface is similar to a loopback interface in that it is not bound to anything. It creates an intercontext circuit between the **VPN1** context and the **local** context. PIM-SM and the MDT default group are configured on this intercontext interface.

Note: The IP address of the intercontext interface, `ic-local`, must be the same as that of the loopback interface in the local context used for BGP peering.

Because the MDT default group is configured in the **VPN1** context on each PE router, this information must be sent to the other PE router. When each PE router discovers that the other PE router is configured for MDTs, with the same MDT group, it sends a PIM join, with the remote PE router's loopback address as the multicast source, and the MDT group as the multicast group. This forms the MDT tree for forwarding traffic from CE router to the backbone.

The configuration for the **PE1** router is as follows:

```
[local]PE1#configure
[local]PE1(config)#service multiple-contexts
[local]PE1(config)#context local
[local]PE1(config-ctx)#no ip domain-lookup
[local]PE1(config-ctx)#interface ic-vpn1 intercontext p2p 1
[local]PE1(config-if)#pim sparse-mode passive
[local]PE1(config-if)#exit
[local]PE1(config-ctx)#interface lo1 loopback
[local]PE1(config-if)#ip address 10.0.0.3/32
[local]PE1(config-if)#pim sparse-mode passive
[local]PE1(config-if)#mdt encapsulation gre
[local]PE1(config-if)#exit
[local]PE1(config-ctx)#interface to_P
[local]PE1(config-if)#ip address 10.1.1.3/24
[local]PE1(config-if)#pim sparse-mode
[local]PE1(config-if)#exit
[local]PE1(config-ctx)#router rip backbone
[local]PE1(config-rip)#redistribute connected
[local]PE1(config-rip)#interface to_P
[local]PE1(config-rip-if)#exit
[local]PE1(config-rip)#exit
[local]PE1(config-ctx)#router mpls
[local]PE1(config-mpls)#interface to_P
[local]PE1(config-mpls-if)#exit
[local]PE1(config-mpls)#exit
[local]PE1(config-ctx)#router ldp
[local]PE1(config-ldp)#interface lo1
[local]PE1(config-ldp)#interface to_P
[local]PE1(config-ldp)#exit
[local]PE1(config-ctx)#router bgp 100
[local]PE1(config-bgp)#neighbor 10.0.0.2 internal
[local]PE1(config-bgp-neighbor)#update-source lo1
[local]PE1(config-bgp-neighbor)#address-family ipv4 unicast
[local]PE1(config-bgp-af)#exit
[local]PE1(config-bgp-neighbor)#address-family ipv4 vpn
[local]PE1(config-bgp-af)#exit
[local]PE1(config-bgp-neighbor)#exit
[local]PE1(config-bgp)#exit
[local]PE1(config-ctx)#pim rp-address 10.1.1.2
[local]PE1(config-ctx)#exit
[local]PE1(config)#context VPN1 vpn-rd 10.0.0.3:1
[local]PE1(config-ctx)#interface ic-local intercontext p2p 1
[local]PE1(config-if)#ip address 10.0.0.3/24
[local]PE1(config-if)#pim sparse-mode
[local]PE1(config-if)#mdt default-group 239.1.1.1
[local]PE1(config-if)#exit
[local]PE1(config-ctx)#interface to_CE1
[local]PE1(config-if)#ip address 11.1.1.2/24
[local]PE1(config-if)#pim sparse-mode
[local]PE1(config-if)#exit
[local]PE1(config-ctx)#router bgp vpn
[local]PE1(config-bgp)#address-family ipv4 unicast
[local]PE1(config-bgp-af)#export route-target 100:1
[local]PE1(config-bgp-af)#import route-target 100:1
[local]PE1(config-bgp-af)#redistribute connected
[local]PE1(config-bgp-af)#exit
[local]PE1(config-bgp)#exit
[local]PE1(config-ctx)#pim rp-address 11.1.1.2
[local]PE1(config-ctx)#exit
[local]PE1(config)#card ge-10-port 4
[local]PE1(config)#port ethernet 4/8
[local]PE1(config-port)#no shutdown
[local]PE1(config-port)#bind interface to_P local
[local]PE1(config-port)#exit
[local]PE1(config)#port ethernet 4/11
[local]PE1(config-port)#no shutdown
[local]PE1(config-port)#bind interface to_CE1 VPN1
[local]PE1(config)#end
```

The configuration for the **P** router is as follows:

```
[local]P#configure
[local]P(config)#context local
[local]P(config-ctx)#interface to_PE1
[local]P(config-if)#ip address 10.1.1.2/24
[local]P(config-if)#pim sparse-mode
[local]P(config-if)#exit
[local]P(config-ctx)#interface to_PE2
[local]P(config-if)#ip address 20.1.1.2/24
[local]P(config-if)#pim sparse-mode
[local]P(config-if)#exit
[local]P(config-ctx)#router rip backbone
[local]P(config-rip)#redistribute connected
[local]P(config-rip)#interface to_PE1
[local]P(config-rip-if)#exit
[local]P(config-rip)#interface to_PE2
[local]P(config-rip-if)#exit
[local]P(config-rip)#exit
[local]P(config-ctx)#router mpls
[local]P(config-mpls)#interface to_PE1
[local]P(config-mpls-if)#exit
[local]P(config-mpls)#interface to_PE2
[local]P(config-mpls-if)#exit
[local]P(config-mpls)#exit
[local]P(config-ctx)#router ldp
[local]P(config-ldp)#interface to_PE1
[local]P(config-ldp)#interface to_PE2
[local]P(config-ldp)#exit
[local]P(config-ctx)#pim rp-address 10.1.1.2
[local]P(config-ctx)#exit
[local]P(config)#card ge-10-port 13
[local]P(config)#port ethernet 13/6
[local]P(config-port)#no shutdown
[local]P(config-port)#bind interface to_PE1 local
[local]P(config-port)#exit
[local]P(config)#port ethernet 13/11
[local]P(config-port)#no shutdown
[local]P(config-port)#bind interface to_PE2 local
[local]P(config)#end
```

The configuration for the **PE2** router is as follows:

```
[local]PE2#configure
[local]PE2(config)#service multiple-contexts
[local]PE2(config)#context local
[local]PE2(config-ctx)#interface ic-vpn1 intercontext p2p 1
[local]PE2(config-if)#pim sparse-mode passive
[local]PE2(config-if)#exit
[local]PE2(config-ctx)#interface lo1 loopback
[local]PE2(config-if)#ip address 10.0.0.2/32
[local]PE2(config-if)#pim sparse-mode passive
[local]PE2(config-if)#mdt encapsulation gre
[local]PE2(config-if)#exit
[local]PE2(config-ctx)#interface to_P
[local]PE2(config-if)#ip address 20.1.1.3/24
[local]PE2(config-if)#pim sparse-mode
[local]PE2(config-if)#exit
[local]PE2(config-ctx)#router rip backbone
[local]PE2(config-rip)#redistribute connected
[local]PE2(config-rip)#interface to_P
[local]PE2(config-rip-if)#exit
[local]PE2(config-rip)#exit
[local]PE2(config-ctx)#router mpls
[local]PE2(config-mpls)#interface to_P
[local]PE2(config-mpls-if)#exit
[local]PE2(config-mpls)#exit
[local]PE2(config-ctx)#router ldp
[local]PE2(config-ldp)#interface lo1
[local]PE2(config-ldp)#interface to_P
[local]PE2(config-ldp)#exit
[local]PE2(config-ctx)#router bgp 100
[local]PE2(config-bgp)#neighbor 10.0.0.3 internal
[local]PE2(config-bgp-neighbor)#update-source lo1
[local]PE2(config--bgp-neighbor)#address-family ipv4 unicast
[local]PE2(config--bgp-af)#exit
[local]PE2(config-bgp-neighbor)#address-family ipv4 vpn
[local]PE2(config-bgp-af)#exit
[local]PE2(config-bgp-neighbor)#exit
[local]PE2(config-bgp)#exit
[local]PE2(config-ctx)#pim rp-address 10.1.1.2
[local]PE2(config-ctx)#exit
[local]PE2(config)#context VPN1 vpn-rd 10.0.0.2:1
[local]PE2(config-ctx)#no ip domain-lookup
[local]PE2(config-ctx)#interface ic-local intercontext p2p 1
[local]PE2(config-if)#ip address 10.0.0.2/24
[local]PE2(config-if)#pim sparse-mode
[local]PE2(config-if)#mdt default-group 239.1.1.1
[local]PE2(config-if)#exit
[local]PE2(config-ctx)#interface to_CE2
[local]PE2(config-if)#ip address 21.1.1.2/24
[local]PE2(config-if)#pim sparse-mode
[local]PE2(config-if)#no logging console
[local]PE2(config-if)#exit
[local]PE2(config-ctx)#router bgp vpn
[local]PE2(config-bgp)#address-family ipv4 unicast
[local]PE2(config-bgp-af)#export route-target 100:1
[local]PE2(config-bgp-af)#import route-target 100:1
[local]PE2(config-bgp-af)#redistribute connected
[local]PE2(config-bgp-af)#exit
[local]PE2(config-bgp)#exit
[local]PE2(config-ctx)#pim rp-address 11.1.1.2
[local]PE2(config-ctx)#exit
[local]PE2(config)#card ge-10-port 1
[local]PE2(config)#port ethernet 1/3
[local]PE2(config-port)#no shutdown
[local]PE2(config-port)#bind interface to_CE2 VPN1
[local]PE2(config-port)#exit
[local]PE2(config)#port ethernet 1/12
[local]PE2(config-port)#no shutdown
[local]PE2(config-port)#bind interface to_P local
[local]PE2(config-port)#end
```

3.5 Remote Multicast Replication

RMR is used to enable IP multicast services. Figure 9 shows the RMR network topology used for the configuration example.

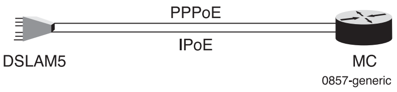

Figure 9 RMR Network Topology

The MC, a SmartEdge router, is connected to an MR, DSLAM5, with PPPoE and IPoE circuits. The PPPoE circuit is created on a 4-port gigabit Ethernet card on slot 14, and an IPoE circuit is created on the **ipoe_to_dslam5** interface, which is bound to a 12-port ethernet card on slot 4. The interface is enabled to forward multicast traffic, and to send and receive the IGMP control messages. The **foo** IGMP service profile is linked to the multicast-enabled **ipoe_to_dslam5** interface.

Subscribers are brought up from the PPPoE circuit. The multicast traffic and the IGMP control messages are forwarded on the IPoE circuit. DSLAM5 replicates the multicast stream for all interested subscribers.

The RMR configuration is as follows:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface ipoe_to_dslam5
[local]Redback(config-if)#ip address 11.1.1.1/24
[local]Redback(config-if)#igmp service-profile foo
[local]Redback(config-if)#multicast output accept-unknown-mac
[local]Redback(config-if)#pim sparse-mode passive
[local]Redback(config-if)#exit
[local]Redback(config)#interface pppoe_to_dslam5 multibind
[local]Redback(config-if)#ip address 192.1.1.1/16
[local]Redback(config-if)#ip pool 192.1.0.0/16
[local]Redback(config-if)#pim sparse-mode passive
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#igmp service-profile foo
[local]Redback(config-igmp-service-profile)#instant-leave
[local]Redback(config-igmp-service-profile)#static-group 224.1.1.1
[local]Redback(config-igmp-service-profile)#exit
[local]Redback(config-ctx)#igmp service-profile bar
[local]Redback(config-igmp-service-profile)#multicast destination ipoe_to_dslam5 local
[local]Redback(config-igmp-service-profile)#exit
[local]Redback(config-ctx)#subscriber name joe
[local]Redback(config-sub)#password test
[local]Redback(config-sub)#ip address pool
[local]Redback(config-sub)#ip igmp service-profile test
[local]Redback(config-sub)#exit
[local]Redback(config-ctx)#pim rp-address 21.1.1.1
[local]Redback(config-ctx)#pim static group 224.1.1.1 source 50.1.1.100 send-join
[local]Redback(config-ctx)#ip access-list 1
[local]Redback(config-access-list)#seq 10 deny ip host 224.1.1.1
[local]Redback(config-access-list)#exit
[local]Redback(config-ctx)#exit
[local]Redback(config)#card ge-10-port 4
[local]Redback(config)#port ethernet 4/2
[local]Redback(config-port)#no shutdown
[local]Redback(config-port)#bind interface ipoe_to_dslam5 local
[local]Redback(config-port)#exit
[local]Redback(config)#card gigaether-4-port 14
[local]Redback(config)#port ethernet 14/1
[local]Redback(config-port)#no shutdown
[local]Redback(config-port)#encapsulation pppoe
[local]Redback(config-port)#bind authentication chap pap context local maximum 8000
[local]Redback(config-port)#end
```

3.6 QoS Adjustments for RMR

The following example shows QoS adjustment configured for static and dynamic bindings.

```
[local]Redback#configure
[local]Redback(config)#context contextA
[local]Redback(config-ctx)#ip access-list HD-channels
[local]Redback(config-access-list)#seq 10 permit igmp 225.0.2.0 0.0.0.255
[local]Redback(config-access-list)#exit
[local]Redback(config-ctx)#ip access-list SD-channels
[local]Redback(config-access-list)#seq 10 permit igmp host 224.0.10.1
[local]Redback(config-access-list)#seq 20 permit igmp 225.0.1.0 0.0.0.255
[local]Redback(config-access-list)#exit
[local]Redback(config-ctx)#igmp group-bandwidth 2000 group-list HD-channels
qos-adjust average-packet-size\ 1200 no-oif
[local]Redback(config-ctx)# igmp group-bandwidth 1000 group-list SD-channels
[local]Redback(config-ctx)#igmp service-profile profile1
[local]Redback(config-igmp-service-profile)#multicast adjust-qos-rate delay-interval 8
[local]Redback(config-igmp-service-profile)#multicast adjust-qos-rate metering minimum-rate 1200
[local]Redback(config-igmp-service-profile)#multicast adjust-qos-rate queuing minimum-rate 1000
[local]Redback(config-igmp-service-profile)#exit
[local]Redback(config-ctx)#interface igmpIf1
[local]Redback(config-if)#igmp service-profile profile1
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#ip igmp service-profile profile1
[local]Redback(config-sub)#qos policy metering metering-policy
```

```
[local]Redback(config-sub)#qos policy queuing  pwfq-policy
[local]Redback(config-sub)#qos profile overhead overhead-profile
[local]Redback(config-sub)#exit
[local]Redback(config-cxt)#exit
[local]Redback(config)#qos policy pwfq-policy pwfq
[local]Redback(config-policy-pwfq)#rate maximum 5000
[local]Redback(config-policy-pwfq)#num-queues 1
[local]Redback(config-policy-pwfq)#queue 0 priority 0 weight 100
[local]Redback(config-policy-pwfq)#exit
[local]Redback(config)#qos policy pwfq-policy pwfq
[local]Redback(config-policy-pwfq)#rate maximum 5000
[local]Redback(config-policy-pwfq)#num-queues 1
[local]Redback(config-policy-pwfq)#queue 0 priority 0 weight 100
[local]Redback(config-policy-pwfq)#exit
[local]Redback(config)#qos policy metering-policy metering
[local]Redback(config-policy-metering)#rate 7000
[local]Redback(config-policy-metering)#burst 100000
[local]Redback(config-policy-metering)#excess-burst 100000
[local]Redback(config-policy-metering)#exit
[local]Redback(config)#qos profile overhead-profile overhead
[local]Redback(config-profile-overhead)#rate-factor 10
[local]Redback(config-profile-overhead)#encaps-access-line ipoa-llc
[local]Redback(config-profile-overhead)#reserved 4
[local]Redback(config-profile-overhead)#type adsl1
[local]Redback(config-profile-overhead)#encaps-access-line pppoa-llc
[local]Redback(config-profile-overhead)#reserved 8
[local]Redback(config-profile-overhead)#exit
[local]Redback(config)#link-group access-lg access
[local]Redback(config-link-group)#encapsulation dot1q
[local]Redback(config-link-group)# qos pwfq scheduling physical-port
[local]Redback(config-link-group)#dot1q pvc 100 encapsulation 1qtunnel
[local]Redback(config-link-group)#dot1q pvc on-demand 100:100 encapsulation pppoe
[local]Redback(config-link-group)#bind authentication chap context contextA
[local]Redback(config-link-group)#exit
[local]Redback(config)#port ethernet 4/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 100
[local]Redback(config-dot1q-pvc)#bind interface igmpIf1 contextA
[local]Redback(config-dot1q-pvc)#exit
[local]Redback(config-port)#dot1q pvc 101
[local]Redback(config-dot1q-pvc)#bind authentication chap contextA
[local]Redback(config-dot1q-pvc)#exit
[local]Redback(config)#port ethernet 4/2
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-portlink-group access-lg
[local]Redback(config-port)#exit
[local]Redback(config)#port ethernet 4/3
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 100
[local]Redback(config-dot1q-pvc)#qos policy metering metering-policy
[local]Redback(config-dot1q-pvc)#qos policy queuing pwfq-policy
[local]Redback(config-dot1q-pvc)#exit
[local]Redback(config-portexit
[local]Redback(configend
```

3.7 Anycast RP

Anycast RP is a mechanism that provides RP redundancy and load-sharing capabilities by allowing the use of multiple RPs within a single multicast domain. Assuming that the sources are evenly spaced around the network, an equal number of sources register with each RP. That is, the process of registering the sources are shared equally by all the RPs in the network.

All routers acting as RPs must be configured with a loopback interface using the same anycast RP address. All downstream routers use that anycast RP address as the IP address for their local RP. To facilitate communication between RPs, each router acting as an RP must also be configured with its own unique IP address, which is used only to send and receive messages from the other RPs.

Figure 10 shows the Anycast RP network topology used for the configuration example.

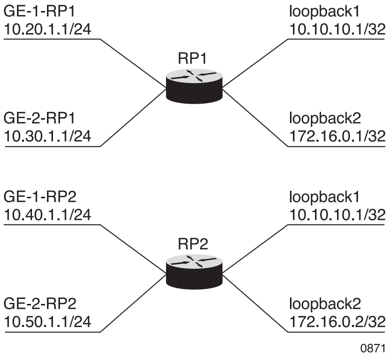

Figure 10 Anycast RP Network Topology

In this configuration example, two routers, **RP1** and **RP2**, are configured for anycast RP. Both routers are configured with a loopback interface, **loopback1**, using the same IP address, which is used as the IP address for the anycast RP set. Both routers are also configured with a loopback interface, **loopback2**, using unique IP addresses. The **loopback2** interface is used to facilitate communication between the two RPs. The other interfaces, **GE-1-RP1**, **GE-2-RP1**, **GE-1-RP2**, and **GE-2-RP2**, are physical interfaces that connect to the network, and are used to send and receive multicast packets.

The configuration for the **RP1** router is as follows:

```
[local]RP1#configure
[local]RP1(config)#context local
[local]RP1(config-ctx)#interface loopback1 loopback
[local]RP1(config-if)#description Anycast-RP-Looback
[local]RP1(config-if)#ip address 10.10.10.1/32
[local]RP1(config-if)#pim sparse-mode
[local]RP1(config-if)#exit
[local]RP1(config-ctx)#interface loopback2 loopback
[local]RP1(config-if)#description Unique-RP-Loopback
[local]RP1(config-if)#ip address 172.16.0.1/32
[local]RP1(config-if)#pim sparse-mode
[local]RP1(config-if)#exit
[local]RP1(config-ctx)#interface GE-1-RP1
[local]RP1(config-if)#ip address 10.20.1.1/24
[local]RP1(config-if)#pim sparse-mode
[local]RP1(config-if)#exit
[local]RP1(config-ctx)#interface GE-2-RP2
[local]RP1(config-if)#ip address 10.30.1.1/24
[local]RP1(config-if)#pim sparse-mode
[local]RP1(config-if)#exit
[local]RP1(config-ctx)#router ospf
[local]RP1(config-ospf)#area 0.0.0.0
[local]RP1(config-ospf-area)#interface GE-2-RP1
[local]RP1(config-ospf-if)#exit
[local]RP1(config-ospf-area)#interface GE-1-RP2
[local]RP1(config-ospf-if)#exit
[local]RP1(config-ospf-area)#interface loopback1
[local]RP1(config-ospf-if)#exit
[local]RP1(config-ospf-area)#interface loopback2
[local]RP1(config-ospf-if)#exit
[local]RP1(config-ospf-area)#exit
[local]RP1(config-ospf)#exit
[local]RP1(config-ctx)#pim anycast-rp 10.10.10.1 172.16.0.1
[local]RP1(config-ctx)#pim anycast-rp 10.10.10.1 172.16.0.2
[local]RP1(config-ctx)#pim rp-address 10.10.10.1
```

The configuration for the **RP2** router is as follows:

```
[local]RP2#configure
[local]RP2(config)#context local
[local]RP2(config-ctx)#interface loopback1 loopback
[local]RP2(config-if)#description Anycast-RP-Looback
[local]RP2(config-if)#ip address 10.10.10.1/32
[local]RP2(config-if)#pim sparse-mode
[local]RP2(config-if)#exit
[local]RP2(config-ctx)#interface loopback2 loopback
[local]RP2(config-if)#description Unique-RP-Loopback
[local]RP2(config-if)#ip address 172.16.0.2/32
[local]RP2(config-if)#pim sparse-mode
[local]RP2(config-if)#exit
[local]RP2(config-ctx)#interface GE-1-RP2
[local]RP2(config-if)#ip address 10.40.1.1/24
[local]RP2(config-if)#pim sparse-mode
[local]RP2(config-if)#exit
[local]RP2(config-ctx)#interface GE-2-RP2
[local]RP2(config-if)#ip address 10.50.1.1/24
[local]RP2(config-if)#pim sparse-mode
[local]RP2(config-if)#exit
[local]RP2(config-ctx)#router ospf
[local]RP2(config-ospf)#area 0.0.0.0
[local]RP2(config-ospf-area)#interface GE-2-RP2
[local]RP2(config-ospf-if)#exit
[local]RP2(config-ospf-area)#interface GE-1-RP2
[local]RP2(config-ospf-if)#exit
[local]RP2(config-ospf-area)#interface loopback1
[local]RP2(config-ospf-if)#exit
[local]RP2(config-ospf-area)#interface loopback2
[local]RP2(config-ospf-if)#exit
[local]RP2(config-ospf-area)#exit
[local]RP2(config-ospf)#exit
[local]RP2(config-ctx)#pim anycast-rp 10.10.10.1 172.16.0.1
[local]RP2(config-ctx)#pim anycast-rp 10.10.10.1 172.16.0.2
[local]RP2(config-ctx)#pim rp-address 10.10.10.1
```

3.8 IGMP Snooping

This section provides examples of configuring IGMP snooping.

3.8.1 IGMP Snooping Configuration—Ethernet Bridge Circuit

This section provides an example of how to configure IGMP snooping on an Ethernet bridge circuit. First, create and configure and IGMP snooping profile. In the following example, the user creates and configures an IGMP snooping profile called sanjose1:

```
[local]Redback#configure
[local]Redback(config)#igmp snooping profile sanjose1
[local]Redback(config-igmp-snooping-profile)#static 233.1.1.1
[local]Redback(config-igmp-snooping-profile)#static 233.1.1.2
[local]Redback(config-igmp-snooping-profile)#static 233.1.1.3
[local]Redback(config-igmp-snooping-profile)#static 233.1.1.4
[local]Redback(config-igmp-snooping-profile)#commit
Transaction committed.
```

Next, apply the IGMP snooping profile to an existing bridge profile. In the following example, the user applies the IGMP snooping profile called sanjose1 to the bridge profile called bridge-prof1:

```
[local]Redback(config)#bridge profile bridge-prof1
[local]Redback(config-bridge-profile)#igmp snooping profile sanjose1
[local]Redback(config-bridge-profile)#commit
```

Then, enable IGMP snooping on a particular Ethernet bridge, as desired. In the following example, the user enables IGMP snooping on a bridge called br-sj-1, in the context sj1:

```
[local]Redback(config)#context sj1
[local]Redback(config-ctx)#bridge br-sj-1
[local]Redback(config-bridge)#igmp snooping
[local]Redback(config-igmp-snooping)#commit
```

Finally, add the bridge profile to an existing bridge:

```
[local]Redback(config)#port ethernet
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 105
[local]Redback(config-dot1q-pvc)#bridge profile bridge-prof1
[local]Redback(config-dot1q-pvc)#commit
```

3.8.2 IGMP Snooping Configuration—VPLS Instance

This section provides an example of how to configure IGMP snooping on a VPLS instance. First, create and configure an IGMP snooping profile. In the following example, the user creates an IGMP and configures an IGMP snooping profile called red1:

```
[local]Redback#configure
[local]Redback(config)#igmp snooping profile red1
[local]Redback(config-igmp-snooping-profile)#access-group acl1
[local]Redback(config-igmp-snooping-profile)#max-groups 100000
[local]Redback(config-igmp-snooping-profile)#mrouter static
[local]Redback(config-igmp-snooping-profile)#receive deny
[local]Redback(config-igmp-snooping-profile)#commit
Transaction committed.
```

Next, apply the IGMP snooping profile to an existing bridge profile. In the following example, the user applies the IGMP snooping profile called `sanjose1 to the bridge profile called red-bridge1:`

```
[local]Redback(config)#bridge profile red-bridge1
[local]Redback(config-bridge-profile)#igmp snooping profile sanjose1
[local]Redback(config-bridge-profile)#commit
```

Finally, add the bridge profile to a VPLS profile (which is attached to a VPLS neighbor). In the following example, the user applies the red-bridge1 bridge profile to a VPLS profile called `red-profile1:`

```
[local]Redback(config)#vpls profile red-profile1
[local]Redback(config-vpls-profile)#neighbor 122.23.52.121
[local]Redback(config-vpls-profile-neighbor)#bridge profile red-bridge1
[local]Redback(config-bridge-profile)#commit
```