Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000

Release 4.1
NN46250-504
Issue 07.01
October 2015

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya may generally make available to users of its products and Hosted Services. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA CHANNEL PARTNER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. “Software” means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. “Designated Processor” means a single stand-alone computing device. “Server” means a Designated Processor that hosts a software application to be accessed by multiple users. “Instance” means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine (“VM”) or similar deployment.

Licence types

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

“Heritage Nortel Software” means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at http://support.avaya.com/LicenseInfo under the link “Heritage Nortel Products” or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants You a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Avaya’s website at: http://support.avaya.com/Copyright or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

Preventing Toll Fraud

“Toll Fraud” is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can

result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Contents

Chapter 1: Introduction

Purpose

This document describes conceptual and procedural information to administer and configure IP Multicast Routing protocols on Avaya Virtual Services Platform 9000. Operations include the following:

- Internet Group Management Protocol (IGMP)
- Protocol Independent Multicast—Source Specific Mode (PIM-SSM)
- Protocol Independent Multicast— Sparse Mode (PIM-SM)
- Multicast MAC Filtering
- Multicast Virtualization

Configure IP multicast routing to transmit data from a source to multiple recipients at the same time. This one-to-many delivery mechanism is similar to broadcasting; however, multicasting transmits data to specific groups, and broadcasting transmits to all devices on the network. Because multicasting transmits only one stream of data to many destinations, multicasting conserves bandwidth

You must configure at least one IP interface on the Avaya Virtual Services Platform 9000. For more information about how to configure interfaces, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

For information about how to configure IP Multicast over Fabric Connect, see *Configuring Avaya Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Related resources

Documentation

See *Documentation Reference for Avaya Virtual Services Platform 9000,* NN46250-100 for a list of the documentation for this product.

Training

Ongoing product training is available. For more information or to register, you can access the website at http://avaya-learning.com/.

Course code	Course title
4D00010E	Knowledge Access: ACIS - Avaya ERS 8000 and VSP 9000 Implementation
5D00040E	Knowledge Access: ACSS - Avaya VSP 9000 Support

Viewing Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

Procedure

- To find videos on the Avaya Support website, go to http://support.avaya.com and perform one of the following actions:
  - In **Search**, type `Avaya Mentor Videos` to see a list of the available videos.
  - In **Search**, type the product name. On the Search Results page, select **Video** in the **Content Type** column on the left.
- To find the Avaya Mentor videos on YouTube, go to www.youtube.com/AvayaMentor and perform one of the following actions:
  - Enter a key word or key words in the **Search Channel** to search for a specific product or topic.
  - Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the website.

Note:

Videos are not available for all products.

Support

Go to the Avaya Support website at http://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes,

downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Searching a documentation collection

On the Avaya Support website, you can download the documentation library for a specific product and software release to perform searches across an entire document collection. For example, you can perform a single, simultaneous search across the collection to quickly find all occurrences of a particular feature. Use this procedure to perform an index search of your documentation collection.

Before you begin

- Download the documentation collection zip file to your local computer.
- You must have Adobe Acrobat or Adobe Reader installed on your computer.

Procedure

1. Extract the document collection zip file into a folder.
2. Navigate to the folder that contains the extracted files and open the file named <*product_name_release*>.pdx.
3. In the Search dialog box, select the option **In the index named <*product_name_release*>.pdx**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
   - Whole Words Only
   - Case-Sensitive
   - Include Bookmarks
   - Include Comments
6. Click **Search**.

   The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance Ranking.

Chapter 2: New in this release

The following sections describe what is new in *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504, for Release 4.1.

Features

See the following sections for information about feature-related changes.

Premier License

Multicast virtualization requires a Premier License. Release 4.1 introduces the Product Licensing and Delivery System (PLDS) as the license order, delivery, and management tool. The switch supports Release 4.1 features in either the Base License or Premier License. For more information on multicast virtualization, see Multicast virtualization on page 33.

For more information on PLDS and licensing, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Other changes

There are no other changes.

Chapter 3: IP multicast fundamentals

IP multicast extends the benefits of Layer 2 multicasting on LANs to WANs. Use multicasting techniques on LANs to help clients and servers find each other. With IP multicast, a source can send information to multiple destinations in a WAN with a single transmission. IP multicast results in efficiency at the source and saves a significant amount of bandwidth.

Overview of IP multicast

IP multicast transmits messages to multiple recipients at the same time. This one-to-many delivery mechanism is similar to broadcasting, except that multicasting transmits to specific groups and broadcasting transmits to all receivers on a network. Because IP multicast transmits only one stream of data to the network where it replicates to many receivers, multicasting saves a considerable amount of bandwidth.

IP multicast services benefit applications such as video conferencing, dissemination of datagram information, and dissemination of mail or news to a large number of recipients.

Multicast protocols use different techniques to discover delivery paths.

A distribution tree is a set of multicast routers and subnetworks that permit the members of a group to receive traffic from a source. The source of the tree depends on the algorithm used by the multicast protocol. The following diagram is an example of a simple distribution tree where S is the multicast source and the arrows indicate the multicast broadcast procedure.

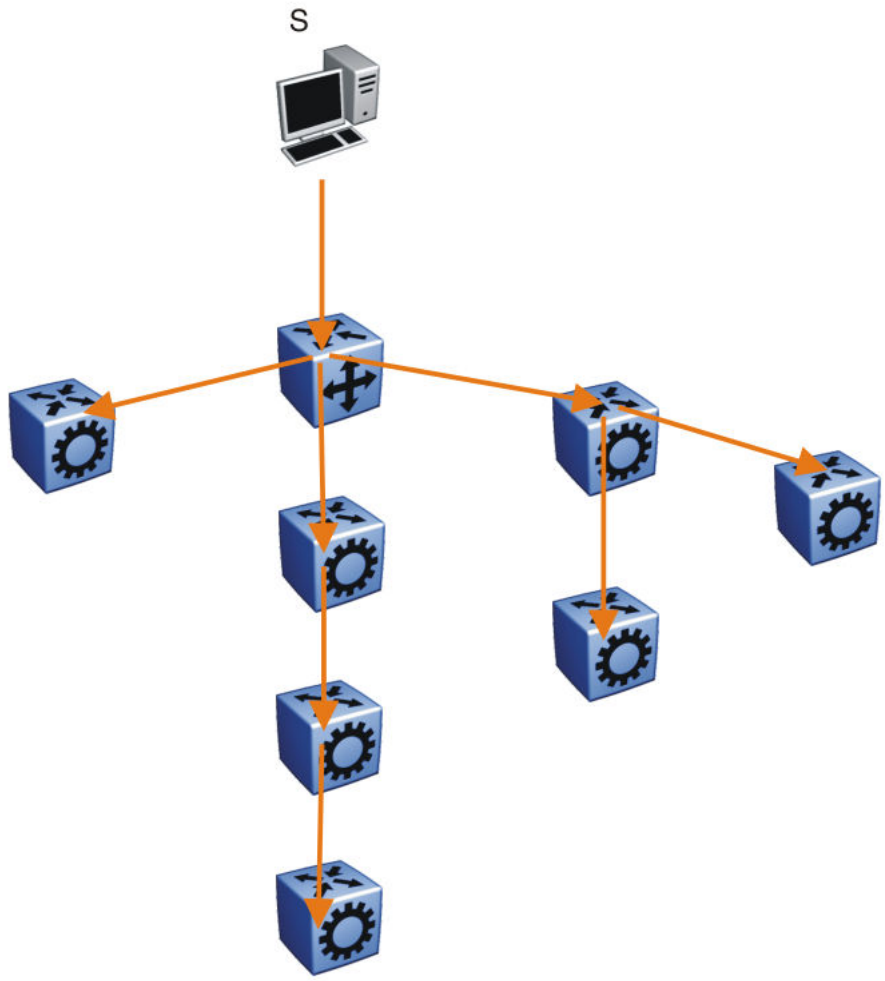

Figure 1: Multicast distribution tree and broadcasting

Broadcast and prune methods use multicast traffic to build the distribution tree. Periodically, the source sends or broadcasts data to the extremities of the internetwork to search for active group members. If no local members of the group exist, the router sends a message to the host, removing itself from the distribution tree, and thus pruning the router.

The following diagram illustrates how the host prunes routers from the distribution tree. First, the router sends a message to the source, after which the pruned routers do not receive multicast data.

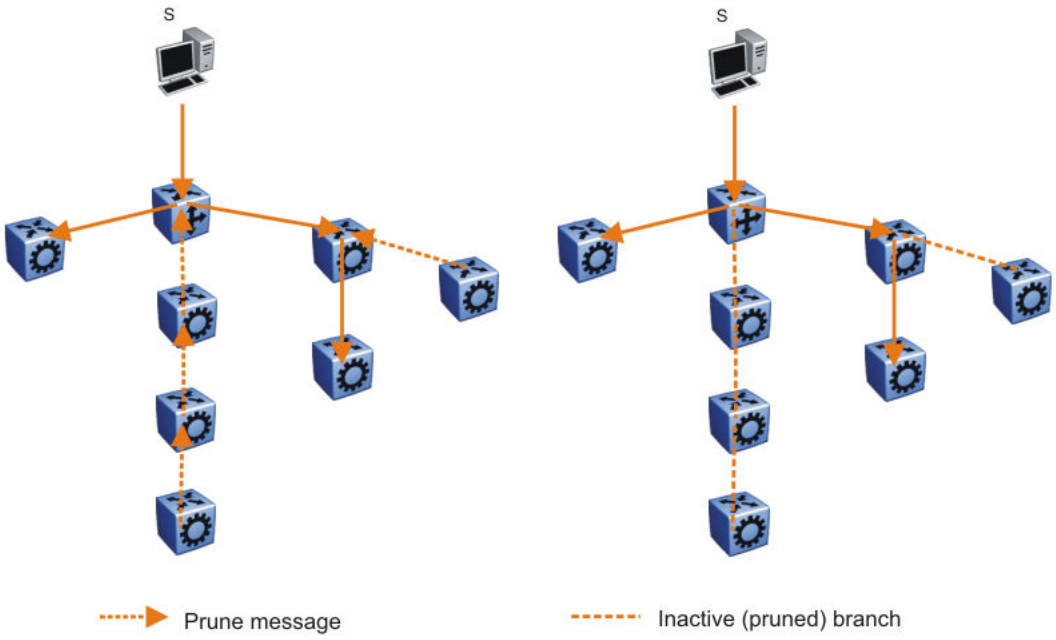

Figure 2: Pruning routers from a distribution tree

Reverse path multicast is based on the concept that a multicast distribution tree is built on the shortest path from the source to each subnetwork that contains active receivers. After a datagram arrives on an interface, the router determines the reverse path to the source of the datagram by examining the routing table of known network sources. If the datagram is not on the optimal delivery tree, the router discards it.

Multicast host groups and their group members enable the IP multicast router to transmit just to those groups interested in receiving the traffic. The Avaya Virtual Services Platform 9000 uses the Internet Group Membership Protocol (IGMP) to learn the existence of host group members on their directly attached subnets. A router communicates with the hosts on a local network by sending IGMP queries. Hosts respond by issuing IGMP reports. For more information about host groups, see Multicast host groups on page 15 and Multicast addresses on page 16. For more information about IGMP, see Internet Group Management Protocol on page 17.

Multicast traffic forwarding transmits frames to all interfaces or subnets for which it receives IGMP reports for the multicast group indicated in the destination IP address. Multicast packets forwarded within the same virtual LAN (VLAN) remain unchanged. The switch does not forward packets to networks that do not use members of the multicast group indicated in the destination IP address.

Multicast host groups

IP multicast is a method for addressing, routing, and delivering a datagram to a collection of receivers called a host group.

Host groups are permanent or transient, with the following characteristics:

- A permanent host group uses a well-known, administratively assigned IP multicast group address. This address is permanent and defines the group. A permanent host group can consist of zero or more members.
- A transient host group exists only as long as members need its services. IP addresses in the multicast range that are not reserved for permanent groups are available for dynamic assignment to transient host groups.

A host system on an IP network sends a message to a multicast group by using the IP multicast address for the group. To receive a message addressed to a multicast group, however, the host must be a member of the group and must reside on a network where that group is registered with a local multicast router.

An IP multicast host group can consist of zero or more members and places no restrictions on its membership. Host members can reside anywhere, they can join and leave the group at any time, and they can be members of more than one group at the same time.

In general, hosts that are members of the same group reside on different networks. However, a range of multicast addresses (224.0.0.x) is reserved for locally-scoped groups. All message traffic for these hosts typically remains on the local network. Hosts that belong to a group in this address range and that reside in different networks do not receive message traffic for each other.

Important:

With the Avaya Virtual Services Platform 9000, you can apply a special set of filters (global filters) to multicast packets. You can create, deny, or accept filters to configure the sources that can receive and send data. For more information about how to configure filters, see *Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000,* NN46250-502.

Multicast addresses

Each host group uses a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP version 4 multicast address is a Class D address (the high-order bits are 1110) from 224.0.0.0 to 239.255.255.255. These addresses are assigned statically for use by permanent groups and dynamically for use by transient groups.

The block of addresses from 224.0.0.1 to 224.0.0.255 is reserved for routing protocols and other low-level protocols. Multicast routers do not forward datagrams with addresses in this range because the time-to-live (TTL) value for the packet is usually 1.

Multicast protocols

You can use the following protocols to enable multicast routing on a Avaya Virtual Services Platform 9000:

- Internet Group Management Protocol (IGMP)—learns the existence of host group members on directly attached subnets.
- Multicast Router Discovery (MRDISC) protocol—discovers multicast routers in a Layer 2 bridged domain configured for IGMP snoop.
- Protocol Independent Multicast (PIM)
  - Sparse Mode (PIM-SM) protocol—suitable for implementation on networks sparsely populated by receivers.
  - Source Specific Multicast (PIM-SSM) protocol—uses a one-to-many model where members can receive traffic from one or more specific sources. This protocol is suitable for television channels and other content-distribution applications.

Static source groups

Use static source groups (or static mroutes) to configure static source-group entries in the PIM-SM, or PIM-SSM multicast routing table. PIM cannot prune these entries from the distribution tree. In other words, even if no receivers for the group exist, the multicast stream for a static source-group

entry stays active. PIM never prunes static forwarding entries. If you no longer need the entries, you must manually delete them.

To configure static source groups, you must first globally enable PIM. If you disable PIM, the switch saves all of the configured static source-group entries and deactivates them. After you re-enable PIM, the switch reactivates the static source groups.

Static source groups ensure that the multicast route (mroute) records remain in the distribution tree. After receivers join the group, they do not experience a delay in receiving multicast data because they do not need to graft onto the group, or start a join process in the case of PIM. This timing is essential for applications where the multicast data must send to a receiver as soon as the receiver joins the group, for example, when a switch delivers television channels to receivers. After the receiver turns the channel, which is equivalent to joining a group, the receiver can view the channel immediately.

Static entries result in continuous traffic if the source is active, even if no receivers exist. However, the system does not forward traffic with a static entry if no receivers exist, but forwards it continuously to the switch where the entry is programmed and crosses intermediate switches on the path.

You can configure static source-group entries for a specific source or subnet. If several sources on the same subnet send traffic to the same group, traffic for all these sources flows continuously when using the subnet configuration.

After you configure static source groups, keep the following points in mind:

- If you disable PIM, the switch deactivates all of the static source groups. After you re-enable PIM, the switch activates the static source groups.
- In PIM-SM configuration, the static source-group feature works for both specific source addresses and subnet addresses by using the SrcSubnetMask field.

  When the network mask is 255.255.255.255, the full source address is used to match the (S,G) which is the specific source case. When the network mask field is a subnet mask for the source, only the source subnet is used to match (S,G)s.
- In PIM-SSM configurations, static source groups have the following limitations:

  - Subnets: SSM static source groups work only with specific IP addresses. Static source groups cannot work with source subnets, so the mask must use a full 32-bit mask, 255.255.255.255, and the source must use a host address.

Internet Group Management Protocol

A host uses IGMP to register group memberships with the local querier router to receive datagrams sent to this router targeted to a group with a specific IP multicast address.

A router uses IGMP to learn the existence of group members on networks to which it directly attaches. The router periodically sends a general query message to each of its local networks. A host that is a member of a multicasting group identifies itself by sending a response.

IGMP queries

When multiple IGMP routers operate on a network, one router is elected to send queries. This elected querier periodically sends host membership queries (also known as general queries) to the attached local subnets. The Avaya Virtual Services Platform 9000 supports queries from all three versions of IGMP.

IGMP host reports

A host that receives a membership query from a local router can respond with a host membership report, one for each multicast group that joins. A host that receives a query delays its reply by a random interval and listens for a reply from other hosts in the same host group. For example, consider a network that includes two host members—host A and host B—of the same multicast group. The router sends out a host membership query on the local network. Both host A and host B receive the query and listen on the network for a host membership report. The delay timer for host B expires first, so it responds to the query with a membership report. Hearing the response, host A does not send a report of its own for the same group.

Each query from a router to a host includes a maximum response time field. IGMP inserts a value n into this field specifying the maximum time in tenths of a second within which the host must issue a reply. The host uses this value to calculate a random value between 0 and n tenths of a second for the period that it waits before sending a response. This calculation is true for IGMP versions 2 and 3. For IGMP version 1, this field is 0 but defaults to a value of 100, that is, 10 seconds.

If at least one host on the local network specifies that it is a member of a group, the router forwards to that network all datagrams that bear the multicast address for the group.

Upon initialization, the host can immediately issue a report for each of its supported multicast groups. The router accepts and processes these asynchronous reports the same as requested reports.

After hosts and routers are in a steady state, they communicate in a way that minimizes the exchange of queries and reports. The designated routers establish a path between the IP multicast stream source and the end stations and periodically query the end stations about whether to continue participation. As long as a client continues to participate, all clients, including nonparticipating end stations on the switch port, receive the IP multicast stream.

Host leave messages

If an IGMPv2 host leaves a group and it is the host that issues the most recent report, it also issues a leave group message. The multicast router on the network issues a group-specific query to determine whether other group members exist on the network. If no host responds to the query, the router assumes that no members belonging to that group exist on that interface.

Fast leave feature

The Avaya Virtual Services Platform 9000 supports a fast leave feature that is useful for multicast-based television distribution applications. Fast leave relies on an alternative leave process where the switch stops sending traffic for the group immediately after it receives a leave message, without issuing a query to check if other group members exist on the network. Fast leave alleviates the network from additional bandwidth demand after a customer changes television channels.

The Avaya Virtual Services Platform 9000 provides several fast leave processes for IP multicast:

- immediate leave with one user for each interface
- immediate leave with several users for each interface

- standard IGMP leave based on a Last Member Query Interval (LMQI), which you can configure in tenths of seconds

Fast leave modifies the IGMP leave processing mechanism on an IGMP interface. After the system receives an IGMP leave message on a fast leave enabled interface, the switch does not send a group-specific query and immediately stops sending traffic to the leaving member (IGMP host) port. Without fast leave, traffic continues to forward until the group times out. This situation wastes bandwidth if no receiver that requires the group traffic exists.

Fast leave mode provides two options of the fast leave mechanism—single-user mode and multiple-users mode:

- Single-user mode: In this mode, the port stops receiving traffic immediately after a group member on that port sends a leave message. Avaya recommends that you use the single-user mode if each interface port connects to only one IGMP host.
- Multiple-users mode: Use this mode if the interface port connects to multiple IGMP hosts. In this case, the port stops receiving traffic after all members leave the IGMP group. The switch removes the leaving IGMP member and, if more group members exist on that port, the switch continues sending traffic to the port.

  When operating in multiple-users mode, the Avaya Virtual Services Platform 9000 must use the correct membership information. To support multiple-users mode, multicast receivers on the same interface cannot use IGMP report suppression. If you must use IGMP report suppression, Avaya recommends that you do not use this mode. Instead, use the LMQI (configurable in units of 1/10ths of seconds) to provide a faster leave process while still sending group-specific queries after the interface receives a leave message.

Fast leave mode applies to all fast-leave enabled IGMP interfaces.

IGMP snoop

The Avaya Virtual Services Platform 9000 provides IP multicast capability when used as a switch. Functioning as a switch, it supports all three versions of IGMP to prune group membership for each port within a VLAN. This feature is IGMP snoop.

Important:

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

Use the IGMP snoop feature to optimize the multicast data flow, for a group within a VLAN, to only those ports that are members of the group. The switch builds a database of group members by listening to IGMP reports from each port. The switch suppresses the reports heard by not forwarding them to ports other than the one receiving the report, thus forcing the members to continuously send their own reports. The switch relays group membership from the hosts to the multicast routers and forwards queries from multicast routers to all port members of the VLAN. Furthermore, the switch forwards multicast data only to the participating group members and to the multicast routers within the VLAN.

The multicast routing functionality can coexist with IGMP snoop on the same switch, but you can configure only one of IGMP snoop or an IP multicast routing protocol, excluding IGMP, on the same VLAN.

Multicast group trace for IGMP snoop

Use this feature to monitor the multicast group trace for an IGMP snoop-enabled Avaya Virtual Services Platform 9000. You can view the multicast group trace from ACLI.

Multicast group trace tracks the data flow path of the multicast streams. Group trace tracks information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port.

IGMP proxy

If a Avaya Virtual Services Platform 9000 receives multiple reports for the same multicast group, it does not transmit each report to the multicast upstream router. Instead, the switch consolidates the reports into a single report and forwards the one report. If you add another multicast group or the system receives a query since it last transmitted the report upstream, the system forwards the report onto the multicast router ports. This feature is IGMP proxy.

IGMP versions

The Avaya Virtual Services Platform 9000 supports IGMPv1, IGMPv2, and IGMPv3. IGMPv1 and IGMPv2 are backward compatible and can exist together on a multicast network. The following list describes the purpose for each version:

- IGMPv1 provides the support for IP multicast routing. IGMPv1 specifies the mechanism to communicate IP multicast group membership requests from a host to its locally attached routers. For more information, see RFC1112.
- IGMPv2 extends the features in IGMPv1 by quickly reporting group membership termination to the routing protocol. This feature is important for multicast groups with highly volatile group membership. For more information, see RFC2236.
- IGMPv3 supports the PIM Source Specific Multicast (SSM) protocol, PIM-SM, and snooping. A host can selectively request or filter traffic from individual sources within a multicast group or from specific source addresses sent to a particular multicast group. Multicast routing protocols use this information to avoid delivering multicast packets from specific sources to networks where there are no interested receivers. For more information, see RFC3376.

For the Virtual Services Platform 9000 implementation of PIM-SSM, each group can use multiple sources.

The following list identifies groups records that a report message includes: -

- current-state record
- source-list-change record
- filter-mode-change record

A current-state record is sent by a system in response to a query received on an interface. It reports the current reception state of that interface, with respect to a single multicast address.

The Record Type of a current-state record has one of the following two values:

- MODE_IS_INCLUDE — Indicates that the interface has a filter mode of include for the specified multicast address. The source address fields in this group record contain the source list of the interface for the specified multicast address.
- MODE_IS_EXCLUDE — Indicates that the interface has a filter mode of exclude for the specified multicast address. The source address fields in this group record contain the source list of the interface for the specified multicast address.

Source-List Change Record — The system sends a source-list-change record after a change of source list occurs that does not coincide with a filter-mode change on the interface for a particular multicast address. The interface on which the change occurs sends a report that includes the record. The record type of a source-list-change record can be one of the following two values:

- ALLOW_NEW_SOURCES — Indicates that the source address [i] fields in this group record contain a list of the additional sources that the system wishes to hear from, for packets sent to the specified multicast address. If the change was to an include source list, these are the addresses that were added to the list. If the change was to an exclude source list, these are the addresses that were deleted from the list.
- BLOCK_OLD_SOURCES — Indicates that the source address [i] fields in this group record contain a list of the sources that the system no longer wishes to hear from, for packets sent to the specified multicast address. If the change was to an include source list, these are the addresses that were deleted from the list; if the change was to an exclude source list, these are the addresses that were added to the list.

  If a change of source list results in both allowing new sources and blocking old sources, then two group records are sent for the same multicast address, one of type ALLOW_NEW_SOURCES and one of type BLOCK_OLD_SOURCES.

Filter Mode — Virtual Services Platform 9000 implements the filter-mode-change record. The system sends a filter-mode-change record whenever the filter mode changes (during a change from include to exclude, or from exclude to include) for a particular multicast address. The interface on which the change occurs sends a report that includes the record. The record type of a filter-mode-change record can be one of the following two values:

- CHANGE_TO_INCLUDE_MODE — Indicates that the interface has changed to include filter mode for the specified multicast address. The source address [i] fields in this group record contain the new source list of the interface for the specified multicast address.
- CHANGE_TO_EXCLUDE_MODE — Indicates that the interface has changed to exclude filter mode for the specified multicast address. The source address [i] fields in this group record contain the new source list of the interface for the specified multicast address.

After you enable IGMPv3, the following actions occur:

- After you change the version on an interface to or from IGMPv3, the switch experiences a disruption to existing multicast traffic on that interface but traffic does recover. Avaya recommends that you do not make this change when the system passes multicast traffic.

IGMP states

Multicast routers implementing IGMPv3 keep one state for each group for every port in every attached network. This group state consists of a filter-mode, a list of sources, and various timers. For each attached network running IGMP, a multicast router records the desired reception state for that network. This state consists of a set of records of the following form:

- multicast address
- group timer
- filter mode (source records)

Each source record is of the form source address or source timer. If all sources within a given group are desired, an empty source record list is kept with filter-mode set to EXCLUDE. This means hosts on this network want all sources for this group to be forwarded. This is the IGMPv3 equivalent to a IGMPv1 or IGMPv2 group join.

Group timer

A group timer represents the time for the filter-mode to expire and switch to INCLUDE mode and is used only when a group is in EXCLUDE mode.

Group timers are updated according to the types of group records received. If a group timer is expiring when a router filter-mode for the group is EXCLUDE means, there are no listeners on the attached network in EXCLUDE mode. At this point, a router will transition to INCLUDE filter-mode.

Source timer

A source timer is maintained for every source record. Source timers are updated according to:

- the type and filter-mode of the group record received
- whenever the source is present in a received record for that group.

If a source timer expires with a router filter-mode for the group of INCLUDE, the router concludes that traffic from this particular source is no longer desired on the attached network, and deletes the associated source record.

If a source record has a running timer with a router filter-mode for the group of EXCLUDE, it means that at least one system desires the source. It should therefore be forwarded by a router on the network. If a source timer expires with a router filter-mode for the group of EXCLUDE, the router informs the routing protocol that there is no receiver on the network interested in traffic from this source. The records are deleted when the group timer expires in the EXCLUDE router filter-mode.

Processing IGMP messages for groups in SSM range

IGMP messages are processed for groups in SSM range in the following scenarios:

1. IGMPv3 interface enabled; PIM-sparse or snooping enabled
   - IGMPv3 reports that contain group records with groups within SSM range are processed with no restrictions.
   - IGMPv2 reports for groups within SSM range translate to IGMPv3 reports with one group record and type IS_EXCLUDE{NULL}. These reports are processed with no restriction as an IGMPv3 report.
   - IGMPv2 leave for groups within SSM range translate to IGMPv3 reports with one group record and type TO_INCLUDE{NULL}. These reports are processed with no restriction as an IGMPv3 report.
2. IGMPv3 interface enabled; PIM-SSM or ssm-snooping enabled
   - IGMPv3 reports that contain group records with groups within SSM range received from members in the EXCLUDE mode are discarded (eg. IS_EXCLUDE and TO_EXCLUDE messages).
   - IGMPv2 reports for groups within SSM range translate to IGMPv3 reports with one group record and type ALLOW{S1,S2,...}. The source list is obtained from the global ssm-map. If there are no sources in the global ssm-map, the message is discarded. These reports are processed with no restriction as an IGMPv3 report.
   - IGMPv2 leave for groups within SSM range translate to IGMPv3 reports with one group record and type BLOCK{S1,S2,...}. The source list is obtained from the global ssm-map. If there are no sources in the global ssm-map, the message is discarded. These reports are processed with no restriction as an IGMPv3 report.

Note:

In order to accept v2 messages, you must enable the compatibility mode on the IGMPv3 interface.

IGMPv3 source-specific forwarding rules

After a multicast router receives a datagram from a source destined to a particular group, the router must decide to forward the datagram to the attached network. The multicast routing protocol uses IGMPv3 information to forward datagrams to all required sources or groups on a subnetwork.

The following table describes the forwarding suggestions that IGMPv3 makes to the routing protocol. The table also identifies the action taken after the source timer expires, based on the filter mode of the group.

Group filter-mode	Source-timer value	Action
INCLUDE	TIMER > 0	Forward the traffic from the source.
INCLUDE	TIMER = 0	Stop forwarding the traffic from the source, and remove the source record. If no more source records exist for the group, delete the group record.
INCLUDE	No source elements	Do not forward the source.
EXCLUDE	TIMER > 0	Forward the traffic from the source.
EXCLUDE	TIMER = 0	Do not forward the traffic from the source. If no more source records exist for the group, delete the group record.
EXCLUDE	No source elements	Forward the traffic from the source.

IGMPv3 explicit host tracking

IGMPv3 explicit host tracking enables the IGMP to track all the source and group members. To track all the source and group members, the sources that are in the include mode hold a list of members who want to receive traffic from that source.

The members that are in the exclude mode are on hold on the reporter list under the port data. By default, IGMPv3 explicit host tracking is disabled.

Important:

If explicit host tracking is enabled, you cannot downgrade the IGMPv3 interface to IGMPv1 or IGMPv2.

For more information on configuring explicit host tracking, see Configuring explicit host tracking on page 103.

IGMPv3 fast leave

When a BLOCK message is received for a source, you must check if the member that sent this message is the last reporter for the source. If it is the last reporter, delete the source. Else, delete the member. No group and source specific queries are sent.

When a LEAVE message is received, you must check if the member that sent this message is the last reporter for the group. If it is the last reporter, switch to INCLUDE mode if sources are available (if no sources are available the port is deleted). Else, delete the member. No group and source specific queries or group specific queries are sent.

Important:

To use IGMPv3 fast leave feature, you must first enable the explicit host tracking feature.

Synchronization of IGMPv3 over SMLT

The implementation of IGMPv3 offers support for IGMPv3 over SMLT. The IST peers must be in sync with the IGMPv3 reports received over SMLT links to ensure effective performance. The IST protocol ensures the infrastructure to send such information from one IST peer to the other.

The synchronization of IGMPv3 members and their advertised sources is different from IGMPv1 and IGMPv2. Because of IGMPv3 compatibility mode, you must consider the IGMP member version. If you have version 1 or 2 members, you must synchronize the IGMP information as IGMPv1 or IGMPv2 reports, so the peer can build an accurate database. In particular, if members with version 1 or 2 exist, the group filter mode is exclude and the exclude source list is empty. Also no v1 or v2 member will be present on any source from include list.

Each member sends IGMP reports in the same manner for all IGMP versions. The sending mechanism depends on the SMLT state.

After an IST peer receives an IGMPv3 report over an SMLT link, it must pass the message to its peer. If the SMLT state is up, the IST peer sends the message encapsulated in an IST IGMPv3 message. If the SMLT state is down, the IST peer sends the message as a plain IGMPv3 report.

In both cases the IGMPv3 message is not altered and the receiving IST peer processes it as expected in SMLT conditions (translating the receiving port to SMLT port if applicable).

Note:

If you enable compatibility mode, and the member sends an IGMPv1 or IGMPv2 report, the message is either an IST IGMPv1 or v2 encapsulated Message or a plain IGMPv1 or IGMPv2 report.

After SMLT up or down events occur, the IST peer must synchronize its IGMPv3 database to its peer, taking into account the new state of the SMLT link.

If you enable IGMP explicit host tracking, each include source stores information for each member that advertises that particular source in an include list. This information is synchronized with the IST peer.

If you do not enable explicit host tracking, each source from include list contains only information related to the last member that sent an IGMPv3 report. Only this information is synchronized with the IST peer.

Backward compatibility

IGMPv3 for PIM-SSM is backward compatible with IGMPv2. You can configure the switch to operate in v3-only mode or in v2-v3 compatibility mode. If you configure the switch to use only v3-only mode, it ignores all v2 and v1 messages except the query message.

If you configure the switch to operate in v2-v3 compatibility mode, the switch supports all IGMPv1, v2, and v3 messages. The switch parses the group address of the messages. If the group address is out of SSM range and it is a v3 message, the switch drops the message; if it is a v2 message, PIM-SM or IGMP snoop processes handle the message.

After the switch receives an IGMPv2 leave message, and the group address in it is within SSM range, the switch sends the group-and-source specific query. If the group address is not within the SSM range, the switch sends the group specific query.

According to RFC3376, the multicast router with IGMPv3 can use one of two methods to handle older query messages:

- If an older version of IGMP is present on the router, the querier must use the lowest version of IGMP present on the network.
- If a router that is not explicitly configured to use IGMPv1 or IGMPv2 hears an IGMPv1 query or IGMPv2 general query, it logs a rate-limited warning.

You can configure if the switch dynamically downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning.

In v2-v3 compatibility mode, an IGMPv2 host can only join if you configure a static entry in SSM map and if the interface operates in PIM-SSM mode or IGMP SSM-Snoop mode.

You can use the compatibility mode with Split MultiLink Trunking (SMLT). One core switch sends an SMLT message to the other core switch after it receives an IGMPv3 message. This action synchronizes the IGMP host information.

Implementation of IGMP

You can enable and disable multicast routing on an interface basis. If you disable multicast routing on an interface, the interface does not generate IGMP queries. If the switch or interface is in IGMP router behavior mode, for example, PIM enabled, you cannot configure IGMP snoop. The switch still learns the group membership and snoops multicast receivers on the switch VLAN or ports.

IGMP Layer 2 Querier

In a Layer 2 multicast network, you can enable Layer 2 querier on one of the switches in the VLAN. IGMP Layer 2 querier provides the IGMP querier function so that the switch can provide the recurring queries that maintain IGMP groups when you do not use multicast routing for multicast traffic.

Overview

In a multicast network, if you only need to use Layer 2 switching for the multicast traffic, you do not need multicast routing. However, you must have an IGMP querier on the network for multicast traffic to flow from sources to receivers. A multicast router provides the IGMP querier function. You can also use the IGMP Layer 2 Querier feature to provide a querier on a Layer 2 network without a multicast router.

The Layer 2 querier function originates queries for multicast receivers, and processes the responses accordingly. On the connected Layer 2 VLANs, IGMP snoop continues to provide services as normal. IGMP snoop responds to queries and identifies receivers for the multicast traffic.

You must enable Layer 2 querier and configure an IP address for the querier before it can originate IGMP query messages. If a multicast router exists on the network, the switch automatically disables the Layer 2 querier.

In a Layer 2 multicast network, enable Layer 2 querier on only one of the switches in the VLAN. A Layer 2 multicast domain supports only one Layer 2 querier. No querier election exists.

IGMP Snooping

IGMP Snooping enables Layer 2 switches in the network to examine IGMP control protocol packets exchanged between downstream hosts and upstream routers.

When Layer 2 switches examine the IGMP control protocol packets, they:

- Generate the Layer 2 MAC forwarding tables used for further switching sessions
- Regulate the multicast traffic to prevent it from flooding the Layer 2 segment of the network

IGMP Layer 2 Querier and IGMP interaction

IGMP Layer 2 Querier uses IGMP to learn which groups have members on each of the attached physical networks, and it maintains a list of multicast group memberships for each attached network and a timer for each membership. In this case, multicast group memberships means the presence of at least one member of a multicast group on a given attached network, not a list of all of the members.

IGMP Layer 2 Querier can assume one of two roles for each of the attached networks:

- Querier
- Non-Querier

After you enable IGMP Layer 2 Querier, the system assumes it is a multicast router, so it sends the General Query, Group Specific/Group, and Source Specific Query when Leave/BLOCK messages are received. IGMP queries are required to maintain an IGMP group.

Note:

Group Specific When Leave does not apply to IGMPv1.

IGMP Layer 2 Querier limitations

The following limitations apply to IGMP Layer 2 Querier.

- IGMP Layer 2 Querier is based on IGMP Snoop. If you disable IGMP Snoop, IGMP Layer 2 Querier does not work until you enable IGMP Snoop and IGMP Layer 2 Querier.
- After you enable IGMP Snoop and IGMP Layer 2 Querier on an interface, if the system receives no IGMP query messages it becomes the querier.

Multicast access control

Multicast access control is a set of features that operate with standard existing multicast protocols. You can configure multicast access control for an IP multicast-enabled port or VLAN with an access control policy that consists of several IP multicast groups.

You can use this feature to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams). For example, in a television distribution application, instead of applying a filter to each channel (multicast group), you can apply a multicast access policy to a range of channels (groups), thereby reducing the total number of

filters and providing a more efficient and scalable configuration. Also, if you want to add or remove television channels from a package, you can modify the multicast access policy; you do not need to change filters for individual VLANs or ports. Multicast access policies contain an ID and a name (for example, PremiumChannels), the list of IP multicast addresses, and the subnet mask.

Multicast access control is not a regular filtering configuration. Multicast access control is for multicast streams and relies on handling multicast control and initial data to prevent hosts from sending or receiving specified multicast streams; it does not use filters. Also, multicast access control provides a list of multicast groups in one configuration using the same routing policy prefix list configuration. For information about prefix lists, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505. You can configure multicast access control and change it dynamically to support changes in the configuration without restarting the protocol. You can change the access capabilities of a user or service subscriber without loss of service.

The following paragraph provides an example of a typical application:

The local cable television company offers three packages; each one includes 35 channels (35 multicast groups). The company configures each package in an access control policy. This policy applies to a set of VLANs or ports to prevent users from viewing the channels on those VLANs. Use the same policy to prevent users from sending traffic to those groups (also known as spoofing) by specifying the deny-tx option for that port. After you define the packages, you can use them for access policy configuration. You can easily change the package by changing the group range, without changing all the port configurations.

The multicast access control functionality applies to an IP multicast application where you must control user access. You can use it in financial-type applications and other enterprise applications, such as multicast-based video conferencing.

Six types of multicast access control policies exist:

- deny-tx
- deny-rx
- deny-both
- allow-only-tx
- allow-only rx
- allow-only-both

The tx policies control the sender and ingress interface for a group; the rx policies control the receivers and egress interface for a group.

deny-tx

Use the deny-tx access policy to prevent a matching source from sending multicast traffic to the matching group on the interface where you configure the deny-tx access policy. Configure this policy on the ingress interface to the multicast source. The deny-tx access policy performs the opposite function of the allow-only-tx access policy. Therefore, the deny-tx access policy and the allow-only-tx access policy cannot exist on the same interface at the same time.

For example, in the following figure, a VLAN 1, the ingress VLAN, uses a deny-tx access policy. This policy prevents multicast traffic sent by Sender from forwarding from VLAN 1 to a receiver, consequently preventing Receiver 1 and Receiver 2 from receiving data from the multicast group. You can create receive-only VLANs, such as VLAN 1, with the deny-tx policy.

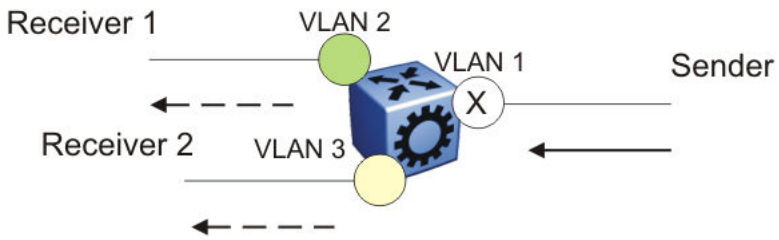

Figure 3: Data flow using deny-tx policy

deny-rx

Use the deny-rx access policy to prevent a matching group from receiving IGMP reports from the matching receiver on the interface where you configure the deny-rx access policy. The deny-rx access policy performs the opposite function of the allow-only-rx access policy. Therefore, the deny-rx access policy and the allow-only-rx access policy cannot exist on the same interface at the same time.

For example, in the following figure, a VLAN 2 uses a deny-rx access policy, preventing IGMP reports sent by Receiver 1 from receiving on VLAN 2. You can deny a multicast group access to a specific VLAN or receiver using the deny-rx policy.

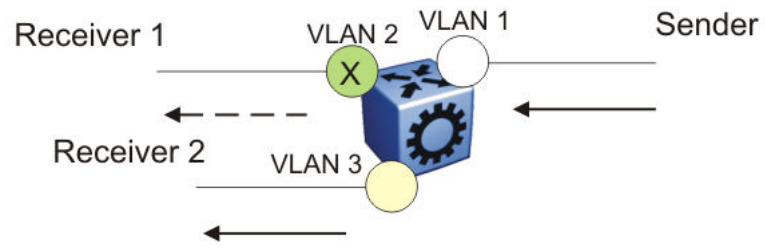

Figure 4: Data flow using deny-rx policy

deny-both

Use the deny-both access policy to prevent a matching IP address from both sending multicast traffic to, and receiving IGMP reports from, a matching receiver on an interface where you configure the deny-both policy. You can use this policy to eliminate all multicast activity for a receiver or source in a specific multicast group. The deny-both access policy performs the opposite function of the allow-only-both access policy. Therefore, the deny-both access policy and the allow-only-both access policy cannot exist on the same interface at the same time.

For example, in the following figure, a VLAN 2 uses a deny-both access policy, preventing VLAN 2 from receiving IGMP reports sent by Receiver 2, and preventing multicast traffic sent by Sender 2 from forwarding from VLAN 2. You can prevent certain VLANs from participating in an activity involving the specified multicast groups with the deny-both policy.

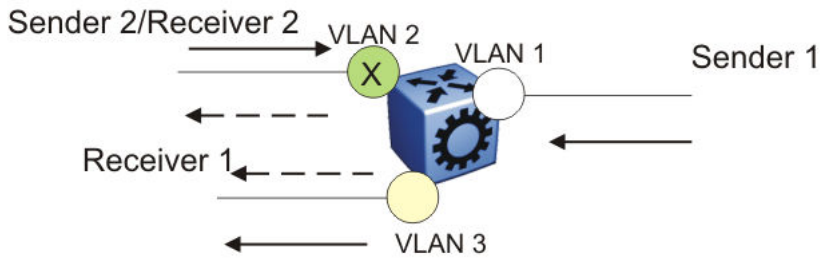

Figure 5: Data flow using deny-both policy

allow-only-tx

Use the allow-only-tx policy to allow only the matching source to send multicast traffic to the matching group on the interface where you configure the allow-only-tx policy. The interface discards all other multicast data it receives. The allow-only-tx access policy performs the opposite function of the deny-tx access policy. Therefore, the allow-only-tx access policy and the deny-tx access policy cannot exist on the same interface at the same time.

allow-only-rx

Use the allow-only-rx policy to allow only the matching group to receive IGMP reports from the matching receiver on the interface where you configure the allow-only-rx access policy. The interface discards all other multicast data it receives. The allow-only-rx access policy performs the opposite function of the deny-rx access policy. Therefore, the allow-only-rx access policy and the deny-rx access policy cannot exist on the same interface at the same time.

allow-only-both

Use the allow-only-both policy to allow only the matching IP address to both send multicast traffic to, and receive IGMP reports from, the matching receiver on the interface where you configure the allow-only-both access policy. The interface discards all other multicast data and IGMP reports. The allow-only-both access policy performs the opposite function of the deny-both access policy. Therefore, the allow-only-both access policy and the deny-both access policy cannot exist on the same interface at the same time.

Host addresses and masks

When you configure multicast access policies, you must specify the host (IP) address and host (subnet) mask of the host to filter (the host that sends multicast traffic).

You can use the host subnet mask to restrict access to a portion of the host network. For example, if you configure the host subnet mask as 255.255.255.255, you use the full host address. To restrict access to a portion of the network of a host, use a subnet mask such as 255.255.255.0. Access control applies to the specified subnet only.

Multicast stream limitation feature

You can configure the multicast stream limitation feature to limit the number of multicast groups that can join a VLAN. By limiting the number of concurrent multicast streams, a service provider can, for example, protect the bandwidth on a specific interface and control access to multicast streams.

Use multicast stream limitation in an environment where you want to limit users to a certain number of multicast streams simultaneously. For example, a television service provider can limit the number of television channels a user can watch at a time. (To a television service provider, a multicast stream is synonymous with a television channel.) If a user purchases a service contract for two single-tuner television receivers, they can use two channels flowing at the same time, but not a third. The service provider can control the bandwidth usage in addition to preventing users from watching more than the allowed number of channels at a point in time.

You can enable the multicast stream limitation feature on the Avaya Virtual Services Platform 9000 by using one of the following methods:

- for each interface—This limitation controls the total number of streams for all clients on this brouter port.
- for each VLAN—This limitation controls the total number of streams for all clients on this VLAN. This method is equivalent to the interface stream limitation.
- for each VLAN port—This limitation controls the number of streams for all clients on this VLAN port. This method is equivalent to the interface port stream limitation.

You can configure the maximum number of streams for each limit independently. After the number of streams meets the limit, the interface drops additional join reports for new streams. The maximum number of streams for each limit is 65535 and the default is 4.

Multicast Router Discovery protocol

The Multicast Router Discovery (MRDISC) protocol can automatically discover multicast-capable routers. By listening to multicast router discovery messages, Layer 2 devices can determine where to send multicast source data and IGMP host membership reports. This feature is useful in a Layer 2 bridging domain that you configure for IGMP snoop.

IGMP multicast router discovery consists of three message types that discover multicast routers on the network:

- Multicast router advertisements: routers advertise that IP multicast forwarding is enabled on an interface.
- Multicast router solicitations: routers solicit a response of multicast router advertisements from all multicast routers on a subnet.
- Multicast router termination messages: a router terminates its multicast routing functions.

Multicast routers send multicast router advertisements periodically on all interfaces where you enable multicast forwarding. Multicast routers also send advertisements in response to multicast router solicitations.

Multicast router solicitations transmit to the IGMP-MRDISC all-routers multicast group that uses a multicast address of 224.0.0.2. Multicast router solicitations do not transmit if a router needs to discover multicast routers on a directly attached subnet.

Multicast router termination messages transmit after a router terminates its multicast routing functions. Other non-IP forwarding devices, such as Layer 2 switches, can send multicast router solicitations to solicit multicast router advertisements.

If you enable IGMP snoop on a Avaya Virtual Services Platform 9000, MRDISC is enabled by default.

Multicast flow distribution over MLT

MultiLink Trunking (MLT) is a mechanism to distribute multicast streams over a multilink trunk and achieve an even distribution of the streams. The distribution is based on source-subnet and group addresses. In applications like television distribution, multicast traffic distribution is particularly important because the bandwidth requirements are substantial when you use a large number of television streams.

Avaya Virtual Services Platform 9000 enables this feature by default and you can not change the configuration.

Distribution algorithm

To determine the port for a particular source-group (S,G) pair, the number of active ports of the multilink trunk is used to MOD the number generated by the XOR of each byte of the masked group address with the masked source address. By default, the group mask and source mask is 255.255.255.255. A byte with a value of 255 in the mask means that the corresponding byte in the group or source address is taken into account when the algorithm is applied.

For example, consider:

If the group address is G[0].G[1].G[2].G[3], the group mask is GM[0].GM[1].GM[2].GM[3], the source subnet address is S[0].S[1].S[2].S[3], and the source mask is SM[0].SM[1].SM[2].SM[3]

Then, the port equals:

(((((G[0] AND GM[0]) xor (S[0] AND SM[0])) xor ((G[1] AND GM[0]) xor (S[1] AND SM[1]))) xor ((G[2] AND GM[2]) xor (S[2] AND SM[2]))) xor ((G[3] AND GM[3]) xor (S[3] AND SM[3]))) MOD (active ports of the MLT)

The algorithm used for traffic distribution causes sequential distribution if the streams are similar to those in the example that follows. Assume that the multilink trunk ports are 3/1 to 3/4, that mask configuration is 0.0.0.0 for the source mask and 0.0.0.255 for the group mask, and that source A.B.C.D sends to the following groups:

X.Y.Z.1

X.Y.Z.2

X.Y.Z.3

.....

X.Y.Z.10

The algorithm chooses link 3/1 for group X.Y.Z.1, 3/2 for group X.Y.Z.2, 3/3 for group X.Y.Z.3, and continues for the remaining ports.

Traffic redistribution

Traffic redistribution distributes the streams on the multilink trunk links if an MLT configuration change occurs. For example, you can add or delete ports.

This feature redistributes active streams according to the distribution algorithm on the multilink trunk links. This redistribution can cause minor traffic interruptions. To minimize the effect of redistribution of multicast traffic on the multilink trunks, the implementation does not move the streams to the appropriate links at the same time. Instead, it redistributes a few streams at every time tick of the system.

To that end, after a multilink trunk port becomes inactive, this feature redistributes all the streams on the multilink trunk ports based on the assignment provided by the distribution algorithm. For more information, see Distribution algorithm on page 31.

By default, redistribution is enabled and you can not change the configuration.

For more information about MLT, see *Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000,* NN46250-503.

Multicast MAC filtering

Some network applications rely on a Layer 2 multicast MAC mechanism to send a frame to multiple hosts for processing. For example, mirroring is one such application. You can direct MAC multicast flooding to a specific set of ports using the multicast MAC filtering feature.

Important:

You can configure multicast MAC filtering only for local addresses on a switch. You cannot use this feature to route traffic between switches (for example, you cannot configure it to forward for interfaces that are not local).

The multicast MAC is a MAC address where the least significant bit of the most significant byte is 1. The multicast MAC filtering feature is available for Layer 2. Because the feature is also effective for IP routed traffic, however, Layer 3 functionality is available as well. (This filtering does not apply to Bridge Protocol Data Units (BPDU).

In Layer 2, a multicast MAC address generally floods to all ports in the VLAN. With multicast MAC filtering, you can define a separate flooding domain for a multicast MAC address, which is a subset of the ports on a VLAN.

In Layer 3, you must configure an Address Resolution Protocol (ARP) entry for routed traffic that maps the unicast IP address to the multicast MAC address and lists the delivery ports for data destined for that IP or multicast MAC address.

To perform multicast MAC filtering, create the VLAN, and then manually define a flooding domain (that is, MAC address and port list) for a specific multicast address. When you specify the multicast MAC flooding domain, you must indicate the ports or multilink trunks to consider for multicast traffic. The flooding is based on whether the specified ports are active members in the VLAN.

Multicast virtualization

Multicast provides simplified extension of internal video and data delivery to remote locations.

Virtualized multicast enables multiple VPN routing instances on devices and supports various unicast routing protocols so that you can provide the services of many virtual routers from one physical device.

You must purchase and install a Premier License to use the VRF Lite feature and, therefore, multicast virtualization.

You can configure multicast routing support with the Virtual Routing and Forwarding (VRF) Lite feature and you can use VRF Lite to emulate many virtual routers with one router.

Multicast virtualization support includes:

- IGMP snooping
- IGMP in Layer 2 virtual services networks (VSN)
- IGMP in Layer 3 VSNs

To implement multicast virtualization, you must perform the following tasks:

1. Create a VRF. For more information about how to create and configure a VRF, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.
2. Create a VLAN and associate it with the VRF.
3. Enable one of the following: IGMP snooping on the VLAN, Layer 2 VSN, or Layer 3 VSN.

If you use IGMP snooping on the VLAN, ensure the IGMP version on the multicast hosts or other network devices is either the same as the version on the VLAN, or enable compatibility mode.

Multicast virtualization does not support PIM. Virtual Services Platform 9000 supports IGMP with PIM only in the Global Router. For more information about IGMP in Layer 2 and Layer 3 VSNs, see *Configuring Avaya Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

High Availability (HA)

IGMP snooping, Layer 2 VSN IP Multicast over Fabric Connect, and Layer 3 VSN IP Multicast over Fabric Connect all support full HA.

In full HA implementation, both the configuration and runtime application data tables exist on the master CPU and the secondary CPU. The master CPU automatically updates the forwarding tables of the secondary CPU in real time.

VRF Lite background

VRF Lite provides independent IPv4 forwarding instances and independent routing instances (contexts), which can reside on the same or different modules, VLANs, and ports.

While forwarding and routing instances are mapped to IP interfaces, incoming traffic is classified into a VLAN and IP interface and, depending on the IP interface, routed context traffic is forwarded.

Scalability and performance

Virtual Services Platform 9000 supports:

- IGMP instances on 64 VRFs

- PIM instance only on the GRT

Protocol Independent Multicast-Sparse Mode

PIM-SM, as defined in RFC2362, supports multicast groups spread out across large areas of a company or the Internet. PIM-SM sends multicast traffic only to routers that specifically join a multicast group. This technique reduces traffic flow over WAN links and overhead costs for processing unwanted multicast packets.

Dense-mode protocols use a flood-and-prune technique, which is efficient with densely-populated receivers. However, for sparsely populated networks, PIM-SM is more efficient because it sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic.

PIM-SM is independent of a specific unicast routing protocol, but it does require the presence of a unicast routing protocol, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). PIM-SM uses the information from the unicast routing table to create and maintain multicast trees that enable PIM-enabled routers to communicate.

Typically, a PIM-SM network consists of several multipoint data streams, each targeted to a small number of LANs in the internetwork. For example, customers whose networks consist of multiple hosts on different LANs can use PIM-SM to simultaneously access a video data stream, such as video conferencing, on a different subnet.

Important:

In some cases, PIM stream initialization can take several seconds.

Hosts

A host is a source, a receiver, or both:

- A source, also known as a sender, sends multicast data to a multicast group.
- A receiver receives multicast data from one or several sources that sends data to a multicast group.

PIM-SM domain

PIM-SM operates in a domain of contiguous routers on which PIM-SM is enabled.

Each PIM-SM domain requires the following routers:

- designated router (DR)
- rendezvous point (RP) router
- bootstrap router (BSR)

Although a PIM-SM domain can use only one active RP router and one active BSR, you can configure additional routers as a candidate RP (C-RP) router and as a candidate BSR (C-BSR). Candidate routers provide backup protection in case the primary RP router or BSR fails.

As a redundancy option, you can configure several RPs for the same group in a PIM domain. As a load sharing option, you can have several RPs in a PIM-SM domain map to different groups. Avaya Virtual Services Platform 9000 devices use the hash function defined in the PIM-SM standard to elect the active RP.

Designated router

The designated router (DR), the router with the highest IP address on a LAN, performs the following tasks:

- sends register messages to the RP router on behalf of directly connected sources
- sends join and prune messages to the RP router on behalf of directly connected receivers
- maintains information about the status of the active RP router for local sources in each multicast group

Important:

The DR is not a required configuration. Switches act automatically as the DR for directly attached sources and receivers.

Rendezvous point router

PIM-SM builds a shared multicast distribution tree within each domain, and the RP router is at the root of this shared tree. Although you can physically locate the RP anywhere on the network, it must be as close to the source as possible. Only one active RP router exists for a multicast group.

At the RP router, receivers meet new sources. Sources use the RP to identify themselves to other routers on the network; receivers use the RP to learn about new sources.

The RP performs the following tasks:

- registers a source that wants to announce itself and send data to group members
- joins a receiver that wants to receive data for the group
- forwards data to group

Candidate rendezvous point router

You can configure a set of routers as C-RP routers that serve as backup to the RP router. If an RP fails, all the routers in the domain apply the same algorithm to elect a new RP from the group of C-RP routers. To make sure that the routers use a complete list of C-RP routers, the C-RP router periodically sends unicast advertisement messages to the BSR. The most common implementation is to configure a PIM-SM router as both a C-RP router and a C-BSR.

Avaya Virtual Services Platform 9000 devices use the hash function defined in the PIM-SM standard to elect the active RP.

Static rendezvous point router

You can configure a static entry for an RP router with static RP. This feature avoids the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism. Static RP-enabled switches cannot learn about RPs through the BSR because the switch loses all dynamically learned BSR information and ignores BSR messages. After you configure static RP entries, the switch adds them to the RP set as if they were learned through the BSR.

Important:

In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interface configured as an RP.

When you configure a PIM static RP in a switch, the next hop of the unicast route toward the PIM static RP must be a PIM neighbor. The PIM protocol fails to work, due to a route change, if the next hop toward an already configured static RP becomes a non-PIM neighbor. If a PIM neighbor cannot reach the configured RP, the RP does not activate and its state remains invalid.

A static RP-enabled Avaya Virtual Services Platform 9000 can communicate with switches from other vendors that do not use the BSR mechanism. Some vendors use either early implementations of PIM-SM v1 that do not support the BSR or proprietary mechanisms like the Cisco Auto-RP. For a network to work properly with static RP, you must map all the switches in the network (including switches from other vendors) to the same RP or RPs, if several RPs exist in the network.

To avoid a single point of failure, you can also configure redundant static RPs.

Use the static RP feature when you do not need dynamic learning mode, typically in small networks, or for security reasons, where RPs are forced to devices in the network so that they do not learn other RPs.

Static RP configuration considerations

Before you can configure a static RP, you must enable PIM-SM and enable static RP.

After you meet these prerequisites, keep in mind the following configuration considerations:

- You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.
- All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.
- Static RPs do not age, that is, they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, it does not know about the static RP unless you configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP.
- In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interfaces configured as an RP.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of Avaya and other vendor switches across the network, you must ensure that all switches and routers use the same active RP because other vendors can use different algorithms to elect the active RP. Avaya Virtual Services Platform 9000 devices use the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.

  Important:

  To reduce convergence times, Avaya recommends that you create only one static RP for each group. The more static RPs you configure for redundancy, the more time PIM requires to rebuild the mroute table and associate RPs.
- Static RP configured on the switch is active as long as the switch uses a unicast route to the static RP network. If the switch loses this route, the static RP is invalidated and the hash

algorithm remaps all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm remaps the affected groups.

Bootstrap router

The BSR receives RP router advertisement messages from the candidate RPs. The BSR adds the RP router with its group prefix to the RP set. Only one BSR exists for each PIM-SM domain.

The BSR periodically sends bootstrap messages containing the complete RP set to all routers in the domain. The BSR ensures that all PIM-SM routers send join, prune, and register packets.

Within a PIM-SM domain, you can configure a small set of routers as C-BSRs. The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

Important:

Configure C-BSRs on routers that are central to all candidate RPs.

Shared trees and shortest-path trees

A PIM-SM domain uses shared trees and shortest-path trees to deliver data packets to group members. This section describes both trees.

Shared trees

Group members in a PIM-SM domain receive the first packet of data from sources across a shared tree. A shared tree consists of a set of paths that connect all members of a multicast group to the RP. PIM creates a shared tree when sources and receivers send messages toward the RP.

Shortest-path trees

After receiving a certain number of packets from the RP, the DR switches from a shared tree to an SPT. Switching to an SPT creates a direct route between the receiver and the source. The Avaya Virtual Services Platform 9000 switches to the SPT after it receives the first packet from the RP.

Figure 6: Shared tree and shortest-path tree on page 38 shows a shared tree and an SPT.

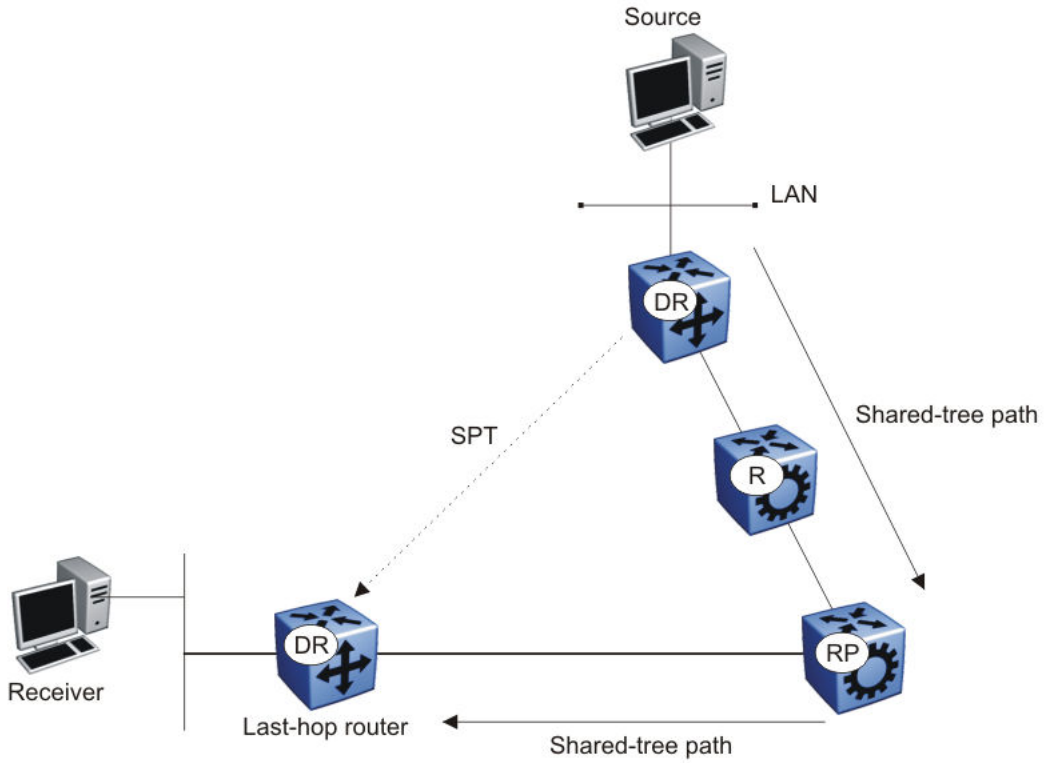

Figure 6: Shared tree and shortest-path tree

Receiver joining a group

The following steps describe how a receiver joins a multicast group:

1. A receiver multicasts an IGMP host membership message to the group that it wants to join.
2. After the last-hop router (the DR), normally the PIM router with the highest IP address for that VLAN, receives the IGMP message for a new group join, the router looks up the associated elected RP with responsibility for the group.
3. After it determines the RP router for the group, the last-hop router creates a (*,G) route entry in the multicast forwarding table and sends a (*,G) join message to the RP. After the last-hop router receives data packets from the RP, if the multicast packet arrival rate exceeds the DR threshold, the last-hop router switches to the SPT by sending an (S,G) join message to the source. (S denotes the source unicast IP address, and G denotes the multicast group address.)
4. If the last-hop router switches to the SPT , the following actions occur:
   - All intermediate PIM routers along the path to the source create the (S,G) entry.
   - To trim the shared tree, the router sends an (S,G) prune message to the RP.

Receiver leaving a group

Before it leaves a multicast group, a receiver sends an IGMP leave message to the DR. If all directly connected members of a multicast group leave or time out, and no downstream members remain, the DR sends a prune message upstream and PIM-SM deletes the route entry after that entry times out.

When the system ages PIM mroutes, it does not clear the (S,G) entry for an inactive route immediately after the expiration period. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.

Source sending packets to a group

The following steps describe how a source sends multicast packets to a group:

1. A source directly attached to a VLAN bridges the multicast data to the DR. The DR for the VLAN (the router with the highest IP address) encapsulates each packet in a register message and sends a unicast message directly to the RP router to distribute to the multicast group.
2. If a downstream group member chooses to receive multicast traffic, the RP router sends a join or prune message toward the source DR and forwards the data down the RP tree after it obtains the data natively.
3. After the receiver DR obtains the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.
4. If no downstream members want to receive multicast traffic, the RP router sends a register-stop message (for the source) to the DR.

   The DR starts the register suppression timer after it receives the first register-stop message. During the register suppression timeout period (the default is 60 seconds), the following events occur:

   - The DR for the source sends a probe packet to the RP router before the register suppression timer expires. The probe packet prompts the RP router to determine whether new downstream receivers joined the group.
   - If no new receivers joined the group, the RP router sends another register-stop message to the DR for the source, and its register suppression timer restarts.
   - After the RP router no longer responds with a register-stop message to the source DR probe message, the register suppression timer expires and the DR sends encapsulated multicast packets to the RP router. The RP router uses this method to tell the DR that new members joined the group.

The RP sends a register-stop message to the DR immediately after it receives the first multicast data packet.

Required elements for PIM-SM operation

For PIM-SM to operate, the following elements must exist in the PIM-SM domain:

- You must enable an underlying unicast routing protocol for the switch to provide routing table information to PIM-SM.
- You must configure an active BSR to send bootstrap messages to all PIM-v2 configured switches and routers to enable them to learn group-to-RP mapping. If you configure several BSRs in a network, an active BSR is elected based on priority and IP address (if priority is equal, the BSR with the higher IP address is elected).
- You must include an RP to perform the following tasks:
  - manage one or several IP multicast groups
  - become the root for the shared tree to these groups
  - accept join messages from receiver switches for groups that it manages
  - elect an active RP based on priority and IP address (if priority is equal, the RP with the higher IP address is elected)

PIM-SM simplified example

Figure 7: PIM-SM simplified example on page 41 shows a simplified example of a PIM-SM configuration.

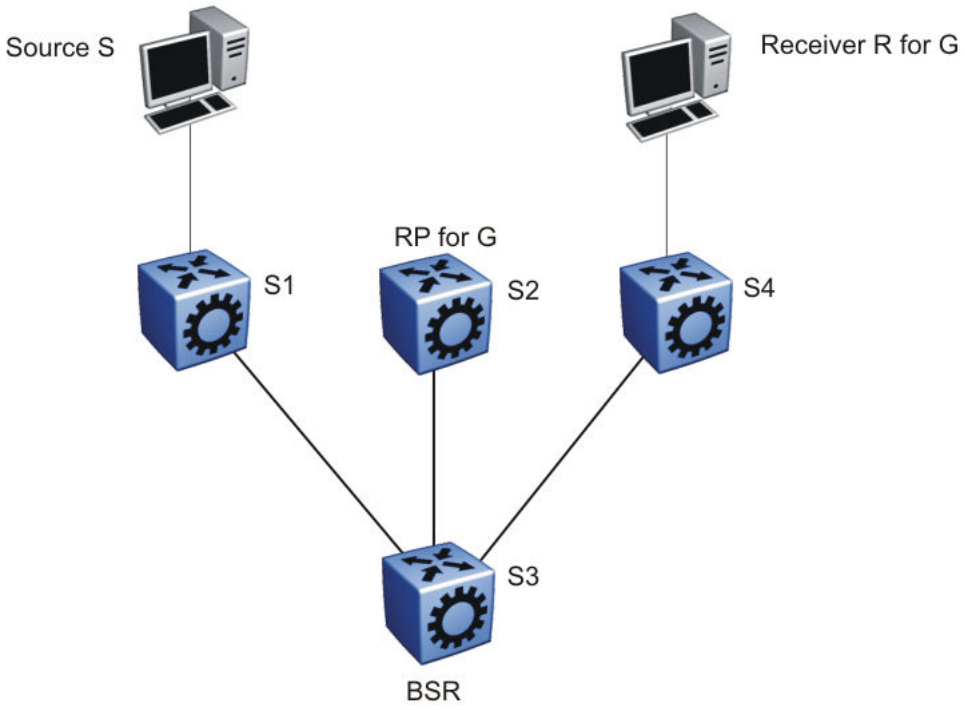

Figure 7: PIM-SM simplified example

In the sample configuration, the following events occur:

1. The BSR distributes RP information to all switches in the network.
2. R sends an IGMP membership report to S4.
3. Acting on this report, S4 sends a (*,G) join message to RP.
4. S sends data to G.
5. The DR (S1 in this example) encapsulates the data that it unicasts to RP (S2) in register messages.
6. S2 decapsulates the data, which it forwards to S4.
7. S4 forwards the data to R.
8. If the packet rate exceeds the DR threshold, S4 sends S1 an (S,G) join message.
9. S1 forwards data to S4. After S4 receives data from S1, it prunes the stream from the RP.

Important:

Figure 7: PIM-SM simplified example on page 41 is a simplified example and is not the best design for a network if you locate the source and receiver as shown. In general, place RPs as close as possible to sources.

PIM-SM static source groups

You can configure static source groups (or static mroutes) as static source-group entries in the PIM-SM multicast routing table. PIM-SM cannot prune these entries from the distribution tree. For more information about static source groups, see Static source groups on page 16.

Join and prune messages

The DR sends join and prune messages from a receiver toward an RP for the group to either join the shared tree or remove (prune) a branch from it. A single message contains both a join and a prune list. This list includes a set of source addresses that indicate the shortest-path trees or the shared trees that the host wants to join. The DR sends join and prune messages hop-by-hop to each PIM router on the path to the source or the RP.

Register and register-stop messages

The DR sends register messages to the RP for a directly connected source. The register message informs the RP of a new source, causing the RP to send join or prune messages back toward the DR of the source, which forwards the data down the RP tree after it obtains the data natively. After the receiver DR obtains the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.

The DR stops sending encapsulated packets to the RP after it receives a register-stop message. This traffic stops without delay because the RP sends a register-stop message immediately after it receives the first multicast data packet, and joins the shortest-path tree.

PIM-SMLT

IP multicast routing support with Split MultiLink Trunking (SMLT) builds a virtual switch that represents the two switches of the split multilink trunk core.

When switches use PIM in the core, they need to exchange protocol-related updates as part of the interswitch trunking (IST) protocol. IST hides the fact that the edge switch attaches to two physical switches.

PIM-SMLT can work in triangular, square, and full mesh configurations with Layer 3 IP multicast.

However, Avaya does not support PIM-SSM in square or full mesh SMLT topologies.

The following rules apply:

- If a VLAN receives traffic from the IST link, it cannot forward on the split multilink trunk link or the edge for the same VLAN.
- If one side of the SMLT link toward the receiver is down, such that the traffic cannot be forwarded directly down the SMLT link from the router on which traffic is ingressing, the IST Peer MUST forward that traffic it receives over the IST link down its side of the SMLT toward the receiver. The decision of whether the IST Peer needs to forward traffic received over the IST to SMLT receivers is made in the datapath, which has full knowledge of the remote SMLT link state.
- Traffic can use the IST to route between VLANs if the forwarding decision for the multicast protocol requires that the other side of the core forwards the multicast traffic (follow the IP multicast routing and forwarding rules for routed traffic). Other VLANs that are not part of SMLT continue to behave in the same way.
- To create a temporary default route pointing to a peer IST, you must enable PIM on the IST VLAN.
- In a scaled multicast environment, if you must reconfigure the members of an MLT link, either SMLT or IST, by removing the ports from the MLT membership list, Avaya recommends that you first shutdown the port by using the **`shutdown`** command at the port configuration level. Let the unicast and multicast traffic subside, and then remove the port from the MLT membership list. If you reconfigure the MLT without first shutting down the port, it can lead to excessive hardware updates to multicast forwarding records and can result in high utilization of the CP.

SMLT provides for fast failover in all cases, but does not provide a functionality similar to Routed SMLT (RSMLT).

Important:

You must enable square SMLT globally before you configure square or full-mesh configurations.

Traffic delay with PIM while restarting peer SMLT switches

If you restart peer SMLT switches, you can lose, or experience a delay in, PIM traffic. The local and remote SMLT links must be up to forward traffic. If a remote SMLT link is down, you can experience a traffic delay.

PIM uses a DR to forward data to receivers on a VLAN. If you restart the DR in an SMLT VLAN, you can lose data because of the following actions:

- If the DR is down, the non-DR switch assumes the role and starts forwarding data.
- After the DR comes back up, it takes priority (higher IP address) to forward data so the non-DR switch stops forwarding data.
- The DR is not ready to forward traffic due to protocol convergence and because it takes time to learn the RP set and create the forwarding path. This situation can result in a traffic delay of 2 to 3 minutes because the DR learns the RP set after Open Shortest Path First (OSPF) converges.

A workaround to this delay is to a configure the static RP router on the peer SMLT switches. This feature avoids the process of selecting an active RP router from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism. After the DR comes back up, traffic resumes as soon as OSPF converges. This workaround reduces the traffic delay to approximately 15 to 65 seconds.

Protocol Independent Multicast-Source Specific Mode

Source Specific Multicast optimizes PIM-SM by simplifying the many-to-many model. Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM only builds source-based SPTs. Whereas PIM-SM always joins a shared tree first, and then switches to the source tree, SSM eliminates the need to start with a shared tree by immediately joining a source through the SPT. SSM avoids using an RP and RP-based shared trees, which can be a potential problem.

Until now only one channel for one group was allowed to exist in ssm map. From now on multiple channels for the members of the SSM group are allowed to be configured in this map.

This configuration is ideal for applications like television channel distribution and other content-distribution businesses. Banking and trade applications can also use SSM as it provides more control over the hosts receiving and sending data over their networks.

When a v2 report in SSM range is received it is translated to an igmpv3 report message with one group record with type ALLOW and the source lists copied from the igmp ssm map static entries and passed to igmpv3 module. When a v2 leave in SSM range is received it is translated to an igmpv3 report message with one group record with type BLOCK and the source lists copied from the igmp ssm map static entries and passed to igmpv3 module. This behaviour is displayed only when PIM-SSM mode is enabled.

SSM applications use IP addresses reserved by the Internet Assigned Numbers Authority (IANA) in the 232/8 range (232.0.0.0 to 232.255.255.255). SSM recognizes packets in this range and controls the behavior of multicast routing devices and hosts that use these addresses. When a source (S) transmits IP datagrams to an SSM destination address (G), a receiver can receive these datagrams by subscribing to the (S,G) channel.

A channel is a source-group (S,G) pair where S is the source that sends to the multicast group and G is an SSM group address. SSM defines channels on an individual or multiple source basis, which enforces the one-to-many concept of SSM applications. In an SSM channel, each group is associated with multiple sources.

SSM features

PIM-SM requires a unicast protocol to forward multicast traffic within the network to perform the Reverse Path Forwarding (RPF) check. PIM-SM uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree that PIM-enabled routers use to communicate. The unicast routing table must contain a route to every multicast source in the network as well as routes to PIM entities like the RPs and BSR.

SSM uses only a subset of the PIM-SM features such as the SPT, DR, and some messages (hello, join, prune, and assert). However, some features are unique to SSM. These features, described in the following sections, are extensions of the IGMP and PIM protocols.

PIM-SSM architecture

The following diagram illustrates how the PIM-SSM architecture requires routers to perform the following actions:

- support IGMPv3 source-specific host membership reports and queries at the edge routers
- initiate PIM-SSM (S,G) joins directly and immediately after receiving an IGMPv3 join report from the designated router
- restrict forwarding to SPTs within the SSM address range by all PIM-SSM routers

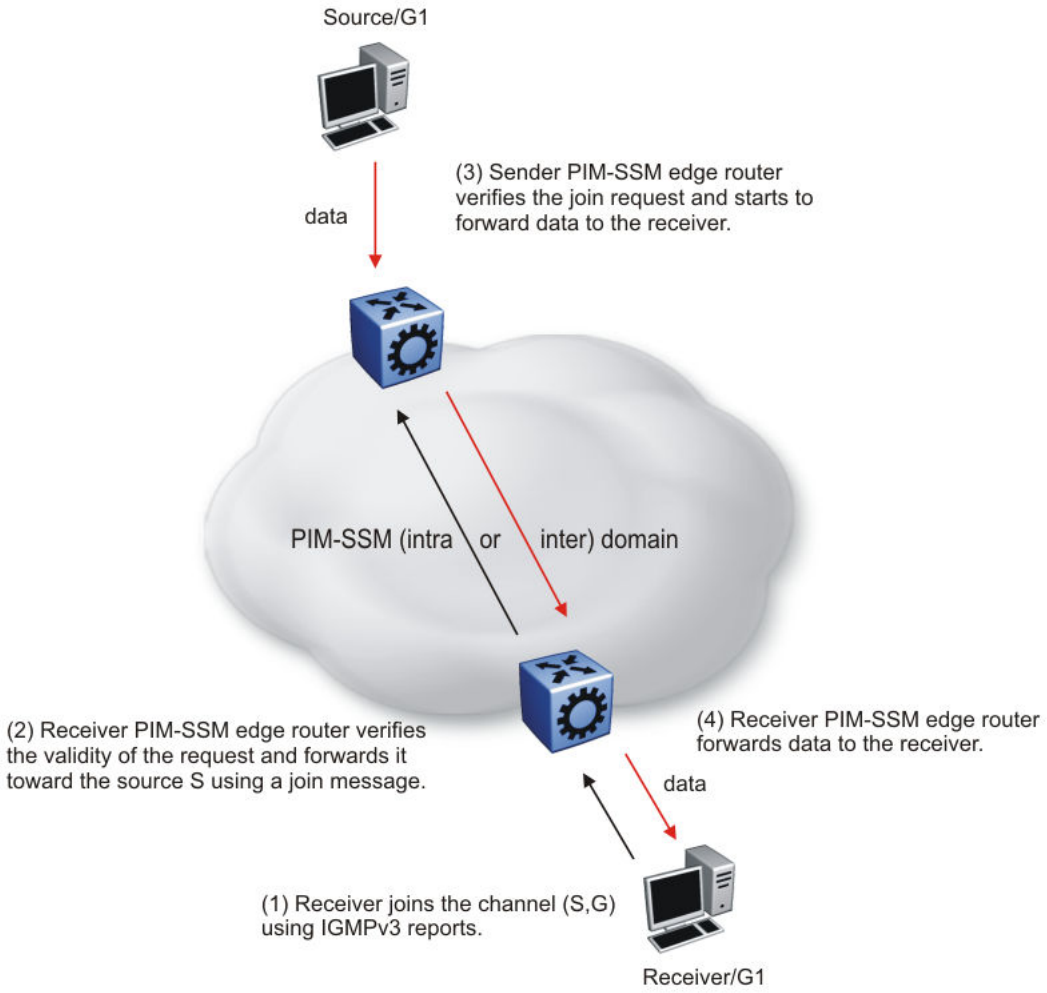

Figure 8: PIM-SSM architecture

The following rules apply to Layer 3 devices with SSM enabled:

- Receive IGMPv3 membership join reports in the SSM range and, if no entry (S,G) exists in the SSM channel table, create one.

- Receive IGMPv2 membership join reports, but only for groups that already use a static (S,G) entry in the SSM channel table.
- Send periodic join messages to maintain a steady SSM tree state.
- Use standard PIM-SM SPT procedures for unicast routing changes, but ignore rules associated with the SPT for the (S,G) route entry.
- Receive prune messages and use standard PIM-SM procedures to remove interfaces from the source tree.
- Forward data packets to interfaces from the downstream neighbors that sent an SSM join, or to interfaces with locally attached SSM group members.
- Drop data packets that do not use an exact-match lookup (S,G) in their forwarding database for S and G.

PIM-SSM static source groups

You can configure static source group entries in the PIM-SSM multicast routing table with static source groups (or static mroutes). PIM-SSM cannot prune these entries from the distribution tree. For more information about static source groups, see Static source groups on page 16.

Implementation of SSM and IGMP

The following sections describe how Avaya Virtual Services Platform 9000 implements PIM-SSM and IGMP.

SSM range

The standard SSM range is 232/8, but you can extend the range to include an IP multicast address. Although you can configure the SSM range, you cannot configure it for all multicast groups (224/4 or 224.0.0.0/240.0.0.0 or 224.0.0.0/255.0.0.0).

You can extend the SSM range to configure existing applications without changing their group configurations.

SSM channel table

You can use the SSM channel to manually configure (S,G) entries that map existing groups to their sending source. These table entries apply to the whole switch, not for each interface, and both IGMPv2 and IGMPv3 hosts use the SSM channel table.

The following rule applies to an SSM channel table for an individual switch:

- You can map one source to multiple groups.
- You can allow multiple sources to the same group.

Important:

Different switches can use different mappings for groups to sources, for example, different channels map differently even if they are on the same network.

SSM and IGMPv2

SSM-configured switches can accept reports from IGMPv2 hosts on IGMPv2 interfaces if the group uses an SSM channel table entry. However, the IGMPv2 host groups must exist in the SSM range defined on the switch, which is 232/8 by default.

- After the SSM switch receives an IGMPv2 report for a group that is in the SSM channel table, it joins the specified source immediately.
- After the SSM switch receives an IGMPv2 report for a group that uses an enabled static SSM channel table entry, it triggers PIM-SSM processing as if it received an equivalent IGMPv3 report.
- After the SSM switch receives an IGMPv2 report for a group out of the SSM range, it processes the report as if it is in PIM-SM mode.

SSM and IGMPv3

The Avaya Virtual Services Platform 9000 supports IGMPv3 for SSM. With IGMPv3, a host can selectively request or filter traffic from sources within the multicast group. IGMPv3 is an interface-level configuration.

Important:

IGMPv3 works without PIM-SSM or SSM-snoop enabled on the interface.

The following rules apply to IGMPv3-enabled interfaces:

- Send only IGMPv3 (source-specific) reports for addresses in the SSM range.
- Accept IGMPv3 reports.
- Drop IGMPv2 reports.

  The IGMPv2 report mentioned in SSM and IGMPv2 on page 47 is processed because it is an IGMPv2 report received on an IGMPv2 interface. If an IGMPv2 interface receives an IGMPv3 report, it drops the report even if PIM-SSM is enabled and the entry is in the SSM channel table. The IGMP versions must match.
- Discard IGMP packets with a group address out of the SSM range.

The Avaya Virtual Services Platform 9000 implements IGMPv3 in one of two modes: dynamic and static.

In dynamic mode, the switch learns about new (S,G) pairs from IGMPv3 reports and adds them to the SSM channel table. If you do not enable dynamic mode and an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM channel table, it ignores the report.

In static mode, you can statically configure (S,G) entries in the SSM channel table. If an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM channel table, it ignores the report. The interface also ignores the report if the group is in the table, but the source or mask does not match what is in the table.

Important:

After you enable IGMPv3, changes to the query interval and robustness values on the querier switch propagate to other switches on the same VLAN through IGMP query.

Both IGMPv2 and IGMPv3 hosts use the SSM channel table:

- An IGMPv2 host (with an IGMPv2 VLAN) must use an existing SSM channel entry if the group is in the SSM range.
- If you enable dynamic learning for an IGMPv3 host, the SSM channel automatically learns the group. Otherwise, the SSM channel also needs a static entry.

The following table summarizes how a switch in PIM-SSM mode works with IGMP if you disable IGMPv3 compatibility. In the following table, references to matching a static SSM channel entry assumes that the entry is enabled. If an entry is disabled, it is treated as though it is disallowed.

Table 1: PIM-SSM interaction with IGMPv2 and v3 with IGMPv3 compatibility disabled

Host	VLAN	SSM range	Action
IGMPv2 host	IGMPv3 VLAN	In or out of range	Drop report.
IGMPv3 host	IGMPv2 VLAN	In or out of range	Drop report.
IGMPv2 host	IGMPv2 VLAN	In range	If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 host	IGMPv2 VLAN	Out of range	Ignore the SSM channel table and process the report as if it is in PIM-SM mode.
IGMPv3 host	IGMPv3 VLAN	Out of range	Process the report.
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic enabled. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and matches an existing SSM channel entry. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and does not match an existing SSM channel entry. Drop report.

The following table summarizes how a switch in PIM-SSM mode works with IGMP if you enable IGMPv3 compatibility.

Table 2: PIM-SSM interaction with IGMPv2 and v3 with IGMPv3 compatibility enabled

Host	VLAN	SSM range	Action
IGMPv2 Host	IGMPv3 VLAN	In range	If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 Host	IGMPv3 VLAN	Out of range	Process the report as in PIM-SM mode.

If an IGMPv3 group report enters the VLAN port and the port must discard one or more of the groups in that packet after the application of IGMP access controls, the port drops the entire packet and does not forward it on to other ports of the VLAN.

If an IGMPv3 interface receives an IGMPv2 or v1 query, the interface backs down to IGMPv2 or v1. As a result, the interface flushes all senders and receivers on the interface.

Configuration limitations

Avaya recommends that you run PIM-SSM on either all switches in the domain or only on the edge routers. If you use a mix of PIM-SSM and PIM-SM switches in the domain, run PIM-SSM on all the edge routers and run PIM-SM on all the core routers.

Important:

A PIM domain with edge routers running PIM-SM and core routers running PIM-SSM does not operate properly. If you prefer or require a mixed PIM-SM and PIM-SSM topology, run PIM-SSM on the edge switches and PIM-SM in the core. Ensure a valid RP configuration exists for groups that exist outside of the SSM range. If a valid RP configuration exists, the SSM switches process the joins in SM mode. If no RP exists, the SSM switches drop the reports.

Static source groups cannot conflict with SSM channels. If you configure a static source group or an SSM channel, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) to different sources or multiple groups to a single source for both static source group and an SSM channel.

PIM passive interfaces

You can configure the PIM interface as active or passive. The default is active. With an active interface, you can configure transmit and receive PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This feature is useful when you use a high number of PIM interfaces and these interfaces connect to end users, not to other switches.

A PIM passive interface does not transmit and drops messages of the following type:

- hello
- join
- prune
- register
- register-stop
- assert
- candidate-RP-advertisement
- bootstrap

If a PIM passive interface receives these types of messages, it drops them and the switch logs a message, detailing the type of protocol message and the IP address of the sending device. These log messages help to identify the device that performs routing on the interface, which is useful if you must disable a device that does not operate correctly.

Important:

A device can send register and register-stop messages to a PIM passive interface, but these messages cannot be sent out of that interface.

The PIM passive interface maintains information about hosts, through IGMP, that are related to senders and receivers, but the interface does not maintain information about PIM neighbors. You can configure a BSR or an RP on a PIM passive interface.

You can also use the PIM passive interface feature as a security measure to prevent routing devices from becoming attached and participating in the multicast routing of the network.

Important:

Before you change the state (active or passive) of a PIM interface, disable PIM on that interface. This action prevents instability in the PIM operations, especially when neighbors exists or the interface receives streams. After you disable PIM, the switch loses traffic for approximately 80 seconds.

Chapter 4: IP multicast basic configuration using ACLI

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts subscribe to multicast services using a host membership protocol. The Internet Group Management Protocol (IGMP) is an example of a host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast–Sparse Mode (PIM–SM) and Protocol Independent Multicast–Source Specific Multicast (PIM–SSM).

Configuring PIM-SM globally

Configure PIM-SM to enable or disable PIM-SM globally on the switch and change default global parameters.

About this task

PIM-SM is the default mode so you do not need to configure the PIM mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable PIM-SM:

   ```
   ip pim enable
   ```

3. Configure the time between bootstrap messages:

   ```
   ip pim bootstrap-period <5-32757>
   ```

4. Configure the timeout to discard data:

   ```
   ip pim disc-data-timeout <5-65535>
   ```

5. Enable the fast join prune interval:

   ```
   ip pim fast-joinprune
   ```

6. Configure the forward cache timeout:

   ```
   ip pim fwd-cache-timeout <10-86400>
   ```

7. Configure the interval for join and prune messages:

   ```
   ip pim join-prune-interval <1-18724>
   ```

8. Specify how long to suppress register messages:

   ```
   ip pim register-suppression-timeout <6-65535>
   ```

9. Specify how often the candidate-rendezvous point (C-RP) sends advertisements:

   ```
   ip pim rp-c-adv-timeout <5-26214>
   ```

10. Configure the polling interval for the routing table manager (RTM):

    ```
    ip pim unicast-route-change-timeout <2-65535>
    ```

11. Verify the configuration changes:

    ```
    show ip pim
    ```

Example

Verify the configuration changes:

```
VSP-9012:1(config)#show ip pim
```

```
================================================================
                    Pim General Group - GlobalRouter
================================================================
PimStat               : disabled
Mode                  : sparse
StaticRP              : disabled
FastJoinPrune         : disabled
BootstrapPeriod       : 60
CRPAdvTimeout         : 60
DiscDataTimeout       : 60
FwdCacheTimeout       : 210
RegSupprTimeout       : 60
UniRouteChangeTimeout : 5
JoinPruneInt          : 60
```

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
bootstrap-period <*5–32757*>	Specifies the interval (in seconds) that the elected bootstrap router (BSR) waits between originating bootstrap messages. The default is 60. To configure this option to the default value, use the default operator with the command.

Table continues…

Variable	Value
disc-data-timeout <*5–65535*>	Specifies how long (in seconds) to discard data until the switch receives the join message from the rendezvous point (RP). An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The default is 60. To configure this option to the default value, use the default operator with the command.
enable	Enables PIM globally on the switch. To configure this option to the default value, use the default operator with the command.
fast-joinprune	Enables the fast join prune interval.
fwd-cache-timeout <*10–86400*>	Specifies the forward cache timeout value. The default is 210. To configure this option to the default value, use the default operator with the command. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.
join-prune-interval <*1–18724*>	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60. To configure this option to the default value, use the default operator with the command.
register-suppression-timeout <6–65535>	Specifies how long (in seconds) the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP. The default is 60. To configure this option to the default value, use the default operator with the command.
rp-c-adv-timeout <5–26214>	Specifies how often (in seconds) a router configured as a C-RP sends C-RP advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected BSR. The default is 60. To configure this option to the default value, use the default operator with the command.
unicast-route-change-timeout <*2–65535*>	Specifies how often (in seconds) the switch polls the RTM for unicast routing information updates for PIM. The default is 5. **Important:** Lowering this value increases how often the switch polls the RTM. This value can affect the performance of the switch, especially when a high volume of traffic flows through the switch. To configure this option to the default value, use the default operator with the command.

Job aid

The following table shows the field descriptions for the **show ip pim** command.

Table 3: show ip pim field descriptions

Field	Description
PimStat	Indicates the status of PIM.
Mode	Indicates the PIM mode.
StaticRP	Indicates the status of static RP.
FastJoinPrune	Indicated the status of the fast join prune interval.
BootstrapPeriod	Indicates the interval between originating bootstrap messages at the elected BSR.
CRPAdvTimeout	Indicates the candidate RP timer (in seconds) for sending C-RP-Adv messages.
DiscDataTimeout	Indicates the time (in seconds) used to discard data until the switch receives the join message from the RP. An IP multicast discard record is created and deleted after the timer expires and after the switch receives a join message.
FwdCacheTimeout	Indicates the PIM forward cache expiry value in seconds. This value ages PIM mroutes.
RegSupprTimeout	Indicates the register-suppression timer in seconds.
UniRouteChangeTimeout	Indicates the frequency at which PIM polls the RTM for routing information updates.
JoinPruneInt	Indicates the join pruning interval in seconds.

Configuring PIM on a VLAN

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you configure PIM on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Create a PIM interface on a VLAN:

```
ip pim enable
```

This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

   ```
   ip pim join-prune-interval <1–18724>
   ```

4. Configure the time between hello messages:

   ```
   ip pim hello-interval <0–18724>
   ```

5. Verify the configuration:

   ```
   show ip pim interface vlan [<1–4084>] [mode]
   ```

Example

Configure the interval for join and prune messages, the time between hello messages, and then verify the configuration.

```
VSP-9012:1(config-if)#ip pim join-prune-interval 60
VSP-9012:1(config-if)#ip pim hello-interval 30
VSP-9012:1(config-if)#show ip pim interface vlan 1
================================================================================
                                  Vlan Ip Pim
================================================================================

VLAN-ID   PIM-ENABLE MODE    HELLOINT  JPINT   CBSRPREF      INTF TYPE
--------------------------------------------------------------------------------
1         disable    sparse  30        60      -1  (disabled) active
```

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
1–4084	Specifies the VLAN ID.
enable	Enables PIM on the local switch interface. To configure this option to the default value, use the default operator with the command.
join-prune-interval <1–18724>	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. To configure this option to the default value, use the default operator with the command.
hello-interval <0–18724>	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default is 30 seconds. To configure this option to the default value, use the default operator with the command.

Job aid

The following table shows the field descriptions for the **`show ip pim interface vlan`** command.

Table 4: show ip pim interface vlan field descriptions

Field	Description
VLAN-ID	Identifies the VLAN.
PIM-ENABLE	Identifies the state of PIM on the VLAN.
MODE	Identifies the configured mode of this VLAN. The valid modes are SSM and Sparse.
HELLOINT	Indicates how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Indicates how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default join or prune interval is 60 seconds.
CBSR PREF	Indicates the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
INTF TYPE	Indicates whether the PIM interface is active or passive.

Configuring PIM on a port

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you configure it on an interface.
- The interface uses a valid IP address.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```

2. Create a PIM interface on a port:

   ```
   ip pim enable
   ```

   This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

```
ip pim join-prune-interval <1–18724>
```

4. Configure the time between hello messages:

```
ip pim hello-interval <0–18724>
```

Example

Configure the interval for join and prune messages and the time between hello messages:

```
VSP-9012(config-if)#ip pim join-prune-interval 60
VSP-9012(config-if)#ip pim hello-interval 30
```

Variable definitions

Use the data in the following table to use the **ip pim** command.

Variable	Value
enable	Enables PIM on the local switch interface. To configure this option to the default value, use the default operator with the command.
join-prune-interval <1–18724>	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. To configure this option to the default value, use the default operator with the command.
hello-interval <0–18724>	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default is 30 seconds. To configure this option to the default value, use the default operator with the command.

Configuring SSM globally

Configure SSM to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Before you begin

- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP and OSPF, see *Configuring OSPF and RIP on Avaya Virtual Services Platform 9000,* NN46250-506.
- Enable PIM globally.

About this task

Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range.

For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure PIM-SSM:

   ```
   ip pim mode ssm
   ```

Configuring IGMP on a VLAN

Configure IGMP for each interface to change default multicasting operations.

Before you begin

- For PIM interfaces, you must enable PIM globally and on the VLAN. For snooping interfaces, do not enable PIM.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Enable IGMP v2-v3 compatibility mode:

   ```
   ip igmp compatibility-mode
   ```

3. Configure the system to downgrade the version of IGMP:

   ```
   ip igmp dynamic-downgrade-version
   ```

4. Configure message intervals and response times:

   ```
   ip igmp last-member-query-interval <0-255> [query-interval <1-65535>] [query-max-response <0-255>]
   ```

5. Configure expected packet loss and IGMP version:

   ```
   ip igmp robust-value <2-255> [version <1-3>]
   ```

6. Add multicast router ports:

   ```
   ip igmp mrouter {slot/port[-slot/port][,...]}
   ```

7. Enable proxy-snoop:

   ```
   ip igmp proxy
   ```

8. Enable router alert:

   ```
   ip igmp router-alert
   ```

9. Enable snooping:

   ```
   ip igmp snooping
   ```

10. Enable SSM-snooping:

    ```
    ip igmp ssm-snoop
    ```

Example

Enter VLAN Interface Configuration Mode for VLAN 1:

```
VSP-9012:1(config)#interface vlan 1
```

Configure the last member query interval to 15 tenths of a second (equal to 1.5 seconds).

```
VSP-9012:1(config-if)#ip igmp last-member-query-interval 15
```

Configure the query interval to 100 seconds.

```
VSP-9012:1(config-if)#ip igmp query-interval 100
```

Configure the query maximum response time to 15 tenths of a second (equal to 1.5 seconds).

```
VSP-9012:1(config-if)#ip igmp query-max-response 50
```

Configure the robustness value to 4 seconds.

```
VSP-9012:1(config-if)#ip igmp robust-value 4
```

Enable proxy snoop for the VLAN.

```
VSP-9012:1(config-if)#ip igmp proxy
```

Enable snoop for the VLAN.

```
VSP-9012:1(config-if)#ip igmp snooping
```

Enable support for SSM on the snoop interface.

```
VSP-9012:1(config-if)#ip igmp ssm-snoop
```

Enable IGMPv3.

```
VSP-9012:1(config-if)#ip igmp version 3
```

Variable definitions

Use the data in the following table to use the **`ip igmp`** command.

Variable	Value
access-list *WORD<1–64>* {A.B.C.D/X} <eny-tx|deny-rx|deny-both|allow-only-tx| allow-only-rx|allow-only-both>	Specifies the name of the access list from 1–64 characters.

Table continues…

Variable	Value
	Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host. Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic

compatibility-mode	Activates v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command:
	`default ip igmp compatibility-mode`
	, or use the no option to disable compatibility mode:
	`no ip igmp compatibility-mode`

dynamic-downgrade-version	Configures if the system downgrades the version of IGMP to handle older query messages. If the system downgrades, the host with IGMPv3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:
	`default ip igmp dynamic-downgrade-version`
	or use the no option to disable downgrade:
	`no ip igmp dynamic-downgrade-version`

igmpv3-explicit-host-tracking	Enables explicit host tracking on IGMPv3. The default state is disabled.
immediate-leave	Enables fast leave on a VLAN.
immediate-leave-members {slot/port[-slot/port] [,...]}	Configures IGMP fast leave members on a VLAN to specify fast-leave-capable ports.
last-member-query-interval <0–255>	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. the default is 10 tenths of a second. Avaya recommends that you configure this value between 3–10 (equal to 0.3 – 1.0 seconds).

Table continues…

Variable	Value
mrdisc [maxadvertinterval <2–180>] [maxinitadvertinterval <2–180>] [maxinitadvertisements <2–15>] [minadvertinterval <3–180>] [neighdeadinterval <2–180>]	Configure the multicast router discovery options to enable the automatic discovery of multicast capable routers. The default parameter values are: • maxadvertinterval: 20 seconds • maxinitadvertinterval: 2 seconds • maxinitadvertisements: 3 • minadvertinterval: 15 seconds • neighdeadinterval: 60 seconds
mrouter {slot/port[-slot/port][,...]}	Adds multicast router ports. {slot/port[-slot/port][,...]} identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).
proxy	Activates the proxy-snoop option globally for the VLAN.
query-interval <1–65535>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.
query-max-response <0–255>	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
robust-value <2–255>	Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.
router-alert	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. **Important:** To maximize network performance, Avaya recommends that you configure this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
snoop-querier	Enables the IGMP Layer 2 Querier feature on the VLAN. The default is disabled.

Table continues…

Variable	Value
snoop-querier-addr {A.B.C.D}	Specifies the IGMP Layer 2 Querier source IP address.
snooping	Activates the snoop option for the VLAN.
ssm-snoop	Activates support for PIM-SSM on the snoop interface.
static-group {A.B.C.D} {A.B.C.D} {port[slot/port[-slot/ port][,...]]} [static|blocked]	Configures IGMP static members to add members to a snoop group. {A.B.C.D} {A.B.C.D} indicates the IP address range of the selected multicast group. {port[slot/port[-slot/ port][,...]]} adds ports to a static group entry. [static|blocked] configures the route to static or blocked.
stream-limit stream-limit-max-streams *<0-65535>*	Configure multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. The default is 4.
stream-limit-group {slot/port[-slot/port] [,...]} enable max-streams *<0-65535>*	Configure multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN. The default max-streams value is 4.
version <1–3>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

Configuring IGMP ports

Configure IGMP for each interface to change default multicasting operations.

Before you begin

- You must globally enable PIM.
- You must enable PIM on the port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```

2. Enable IGMP v2-v3 compatibility mode:

   ```
   ip igmp compatibility-mode
   ```

3. Configure the system to downgrade the version of IGMP:

   ```
   ip igmp dynamic-downgrade-version
   ```

4. Configure message intervals and response times:

   ```
   ip igmp last-member-query-interval <0-255> [query-interval <1-65535>] [query-max-response <0-255>]
   ```

5. Configure expected packet loss and IGMP version:

   ```
   ip igmp robust-value <2-255> [version <1-3>]
   ```

6. Configure IGMP for a specific port:

   ```
   ip igmp port {slot/port[-slot/port][,...]}
   ```

7. Enable router alert:

   ```
   ip igmp router-alert
   ```

Example

Configure message intervals and response times:

```
VSP-9012(config-if)#ip igmp last-member-query-interval 30 query-interval 60 query-max-response 90
```

Configure expected packet loss and IGMP version:

```
VSP-9012(config-if)#ip igmp robust-value 2 version 3
```

Configure IGMP for a specific port:

```
VSP-9012(config-if)#ip igmp port 4/1
```

Enable router alert:

```
VSP-9012(config-if)#ip igmp router-alert
```

Variable definitions

Use the data in the following table to use the **`ip igmp`** command.

Variable	Value
access-list *WORD<1–64>* {A.B.C.D/X} <eny-tx|deny-rx|deny-both|allow-only-tx| allow-only-rx|allow-only-both>	Specifies the name of the access list from 1–64 characters. Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host. Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
compatibility-mode	Activates v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2.

Table continues…

Variable	Value

	To use the default configuration, use the default option in the command:
	`default ip igmp compatibility-mode`
	, or use the no option to disable compatibility mode:
	`no ip igmp compatibility-mode`

dynamic-downgrade-version	Configures if the system downgrades the version of IGMP to handle older query messages. If the system downgrades, the host with IGMPv3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:
	`default ip igmp dynamic-downgrade-version`
	or use the no option to disable downgrade:
	`no ip igmp dynamic-downgrade-version`

igmpv3-explicit-host-tracking	Enables explicit host tracking on IGMPv3. The default state is disabled.
immediate-leave	Enables fast leave on a port.
last-member-query-interval <0–255>	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. the default is 10 tenths of a second. Avaya recommends that you configure this value between 3–10 (equal to 0.3 – 1.0 seconds).
port {slot/port[-slot/port][,...]}	Configures IGMP for a specific port.
query-interval <1–65535>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.
query-max-response <0–255>	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
robust-value <2–255>	Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.

Table continues…

Variable	Value
router-alert	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. **Important:** To maximize network performance, Avaya recommends that you configure this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
stream-limit stream-limit-max-streams *<0-65535>*	Configure multicast stream limitation on a port to limit the number of concurrent multicast streams on the port. The default is 4.
version <1–3>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

Configuring IGMP on a VRF

You configure IGMP on a VRF instance the same way you configure IGMP for the Global Router, except that you must use VRF Router Configuration mode.

About this task

Use the VRF Lite feature with multicast routing protocols to create multiple virtual multicast routers

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

   ```
   enable
   configure terminal
   router vrf WORD<1-16>
   ```

2. Enable SSM dynamic learning:

   ```
   ip igmp ssm dynamic-learning
   ```

3. Configure the range group:

   ```
   ip igmp ssm group-range {A.B.C.D/X}
   ```

4. Enable the SSM map table for all static entries:

```
ip igmp ssm-map all
```

5. Create a static entry for a specific group:

   ```
   ip igmp ssm-map {A.B.C.D} {A.B.C.D} enable
   ```

6. Enable the generation of IGMP traps:

   ```
   ip igmp generate-trap
   ```

7. Enable the generation of IGMP log messages:

   ```
   ip igmp generate-log
   ```

8. Configure the fast leave mode:

   ```
   ip igmp immediate-leave-mode {multiple-user|one-user}
   ```

Example

For the VRF Red context, configure a new IP multicast group address and create an SSM map table entry for the multicast group and the source at 192.32.99.151. Configure the administrative state to enable all the static SSM map table entries.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
VSP-9012:1(config)#router vrf red
VSP-9012:1(router-vrf)#ip igmp ssm group-range 232.1.1.10/32
VSP-9012:1(router-vrf)#ip igmp ssm-map 232.1.1.10 192.32.99.151
VSP-9012:1(router-vrf)#ip igmp ssm-map all
```

Variable definitions

Use the data in the following table to use the **`ip igmp`** command on a VRF.

Variable	Value
generate-log	Enables the generation of IGMP log messages. The default is disabled.
generate-trap	Enables the generation of IGMP traps. The default is disabled.
immediate-leave-mode {multiple-user|one-user}	• multiple-user: Removes (from the group) the IGMP member who sent the leave message. The default is multiple-user. • one-user: Removes all group members on a fast leave enabled interface port after receiving the first leave message from a member.
ssm dynamic-learning	Enables dynamic learning from IGMPv3 reports. The default is enabled.
ssm group-range {A.B.C.D/X}	Changes the SSM range group to define the SSM range. The SSM range parameter extends the

Table continues…

Variable	Value
	default SSM range of 232/8 to include an IP multicast address. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.
ssm-map <all | {A.B.C.D} {A.B.C.D} enable	Creates a static SSM channel table entry by specifying the group and source IP addresses. The IP address is an IP multicast address within the SSM range. The source IP address is an IP host address that sends traffic to the group. Enables the administrative state for a specific entry (group). This variable does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry.

Chapter 5: IP multicast basic configuration using EDM

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts use a host membership protocol to subscribe to multicast services. The Internet Group Management Protocol (IGMP) is an example of a host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast–Sparse Mode (PIM–SM) and Protocol Independent Multicast–Source Specific Multicast (PIM–SSM).

Important:

To configure multicast-related protocols for a VRF you must first select and launch the VRF. Selecting and launching a VRF context view on page 68 explains how to select and launch a VRF context view.

Selecting and launching a VRF context view

Use this procedure to switch to a VRF context view and launch it so that you can view and configure features for the VRF instance.

About this task

Global Router is the default view at log in. You can configure both Global Router (GRT) and Virtual Routing and Forwarding (VRF) instances when you launch a VRF context view.

You can open only five tabs for each EDM session.

Important:

If you log out from the GRT view, the system generates a warning: all tabs close and your session terminates. If you close a VRF view tab, you close only that view.

Note:

The Set VRF Context view function is not available to users in a service provider deployment where only a tenant VRF view is assigned. If you use a tenant VRF view, Avaya recommends that you use the applicable EDM plugin with COM to access EDM. COM provides VRF mapping and Role-Based Access Control.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VRF Context View**.
2. Click **Set VRF Context View**.
3. Click the **VRF** tab.
4. Select a context to view.
5. Click **Launch VRF Context view**.

   A new browser tab appears that contains the selected VRF view.

VRF field descriptions

Use the data in the following table to use the **VRF** tab.

Name	Description
Id	Shows the unique VRF ID.
Name	Shows the name of the virtual router.
ContextName	Shows the SNMPv3 context name that denotes the VRF context and logically separates the MIB module management.

Enabling PIM-SM globally

Enable PIM-SM to offer multicasting services. After you enable PIM-SM globally and on a particular interface, the IGMP parameters take effect.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Globals** tab.
4. Click **sm** (sparse mode).

   Important:

   You can use static RP if you enable SSM for groups outside the SSM range.
5. Select the **Enable** check box.
6. Click **Apply**.

   The following message appears:

```
Are you sure you want to change the PIM mode? The traffic will not
be stopped immediately. Do you wish to continue?
```

7. Click **Yes**.

Globals field descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Mode	Configures the mode on the routing switch: sm (Sparse Mode) or ssm (Source Specific Multicast).
Enable	Enables or disables PIM.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The range is from 1–18724 and the default is 60 seconds.
RegisterSuppTimer	Specifies how long (in seconds) the designated router suppresses sending registers to the rendezvous point (RP). The timer starts after the designated router receives a register-stop message from the RP. The range is from 6–65535 and the default is 60 seconds.
UniRouteChgTimeOut	Specifies how often (in seconds) the switch polls the routing table manager for unicast routing information updates for PIM. The range is from 2–65535 and the default is 5 seconds. **Important:** If you lower this value, it increases how often the switch polls the routing table manager. This value can affect the performance of the switch, especially if a high volume of traffic flows through the switch.
DiscardDataTimeOut	Specifies how long (in seconds) to discard data until the switch receives a join message from the RP. An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The range is from 5–65535 and the default is 60 seconds.
CRPADVTimeOut	Specifies how often (in seconds) a router configured as a candidate rendezvous point router (C-RP) sends advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected bootstrap router (BSR). The range is from 5–26214 and the default is 60 seconds.
BootStrapPeriod	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. The range is from 5–32757 and the default is 60 seconds.

Table continues…

Name	Description
StaticRP	Enables or disables the static RP feature. You can use static RP to configure a static entry for an RP. A static RP permits communication with switches from other vendors that do not use the BSR mechanism.
FwdCacheTimeOut	Specifies the PIM forward cache expiry value in seconds. This value ages PIM mroutes in seconds. The range is from 10–86400 and the default value is 210. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.
FastJoinPrune	Enables or disables the PIM fast join prune feature.

Enabling PIM on a port

Enable PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you enable it on an interface.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **PIM** tab.
5. Select the **Enable** check box.
6. Click **Apply**.

PIM field descriptions

Use the data in the following table to use the **PIM** tab.

Name	Description
Enable	Enables (true) or disables (false) PIM for the specified port.
Mode	Displays the mode currently running on the routing switch.
IntfType	Indicates the interface type as active or passive.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724 seconds.

Table continues…

Name	Description
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724 seconds.
CBSRPreference	Configures the preference for this local interface to become a candidate BSR (C-BSR). The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. The range is -1-255.

Enabling SSM globally

Enable Source Specific Multicast (SSM) to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

Before you begin

- Configure a unicast protocol, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP and OSPF, see *Configuring OSPF and RIP on Avaya Virtual Services Platform 9000,* NN46250-506.
- Enable PIM globally.

Important:

After you enable PIM in SSM mode, the IGMP parameters take effect. To take full advantage of SSM, enable IGMPv3 if hosts that attach to the switch run IGMPv3 or configure the SSM table.

About this task

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Globals** tab.
4. Click **ssm** (source specific multicast).
5. Select the **Enable** check box.
6. Click **Apply**.

   The following message appears:

```
Are you sure you want to change the PIM mode? The traffic will not
be stopped immediately. Do you wish to continue?
```

7. Click **Yes**.

Enabling PIM on a VLAN interface

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you enable it on an interface.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select the VLAN ID that you want to configure with PIM.
5. Click **IP**.
6. Click the **PIM** tab.
7. Select the **Enable** check box.
8. Click **Apply**.

PIM field descriptions

Use the data in the following table to use the **PIM** tab.

Name	Description
Enable	Enables (true) or disables (false) PIM.
Mode	Displays the mode that currently runs on the switch. The valid modes are SSM and Sparse. This variable is a read-only field.
IntfType	Specifies the type of interface: active or passive.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724.

Table continues…

Name	Description
CBSRPreference	Configures the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. The range is -1-255.

Configuring IGMP parameters on a port

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Before you begin

- For IGMP parameters to take effect, enable PIM on the interface.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **IGMP** tab.
5. Edit the appropriate values.

   Note:

   To use the fast leave feature on IGMP, enable explicit-host-tracking.
6. Click **Apply**.

IGMP field descriptions

Use the data in the following table to use the **IGMP** tab.

Name	Description
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds).

Table continues…

Name	Description
	Important: You must configure this value lower than the QueryInterval.
Robustness	Configure this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect the network to lose query packets, increase the robustness value. The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.
LastMembQueryIntvl	Configures the maximum response time (in 1/10 seconds) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Avaya recommends that you configure this parameter to values greater than 3. If you do not require a fast leave process, Avaya recommends that you use values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
Version	Configures the version of IGMP (1, 2 or 3) that you want to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this port.
Maximum Number Of Stream	Configures the maximum number of streams this port permits. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This variable is a read-only value.
DynamicDowngradeEnable	Configures if the Virtual Services Platform 9000 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected (enabled), which means the switch downgrades to the oldest version of IGMP on the network.
CompatibilityModeEnable	Enables or disables v2-v3 compatibility mode. The default value is clear (disabled), which means IGMPv3 is not compatible with IGMPv2.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.

Configuring IGMP parameters on a VLAN

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Before you begin

- For IGMP parameters to take effect, enable PIM-SM on the interface.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Select **IGMP**.
7. Configure the relevant variables.
8. Click **Apply**.

IGMP field descriptions

Use the data in the following table to use the **IGMP** tab.

Name	Description
QueryInterval	Configures the frequency (in seconds) at which the IGMP host query packets transmit on the interface. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds.) **Important:** You must configure this value lower than the QueryInterval.
Robustness	Configure this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect the network to lose query packets, increase the robustness value.

Table continues…

Name	Description
	The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.
LastMembQueryIntvl	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Avaya recommends that you configure this parameter to values greater than 3. If you do not require a fast leave process, Avaya recommends that you use values greater than 10. (The value 3 is equal to 0.3 seconds, and 10 is equal to 1 second.)
SnoopEnable	Enables or disables snoop.
SsmSnoopEnable	Enables or disables support for PIM Source Specific multicast on the snoop interface.
ProxySnoopEnable	Enables or disables proxy snoop.
Version	Configures the version of IGMP (1, 2, or 3) that you want to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables or disables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this VLAN.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this VLAN. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.
FastLeavePortMembers	Selects the ports that are enabled for fast leave.
SnoopMRouterPorts	Selects the ports in this interface that provide connectivity to an IP multicast router.
DynamicDowngradeEnable	Configures if the Virtual Services Platform 9000 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected (enabled), which means the switch downgrades to the oldest version of IGMP on the network.
CompatibilityModeEnable	Enables or disables v2-v3 compatibility mode. The default value is clear (disabled), which means IGMPv3 is not compatible with IGMPv2.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.

Table continues…

Name	Description
SnoopQuerierEnable	Enables snoop querier. The default is disabled. When you enable IGMP Layer 2 Querier, Layer 2 switches in your network can snoop IGMP control packets exchanged with downstream hosts and upstream routers. The Layer 2 switches then generate the Layer 2 MAC forwarding table, used for switching sessions and multicast traffic regulation, and provide the recurring queries required to maintain IGMP groups. Enable Layer 2 Querier on only one node in the VLAN.
SnoopQuerierAddr	Specifies the pseudo IP address of the IGMP snoop querier. The default IP address is 0.0.0.0.

Chapter 6: PIM configuration using ACLI

This section provides the commands you can use to configure Protocol Independent Multicast (PIM) on Avaya Virtual Services Platform 9000. PIM provides two modes: Sparse Mode (SM) and Source Specific Multicast (SSM).

Before you begin

- Configure an IP interface. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.
- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM. For more information about RIP and OSPF, see *Configuring OSPF and RIP on Avaya Virtual Services Platform 9000,* NN46250-506.
- Enable PIM-SM globally.
- Enable PIM-SM on individual interfaces.
- You must first configure and enable PIM on the circuitless IP interface before you can utilize that interface as a candidate rendezvous point (RP). To configure PIM-SM RP for a circuitless IP interface, see Configuring a candidate rendezvous point on page 83.

Changing the interface status to passive

Change the PIM interface status to passive to deny PIM control traffic on the interface.

Before you begin

- The PIM interface is disabled.

About this task

The command you use depends on the required administrative state of the interface (enable or disable).

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```
   or `interface vlan <1-4084>`

2. Create a passive interface and enable it simultaneously:

   ```
   ip pim passive
   ```

3. Create a passive interface in the disabled state:

   ```
   ip pim interface-type passive
   ```

   You must manually enable the interface.

4. Enable a disabled interface:

   ```
   ip pim enable
   ```

Variable definitions

Use the data in the following table to use the **ip pim** command.

Variable	Value
active	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active. To configure this option to the default value, use the default operator with the command.
passive	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active. To configure this option to the default value, use the default operator with the command.

Changing the interface status to active

Change the PIM interface status to active to allow PIM control traffic on the interface.

Before you begin

- The PIM interface is disabled.

About this task

The command you use depends on the required administrative state of the interface (enable or disable).

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```
   or `interface vlan <1-4084>`

2. Create an active interface in the disabled state:

   ```
   ip pim interface-type active
   ```

   You must manually enable the interface.

3. Create an active interface and enable it simultaneously:

   ```
   ip pim active
   ```

   OR

   ```
   ip pim enable
   ```

   The second command enables an active interface only if this is the first PIM interface you create on the port or VLAN or you created an active interface in the disabled state. If you already created a passive interface in the disabled state, the second command enables that passive interface.

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
active	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to

Table continues…

Variable	Value
	other switches. The default is active. To configure this option to the default value, use the default operator with the command.
passive	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active. To configure this option to the default value, use the default operator with the command.

Configuring the PIM virtual neighbor

Configure a PIM virtual neighbor if the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the PIM virtual neighbor:

   ```
   ip pim virtual-neighbor <A.B.C.D> <A.B.C.D>
   ```

Example

Configure the PIM virtual neighbor:

```
VSP-9012(config)#ip pim virtual-neighbor 2.2.2.245 3.3.3.245
```

Variable definitions

Use the data in the following table to use the **`ip pim virtual-neighbor`** command.

Variable	Value
{A.B.C.D}	The first IP address indicates the IP address of the selected interface. The second IP address indicates the IP address of the neighbor.

Configuring a candidate rendezvous point

Configure a candidate rendezvous point (C-RP) to serve as backup to the RP router.

About this task

You can configure only one interface on an Avaya Virtual Services Platform 9000 for multiple groups. You cannot configure multiple interfaces for multiple groups.

With the mask value, you can configure a C-RP router for several groups in one configuration.

For example, if you use a C-RP configuration with a group address of 224.0.0.0 and a group mask of 240.0.0.0, you can configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Add a candidate rendezvous point:

   ```
   ip pim rp-candidate group <A.B.C.D> <A.B.C.D> rp <A.B.C.D>
   ```

3. Remove a candidate rendezvous point:

   ```
   no ip pim rp-candidate group <A.B.C.D> <A.B.C.D>
   ```

4. Display information about the candidate rendezvous points for the PIM-SM domain:

   ```
   show ip pim rp-candidate
   ```

Example

Add a candidate rendezvous point:

```
VSP-9012(config)#ip pim rp-candidate group 224.1.1.0 255.255.255.0 rp
30.1.1.1
```

Variable definitions

Use the data in the following table to use the **`ip pim rp-candidate`** command.

Variable	Value
group {A.B.C.D} {A.B.C.D}	Specifies the IP address and the address mask of the multicast group. After the IP address and group mask are combined, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
rp {A.B.C.D}	Specifies the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Job aid

The following table shows the field descriptions for the **`show ip pim rp-candidate`** command.

Table 5: show ip pim rp-candidate field descriptions

Field	Description
GRPADDR	Displays the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Displays the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
RPADDR	Displays the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Configuring static RP

Configure a static RP to ignore the bootstrap router (BSR) mechanism and use the statically configured RPs only mechanism.

Before you begin

- Enable PIM-SM globally.

About this task

Static RP-enabled switches use this feature to communicate with switches from other vendors that do not use the BSR.

Important:

You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.

All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable static RP:

   ```
   ip pim static-rp
   ```

The following message appears:

```
WARNING: RP information learnt dynamically through BSR functionality will be lost.
Do you wish to enable Static RP? (y/n) ?
```

3. Enter `y`.
4. Configure a static RP entry:

   ```
   ip pim static-rp {A.B.C.D/X} {A.B.C.D}
   ```

5. Configure all the switches in the network (including switches from other vendors) to map to the same RP.
6. Display information about the candidate rendezvous points for the PIM-SM domain:

   ```
   show ip pim static-rp
   ```

Example

```
VSP-9012:1(config)# ip pim static-rp 239.255.0.0/255.255.0.0 100.1.1.1
```

Variable definitions

Use the data in the following table to use the **`ip pim static-rp`** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and address mask of the multicast group. When combined, the IP address and address mask identify the range of the multicast addresses that the RP handles.
{A.B.C.D}	Specifies the IP address of the static RP.

Configuring a candidate BSR on a port

Configure additional routers as candidate BSRs (C-BSR) to provide backup protection in the event that the primary BSR fails. PIM-SM cannot run without a BSR.

Before you begin

- Static RP is disabled.

About this task

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port][,...]}
```

2. Configure a candidate BSR:

```
ip pim bsr-candidate preference <0–255>
```

Example

Configure a candidate BSR:

```
VSP-9012(config-if)#ip pim bsr-candidate preference 2
```

Variable definitions

Use the data in the following table to use the **`ip pim bsr-candidate`** command.

Variable	Value
preference *<0–255>*	Activates the C-BSR on this interface and configures its preference value, from 0–255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. To set this option to the default value, use the default operator with the command.

Configuring a candidate BSR on a VLAN

Configure additional routers as candidate BSRs (C-BSR) to provide backup protection in the event that the primary BSR fails. PIM-SM cannot run without a BSR.

Before you begin

- Static RP is disabled.

About this task

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1–4084>
```

2. Configure a candidate BSR on a VLAN:

   ```
   ip pim bsr-candidate preference <0–255>
   ```

Example

Configure a candidate BSR on a VLAN:

```
VSP-9012(config-if)#ip pim bsr-candidate preference 5
```

Variable definitions

Use the data in the following table to use the **`ip pim bsr-candidate`** command.

Variable	Value
preference *<0–255>*	Activates the C-BSR on this interface and configures its preference value, from 0–255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. To configure this option to the default value, use the default operator with the command.

Enabling square-SMLT globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

About this task

Important:

The following command also activates full-mesh configurations.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable square-SMLT:

   ```
   multicast smlt-square
   ```

Chapter 7: PIM configuration using EDM

Avaya Virtual Services Platform 9000 supports two modes of Protocol Independent Multicast (PIM): Sparse Mode (SM) and Source Specific Multicast (SSM).

- PIM-SM supports multicast groups spread out across large areas of a company or the Internet.
- PIM-SSM optimizes PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Before you begin

- Before you can configure PIM-SM, you must configure an IP interface. For more information, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing, NN46250-505*.
- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM. For more information about RIP and OSPF, see *Avaya Virtual Services Platform 9000 Configuration — OSPF and RIP, NN46250-506*
- Enable PIM-SM globally.
- Enable PIM-SM on individual interfaces.
- Configure one or more rendezvous points (RP) for the groups that multicast applications use in the network.

  Important:

  If you configure the rendezvous point (RP) to be the address of a circuitless IP (CLIP) interface, then you must first configure and enable PIM on the CLIP interface before you can utilize that interface as a candidate RP. To configure a PIM-SM RP for a circuitless IP interface, see Configuring a candidate RP on page 96.
- Configure one or more bootstrap routers (BSR) to propagate RP information to all switches in the network.

Enabling static RP

Enable static RP to avoid the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.

3. Click the **Globals** tab.
4. Select **sm** (sparse mode).
5. Select **Enable**.
6. Select **Static RP**.
7. Click **Apply**.

   The following message appears:

   ```
   RP information learnt dynamically through BSR functionality will be
   lost. Do you wish to enable Static RP?
   ```

8. Click **Yes**.

   The following message appears:

   ```
   Are you sure you want to change the PIM mode? The traffic will not
   be stopped immediately. Do you wish to continue?
   ```

9. Click **Yes**.

Configuring a static RP

Configure a static RP to ignore the BSR mechanism and use the statically configured RPs only. A static RP-enabled switch uses this feature to communicate with switches from other vendors that do not use the BSR mechanism.

Before you begin

- Before you can configure a static RP, you must enable the following:
  - PIM-SM
  - static RP

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Static RP** tab.
4. Click **Insert**.
5. Type the required information in each box.
6. Click **Insert**.

Static RP field descriptions

Use the data in the following table to use the **Static RP** tab.

Name	Description
GroupAddress	Configures the IP address of the multicast group. When combined with the group mask, this value identifies the range of the multicast addresses that the RP handles.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the range of the multicast addresses that the RP handles.
Address	Configures the IP address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid if the switch uses a unicast route to the network for the static RP and is invalid otherwise.

Job aid

Keep in mind the following configuration considerations:

- You cannot configure a static RP-enabled switch as a BSR or as a candidate RP (C-RP) router.
- All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.
- Static RPs do not age; they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, it does not know about the static RP unless you configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP.
- In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interfaces as an RP.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of Avaya and other vendor switches across the network, ensure that all switches or routers use the same active RP because other vendors use different algorithms to elect the active RP. The Avaya Virtual Services Platform 9000 uses the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.
- Static RP on the switch is active as long as the switch uses a unicast route to the network for the static RP. If the switch loses this route, the static RP is invalidated, and the hash algorithm is invoked to remap all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm is invoked to remap the affected groups.

Viewing the active RP

Perform this procedure to show information about the active RP for all the running multicast groups on the switch.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Active RP** tab.

Active RP field descriptions

Use the data in the following table to use the **Active RP** tab.

Name	Description
GroupAddress	Shows the IP address of the multicast group.
Address	Shows the IP address of the RP router. This address must be one of the local PIM-SM enabled interfaces.
Priority	Shows the priority of the RP.

Configuring a candidate bootstrap router

Configure routers as candidate bootstrap routers (C-BSR) to provide backup protection in case the primary BSR fails. PIM-SM cannot operate without a BSR. A PIM-SM domain can use only one active BSR.

About this task

The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **PIM** tab.
5. Click **Enable**.
6. In the **CBSRPreference** box, type the preference.

The C-BSR with the highest BSR-preference and address becomes the active BSR. The default is –1, which indicates that the current interface is not a C-BSR.

7. Click **Apply**.

Viewing current BSR information

View the current BSR information to review the configuration.

Before you begin

- You must disable static RP.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Current BSR** tab.

Current BSR field descriptions

Use the data in the following table to use the **Current BSR** tab.

Name	Description
Address	Shows the IP address of the current BSR for the local PIM domain.
FragmentTag	Shows a randomly generated number that distinguishes fragments that belong to different bootstrap messages. Fragments that belong to the same bootstrap message carry the same fragment tag.
HashMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. The hashmask allows a small number of consecutive groups to always hash to the same RP.
Priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
BootStrapTimer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

Changing VLAN interface type

Change the state (active or passive) of PIM on a VLAN interface.

Before you begin

- Before you change the state of PIM on a VLAN interface, you must first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when the interface receives streams.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select the VLAN ID that you want to configure with PIM.
5. Click **IP**.
6. Click the **PIM** tab.
7. Clear the **Enable** check box.
8. Click **Apply**.
9. Select **active** or **passive**.
10. Click **Apply**.
11. Reenable PIM on the VLAN interface.

Editing PIM interface parameters

Edit PIM parameters for an interface to customize the PIM configuration.

Before you begin

- Before you change the state (active or passive) of a PIM interface, first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when the interface receives streams.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Interfaces** tab.
4. Edit the fields by double-clicking on them, and then select or type the new value.
5. Click **Apply**.

Interfaces field descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IfIndex	Shows the interface Index. This variable is a read-only field.
Address	Shows the IP address of the PIM interface. This variable is a read-only field.
NetMask	Shows the network mask for the IP address of the PIM interface. This variable is a read-only field.
Mode	Shows the configured mode of this interface. The valid modes are SSM and sparse. This variable is a read-only field.
InterfaceType	Specifies if the interface is active or passive.
DR	Shows the router with the highest IP address on a LAN designated to perform these tasks.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default is 30 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds.
CBSRPreference	Configures the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
OperState	Indicates the status of PIM on this interface: enabled or disabled.

Configuring the PIM virtual neighbor

Configure a PIM virtual neighbor if the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Virtual Neighbors** tab.
4. Click **Insert**.
5. Specify the IP address of the virtual neighbor.
6. Specify the interface index for the PIM interface.
7. Click **Insert**.

Virtual Neighbors field descriptions

Use the data in the following table to use the **Virtual Neighbors** tab.

Name	Description
Address	Specifies the IP address of the neighbor.
IfIndex	Specifies the IP address of the PIM interface.

Viewing PIM-SM neighbor parameters

View PIM-SM neighbor parameters to troubleshoot connection problems or review the configuration.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Neighbors** tab.

Neighbors field descriptions

Use the data in the following table to use the **Neighbors** tab.

Name	Description
Address	Shows the IP address of the PIM neighbor.
IfIndex	Shows the slot and port number or VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Shows the time since this neighbor became a neighbor of the local router.
ExpiryTime	Shows the time remaining before the neighbor expires.

Viewing RP set parameters

View the RP set to see a list of rendezvous point addresses. The BSR constructs this list from C-RP advertisements, and then distributes it to all PIM routers in the PIM domain for the BSR. View the parameters for troubleshooting purposes.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **RP Set** tab.

RP Set field descriptions

Use the data in the following table to use the **RP Set** tab.

Name	Description
GroupAddress	Shows the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GroupMask	Shows the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
Address	Shows the IP address of the C-RP router.
HoldTime	Shows the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
ExpiryTime	Shows the time remaining before this C-RP router times out.

Configuring a candidate RP

Configure a C-RP router to add it to the RP Set.

About this task

You can configure only one interface on a Avaya Virtual Services Platform 9000 for multiple groups; that is, you cannot configure multiple interfaces for multiple groups.

Using the GroupMask value, you can configure a candidate RP for several groups in one configuration. For example, if you use a C-RP configuration with a GroupAddress value of 224.0.0.0 and a GroupMask of 240.0.0.0, you can configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Candidate RP** tab.
4. Click **Insert**.
5. Type the required information in each box.
6. Click **Insert**.

Candidate RP field descriptions

Use the data in the following table to use the **Candidate RP** tab.

Name	Description
GroupAddress	Configures the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
InterfaceAddress	Configures the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Enabling square-SMLT globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

About this task

Important:

The following command also activates full-mesh configurations.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**
2. Click **Multicast**.
3. Click the **Globals** tab.
4. Select **MulticastSquareSmltEnable**.

   Clear this check box if you want to disable square-SMLT globally.
5. Click **Apply**.

Chapter 8: IGMP configuration using ACLI

Hosts use the Internet Group Management Protocol (IGMP) to report their IP multicast group memberships to neighboring multicast routers. Configure IGMP on an individual interface basis.

- Complete one of the following tasks:
  - Configure IGMP on a Layer 2 interface by enabling IGMP snoop.
  - Configure IGMP on a Layer 3 interface by enabling multicast routing, for example, Protocol Independent Multicast-Sparse Mode (PIM-SM) or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

Configuring multicast stream limitation on an Ethernet port

Configure multicast stream limitation on an Ethernet port to limit the number of concurrent multicast streams on the port. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the port drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```

2. Enable multicast stream limitation and configure the maximum number of allowed streams:

   ```
   ip igmp stream-limit stream-limit-max-streams <0-65535>
   ```

3. If stream-limit is already enabled on the interface, change the maximum number of allowed streams:

```
ip igmp stream-limit stream-limit-max-streams <0-65535>
```

4. Display multicast stream limitation information for the ports on a specific interface:

```
show ip igmp stream-limit interface
```

Example

Enable multicast stream limitation on the Ethernet port.

```
VSP-9012:1(config-if)# ip igmp stream-limit
```

Configure the maximum number of allowed streams to 8.

```
VSP-9012:1(config-if)# ip igmp stream-limit stream-limit-max-streams 8
```

Variable definitions

Use the data in the following table to use the **ip igmp stream-limit-max-streams** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this port. The range is from 0–65535 and the default is 4.

Job aid

The following tables show the field descriptions for the **show ip igmp stream-limit interface** command.

Table 6: show ip igmp stream-limit interface field descriptions

Field	Description
INTERFACE	Indicates the interface IP address.
MAX STREAMS	Indicates the maximum number of streams.
NUM STREAMS	Indicates the current number of streams.

Configuring multicast stream limitation on a VLAN

Configure multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the VLAN drops joins to new streams. A service provider uses this feature to

control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Enable multicast stream limitation and configure the maximum number of allowed streams:

   ```
   ip igmp stream-limit stream-limit-max-streams <0-65535>
   ```

3. If stream-limit is already enabled on the VLAN, change the maximum number of allowed streams:

   ```
   ip igmp stream-limit stream-limit-max-streams <0-65535>
   ```

4. Display multicast stream limitation information for the ports on a specific interface:

   ```
   show ip igmp stream-limit port
   ```

Example

Enable multicast stream limitation.

```
VSP-9012:1(config-if)# ip igmp stream-limit
```

Configure the maximum number of allowed streams to 8.

```
VSP-9012:1(config-if)# ip igmp stream-limit stream-limit-max-streams 8
```

Variable definitions

Use the data in the following table to use the **`ip igmp stream-limit`** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this VLAN. The range is from 0–65535 and the default is 4.

Job aid

The following tables show the field descriptions for the **`show ip igmp stream-limit port`** command.

Table 7: show ip igmp stream-limit port field descriptions

Field	Description
INTERFACE	Indicates the interface IP address.
PORT	Indicates the port for the VLAN.
MAX STREAMS	Indicates the maximum number of streams.
NUM STREAMS	Indicates the current number of streams.

Configuring VLAN multicast stream limitation members

Configure multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable

   configure terminal

   interface vlan <1–4084>
   ```

2. Configure multicast stream limitation members on a VLAN:

   ```
   ip igmp stream-limit-group {slot/port[-slot/port][,...]} enable max-streams <0–65535>
   ```

Example

Enable multicast stream limitation on ports 3/3 to 3/8 and configure the maximum allowed number of streams to 6 for this interface.

```
VSP-9012:1(config-if)# ip igmp stream-limit-group 3/3-3/8 max-streams 6
```

Variable definitions

Use the data in the following table to use the **`ip igmp stream-limit-group`** command.

Variable	Value
<0–65535>	Configures the maximum number of allowed streams for the specified ports on this VLAN. The range is from 0–65535 and the default is 4.
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Configuring multicast router discovery options

Configure the multicast router discovery options to enable the automatic discovery of multicast-capable routers.

About this task

Important:

Avaya Virtual Services Platform 9000 does not support the Multicast Router Discovery (MRDISC) protocol on brouter ports.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Enable multicast router discovery:

   ```
   ip igmp mrdisc
   ```

3. Configure the maximum advertisement intervals between successive advertisements:

   ```
   ip igmp mrdisc maxadvertinterval <2-180> maxinitadvertinterval <2-180>
   ```

4. Configure the maximum advertisements after initialization:

   ```
   ip igmp mrdisc maxinitadvertisements <2-15>
   ```

5. Configure the minimum advertisement interval between successive advertisements:

   ```
   ip igmp mrdisc minadvertinterval <3-180>
   ```

6. Configure the time allowed before a neighbor is declared dead:

   ```
   ip igmp mrdisc neighdeadinterval <2-180>
   ```

Example

Configure the maximum advertisement intervals between successive advertisements:

```
VSP-9012(config-if)#ip igmp mrdisc maxadvertinterval 30
maxinitadvertinterval 5
```

Configure the maximum advertisements after initialization:

```
VSP-9012(config-if)#ip igmp mrdisc maxinitadvertisements 8
```

Configure the minimum advertisement interval between successive advertisements:

```
VSP-9012(config-if)#ip igmp mrdisc minadvertinterval 30
```

Configure the time allowed before a neighbor is declared dead:

```
VSP-9012(config-if)#ip igmp mrdisc neighdeadinterval 60
```

Variable definitions

Use the data in the following table to use the **`ip igmp mrdisc`** command.

Variable	Value
maxadvertinterval <*2–180*>	Configures the maximum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration, and then reset the switch. To configure this option to the default value, use the default operator with the command. The default is 20.
maxinitadvertinterval <2–180>	Configures the maximum number (in seconds) between successive initial advertisements. For this change to take effect, you must save the configuration, and then reset the switch. To configure this option to the default value, use the default operator with the command. The default is 2.
maxinitadvertisements <2–15>	Configures the maximum number of initial multicast advertisements after initialization. For this change to take effect, you must save the configuration, and then reset the switch. To configure this option to the default value, use the default operator with the command. The default is 3.
minadvertinterval <*3–180*>	Configures the minimum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration, and then reset the switch. To configure this option to the default value, use the default operator with the command. The default is 15.
neighdeadinterval <*2–180*>	Configures the multicast router discovery dead interval—the number of seconds the multicast route neighbors for the switch must wait before assuming that the multicast router is down. To configure this option to the default value, use the default operator with the command. The default is 60.

Configuring explicit host tracking

Configure explicit host tracking to track all the source and group members.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable

   configure terminal

   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```
   or `interface vlan <1-4084>`

2. Configure explicit host tracking:

   ```
   ip igmp igmpv3-explicit-host-tracking
   ```

3. Display all the tracked members for a specific group:

   ```
   show ip igmp group group <A.B.C.D> tracked-members [member-subnet
   <A.B.C.D/X>] [source-subnet <A.B.C.D/X>] [port {slot/port[-slot/
   port][,...]}] [vlan <1-4084>]
   ```

4. Display the IGMPv3 specific data:

   ```
   show ip igmp group group <A.B.C.D> detail port {slot/port [-slot/
   port][,...]} vlan <1-4084>
   ```

Example

Configure explicit host tracking:

```
VSP-9012:1(config-if)#ip igmp igmpv3-explicit-host-tracking
```

Display all the tracked members:

```
VSP-9012:1(config-if)#show ip igmp group
```

```
================================================================================
                                Igmp Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
   225.1.1.1       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.2       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.3       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.4       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.5       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.6       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.7       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.8       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.9       V22-4/1         22.22.22.200    178        Dynamic
   225.1.1.10      V22-4/1         22.22.22.200    178        Dynamic
   225.12.12.1     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.2     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.3     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.4     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.5     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.6     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.7     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.8     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.9     V2222-7/16      22.2.2.200      172        Dynamic
   225.12.12.10    V2222-7/16      22.2.2.200      172        Dynamic
   226.1.1.1       V33-4/23        33.33.33.200    173        Dynamic
   226.1.1.2       V33-4/23        33.33.33.200    173        Dynamic
```

```
   226.1.1.3       V33-4/23        33.33.33.200    173         Dynamic
   226.1.1.4       V33-4/23        33.33.33.200    173         Dynamic
   226.1.1.5       V33-4/23        33.33.33.200    173         Dynamic
   226.1.1.6       V33-4/23        33.33.33.200    173         Dynamic
   226.1.1.7       V33-4/23        33.33.33.200    173         Dynamic
   226.1.1.8       V33-4/23        33.33.33.200    173         Dynamic
   226.1.1.9       V33-4/23        33.33.33.200    173         Dynamic
   226.1.1.10      V33-4/23        33.33.33.200    173         Dynamic
   226.22.22.1     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.2     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.3     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.4     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.5     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.6     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.7     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.8     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.9     V3333-7/22      33.3.3.200      173         Dynamic
   226.22.22.10    V3333-7/22      33.3.3.200      173         Dynamic
   228.45.45.45    V222-4/1        122.122.122.200 173         Dynamic
   228.56.56.56    V222-4/1        122.122.122.200 166         Dynamic
   229.1.1.1       V333-4/17       133.133.133.200 172         Dynamic
   229.32.32.32    V222-4/1        122.122.122.200 169         Dynamic
   232.1.1.1       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.2       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.3       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.4       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.5       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.6       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.7       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.8       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.9       V333-4/17       133.133.133.200 170         Dynamic
   232.1.1.10      V333-4/17       133.133.133.200 170         Dynamic
   232.32.32.1     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.2     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.3     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.4     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.5     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.6     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.7     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.8     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.9     V222-4/1        122.122.122.200 165         Dynamic
   232.32.32.10    V222-4/1        122.122.122.200 162         Dynamic
   232.42.42.1     V222-4/1        122.122.122.200 167         Dynamic

65 out of 65 group Receivers displayed

Total number of unique groups 65
```

Display all the tracked members for a specific group:

```
VSP-9012:1(config-if)#show ip igmp group group 232.1.1.1 tracked-members
```

```
================================================================================
                          Members of Channels/Groups - GlobalRouter
================================================================================
INTERFACE       CHANNEL/GROUP         MEMBER           MEMBER_MODE EXP
--------------------------------------------------------------------------------
Vlan333-3/30    */232.1.1.1          133.133.133.200   IS_EXCLUDE  205

Note:
The "*" attached to the interface (if any) indicates that the interface has explicit host
tracking disabled.
```

Display IGMPv3 specific data:

```
VSP-9012:1(config-if)#show ip igmp group group 232.32.32.10 detail
================================================================================
                        Igmp Group Detail - GlobalRouter
================================================================================
Interface:                      Vlan222-4/10
IGMPv3 Group:                   232.32.32.10
Interface Group Mode:           INCLUDE
Interface Compatibility Mode:   IGMP_V3
V2 Host Timer:                  Not Running
V1 Host Timer:                  Not Running
Interface Group Include Source List:
        Source Address  Expires
        133.133.133.200 114
```

Variable definitions

Use the data in the following table to use the **`ip igmp igmpv3-explicit-host-tracking`** command.

Variable	Value
explicit-host-tracking	Enables explicit host tracking on IGMPv3. The default state is disable.
<A.B.C.D>	Specifies the IP address of the group of the tracked member.

Configuring IGMP static members

Configure IGMP static members to add members to a snoop group. You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams are always forwarded to the multicast router within the VLAN, in addition to the ports in this static entry.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Configure interface static members:

   ```
   ip igmp static-group {A.B.C.D} {A.B.C.D} {port[slot/port[-slot/port]
   [,...]]} [static|blocked]
   ```

Example

Configure interface static members:

```
VSP-9012(config-if)#ip igmp static-group 239.1.1.1 239.1.2.1 port 2/1
static
```

Variable definitions

Use the data in the following table to use the **ip igmp static-group** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Indicates the IP address range of the selected multicast group.
port	Adds ports to a static group entry
{slot/port[-slot/port][,...]}	Creates a static group entry. Specifies the port or list of ports that is a member of the VLAN interface being configured to which you want to redirect the multicast stream for this multicast group. Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2). Use the no operator to later remove this configuration.
<static|blocked>	Configures the route to static or blocked.

Configuring SSM dynamic learning and range group

Configure SSM dynamic learning and a range group to enable the IGMPv3 dynamic learning feature and to extend the default SSM range of 232/8 to include an IP multicast address. As new SSM channels are learned, they appear in the SSM channel table.

Before you begin

- To define the range group, you must first disable PIM.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable SSM dynamic learning:

   ```
   ip igmp ssm dynamic-learning
   ```

3. Configure the range group:

   ```
   ip igmp ssm group-range <A.B.C.D/X>
   ```

Example

Define the SSM range group address (234.0.0.0) and mask (255.0.0.0).

```
VSP-9012:1(config)# ip igmp ssm group-range 234.0.0.0/255.0.0.0
```

Enable dynamic learning from IGMPv3 reports.

```
VSP-9012:1(config)# ip igmp ssm dynamic-learning
```

Variable definitions

Use the data in the following table to use the **ip igmp ssm** command.

Variable	Value
{A.B.C.D/X}	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Changing the SSM range group

Change the SSM range group to define the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address.

About this task

Important:

This procedure reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), it also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap router (BSR) is local.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable PIM:

```
no ip pim enable
```

If you forget to disable PIM, the following error message appears:

```
Error: PIM is enabled in SSM mode, disable PIM
```

3. Delete each entry in the SSM channel table:

   ```
   no ip igmp ssm-map [all] [{A.B.C.D} enable]
   ```

   If you forget to delete the SSM channels, the following error message appears:

   ```
   Error: SSM source group table not empty
   ```

4. Configure the new IP multicast group address:

   ```
   ip igmp ssm group-range {A.B.C.D/X}
   ```

5. Enable PIM:

   ```
   ip pim enable
   ```

Example

Configure the new IP multicast group address:

```
VSP-9012(config-if)#ip igmp ssm group-range 232.1.1.10/16
```

Variable definitions

Use the data in the following table to use the **ip igmp ssm group-range** and **ip igmp ssm** commands.

Variable	Value
{A.B.C.D/X}	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Configuring the SSM map table

Configure the SSM map table to map groups to their sending source. SSM maps cannot conflict with static source groups. After you configure an SSM map or a static source group, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) to different sources or multiple sources to the same group for both static source group and an SSM map.

About this task

The consistency check applies to all SSM map entries, even if they are disabled. If you disable an entry, it becomes inactive. If you do not delete the entry, you can reenable it later.

After you disable an SSM map, the Avaya Virtual Services Platform 9000 stops multicast traffic from the specified source to the specified group. You can use this static configuration as a security feature to block traffic from a certain source to a specific group.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable the SSM map table for all static entries:

   ```
   ip igmp ssm-map all
   ```

3. Create a static entry for a specific group:

   ```
   ip igmp ssm-map {A.B.C.D} {A.B.C.D} enable
   ```

Example

Create an SSM map table entry for the multicast group 234.0.1.0 and the source at 192.32.99.151.

```
VSP-9012:1(config)# ip igmp ssm-map 234.0.1.0 192.32.99.151
```

Configure the administrative state to enable all the static SSM map table entries.

```
VSP-9012:1(config)# ip igmp ssm-map all
```

Variable definitions

Use the data in the following table to use the **ip igmp ssm-map** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Creates a static SSM channel table entry by specifying the group and source IP addresses. The IP address is an IP multicast address within the SSM range. The source IP address is an IP host address that sends traffic to the group.
{A.B.C.D} enable	Enables the administrative state for a specific entry (group). This variable does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry.

Configuring multicast access control for an IGMP Ethernet port

Configure multicast access control for an IGMP Ethernet port to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable

   configure terminal

   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```

2. Configure multicast access control:

   ```
   ip igmp access-list WORD<1-64> {A.B.C.D/X} <deny-tx|deny-rx|deny-
   both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

3. Change an existing access list:

   ```
   ip igmp access-list WORD<1-64>> {A.B.C.D/X} mode <deny-tx|deny-rx|
   deny-both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp access-list`** command

Variable	Value
{A.B.C.D/X}	Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
deny-tx|deny-rx|deny-both|allow-only-tx| allow-only-rx|allow-only-both	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
mode	Changes the access control group configuration.
WORD<1–64>	Specifies the name of the access list from 1–64 characters.

Configuring multicast access control for a VLAN

Configure multicast access control for an IGMP VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Configure multicast access control:

   ```
   ip igmp access-list WORD<1-64> {A.B.C.D/X} <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

3. Change an existing access list:

   ```
   ip igmp access-list WORD<1-64> {A.B.C.D/X} mode <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp access-list`** command.

Variable	Value
{A.B.C.D/X}	Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
mode	Changes the access control group configuration.
WORD<1–64>	Specifies the name of the access list from 1–64 characters.

Configuring fast leave mode

Configure fast (immediate) leave mode to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces. Normal IGMP behavior is skipped. Fast leave mode provides one command that controls all IGMP fast leave enabled interfaces.

Before you begin

- You must enable explicit-host-tracking before configuring fast-leave mode for IGMPv3. For more information on enabling explicit-host-tracking, see Configuring explicit host tracking on page 103.

About this task

If a single user connects to an interface, you do not need to track if other users exist on the interface to perform the fast leave. In cases like this, you must change the mode to one-user.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. View the current fast leave mode:

   ```
   show ip igmp sys
   ```

3. Configure fast leave mode:

   ```
   ip igmp immediate-leave-mode <multiple-user|one-user>
   ```

Example

Change the mode to one-user.

```
VSP-9012:1(config)# ip igmp immediate-leave-mode one-user
```

Variable definitions

Use the data in the following table to use the **`ip igmp immediate-leave-mode`** command.

Variable	Value
multiple-user|one-user	multiple-user removes from the group only the IGMP member who sent the leave message. Traffic does not stop if other receivers exist on the interface port. This configuration is the default. one-user removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.

Enabling fast leave mode on a port

Enable fast (immediate) leave mode to specify if a port receives a leave message from a member of a group. If you enable fast leave mode on a port, it uses the global fast leave mode configuration.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```

2. Enable fast leave:

   ```
   ip igmp immediate-leave
   ```

Configuring IGMP fast leave members on a VLAN

Configure IGMP fast leave members on a VLAN to specify fast leave capable ports.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Enable fast leave on the VLAN:

   ```
   ip igmp immediate-leave
   ```

3. Configure fast leave members on a VLAN:

   ```
   ip igmp immediate-leave-members {slot/port[-slot/port][,...]}
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp immediate-leave-members`** command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Enabling IGMP Layer 2 Querier

When no multicast router exists in your network, you can use IGMP Layer 2 Querier to allow the Layer 2 switch to act as a multicast router so that the system can participate in multicast environments where multicast routing is not required.

Before you begin

- You must enable IGMP snooping.

About this task

When you enable IGMP Layer 2 Querier, Layer 2 switches in your network can snoop IGMP control packets exchanged with downstream hosts and upstream routers. The Layer 2 switches then generate the Layer 2 MAC forwarding table, used for switching sessions and multicast traffic regulation, and provide the recurring queries required to maintain IGMP groups.

By default, IGMP Layer 2 Querier is disabled.

Enable Layer 2 Querier on only one node in the VLAN.

On Shortest Path Bridging (SPB) Customer VLANs (CVLAN), IGMP Querier is enabled automatically when you enable snooping on the VLAN. For more information about SPB, see *Configuring Avaya Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Enable IGMP Layer 2 Querier:

   ```
   ip igmp snoop-querier
   ```

Next steps

You must enable the IGMP Layer 2 Querier address. See Enabling IGMP Layer 2 Querier Address on page 115

Enabling IGMP Layer 2 Querier address

To use the IGMP Layer 2 Querier feature you must designate the IGMP Layer 2 Querier source IP address, the address the system uses in the query message.

Before you begin

- Enable IGMP Layer 2 Querier.

About this task

You must configure the IGMP Layer 2 Querier address to an IP address in the IP subnet that IGMP hosts, and to which IGMP snoopers in the VLAN belong.

The default IP address is 0.0.0.0 when the IGMP Layer 2 Querier is disabled.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable

   configure terminal

   interface vlan <1-4084>
   ```

2. Enable the IGMP Layer 2 Querier address:

   ```
   ip igmp snoop-querier-addr {A.B.C.D}
   ```

3. Verify the configuration:

   ```
   show ip igmp snooping [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

Enable the IGMP Layer 2 Querier feature for VLAN 4, and configure the querier address. Verify the configuration.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
VSP-9012:1(config)#interface vlan 4
VSP-9012:1(config-if)#ip igmp snoop-querier
VSP-9012:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
VSP-9012:1(config-if)#show ip igmp snooping

================================================================================
                          Igmp Snooping - GlobalRouter
================================================================================
IFINDEX SNOOP   PROXY   SSM     STATIC           ACTIVE           MROUTER
        ENABLE  SNOOP   SNOOP   MROUTER          MROUTER          EXPIRATION
                ENABLE  ENABLE  PORTS            PORTS            TIME
--------------------------------------------------------------------------------
V2      false   false   false                                     0
V3      false   false   false                                     0
V4      true    false   false                                     0
V200    false   false   false                                     0

IFINDEX SNOOP   SNOOP           DYNAMIC   COMPATIBILITY
        QUERIER QUERIER         DOWNGRADE MODE
        ENABLE  ADDRESS         VERSION
--------------------------------------------------------------------------------
V2      false   0.0.0.0         enable    disable
V3      false   0.0.0.0         enable    disable
V4      true    192.0.2.1       enable    disable
V200    false   0.0.0.0         enable    disable

4 out of 4 entries displayed
```

Chapter 9: IGMP configuration using EDM

Hosts use Internet Group Management Protocol (IGMP) to report their multicast group memberships to neighboring multicast routers.

Before you begin

- Configure IGMP on a Layer 3 interface by first enabling multicast routing, for example, Protocol Independent Multicast-Sparse Mode (PIM-SM), or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

Important:

To configure and use IGMP on a VRF instance you must first select and launch the VRF context.

To select and launch the VRF context, see Selecting and launching a VRF context view on page 68.

Enabling IGMP snoop on a VLAN

Enable IGMP snooping on a VLAN to optimize the multicast data flow for a group within a VLAN to only those that are members of the group that uses IGMP snoop.

About this task

The switch listens to group reports from each port and builds a database of multicast group members for each port. The switch suppresses the reports heard by not forwarding them to other hosts, forcing the members to continuously send their own reports.

The switch relays group membership from the hosts to the multicast routers and forwards queries from multicast routers to all port members of the VLAN. The switch multicasts data only to the participating group members and to the multicast routers within the VLAN.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.

6. Click the **IGMP** tab.
7. Select the **SnoopEnable** check box.
8. Select the **ProxySnoopEnable** check box.
9. For SteamLimtEnable, select **enable**.
10. Click **Apply**.

Configuring IGMP interface static members

Configure IGMP interface static members to add members to a snoop group.

About this task

You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams always forward to the multicast router within the VLAN, in addition to the ports in this static entry.

Important:

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Static** tab.
4. Click **Insert**.
5. Type the appropriate information.
6. Click **Insert**.

Static field descriptions

Use the data in the following table to use the **Static** tab.

Name	Description
IfIndex	Shows the interface where the IGMP entry is enabled.

Table continues…

Name	Description
GrpAddr	Indicates the start of the IP multicast address range of the multicast stream. Within the indicated valid range (224.0.0.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an invalid message.
ToGrpAddr	Indicates the end of the IP multicast address range of the multicast stream. If an address is not entered, the IP address in the GrpAddr field is the single address.
MemberPorts	Specifies the ports to which you want to redirect the multicast stream for this multicast group. The ports must be member ports of the VLAN.
NotAllowedToJoin	Specifies the ports that do not receive the multicast stream for this multicast group.

Configuring the SSM map table

Configure the SSM map table to map groups to their sending source. SSM maps cannot conflict with static source groups. After you configure an SSM map or a static source group, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) or multiple groups to different sources for both static source group and an SSM channel.

About this task

The consistency check applies to all SSM channel entries, even if they are disabled. If you disable an entry, it becomes inactive. If you do not delete the entry, you can reenable it later.

After you disable an SSM map, the Avaya Virtual Services Platform 9000 stops multicast traffic from the specified source to the specified group. You can use this static configuration as a security feature to block traffic from a certain source to a specific group.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Ssm Map** tab.
4. Click **Insert**.
5. Type the IP address for the multicast group and source.
6. Click **Insert**.

   You can change the default status of an SSM map from enable to disable by clicking in the AdminState field.

Ssm Map field descriptions

Use the data in the following table to use the **Ssm Map** tab.

Name	Description
IpMulticastGrp	Specifies an IP multicast address that is within the SSM range.
IpSource	Specifies the IP address of the source that sends traffic to the group.
LearningMode	Displays whether the entry is statically configured (Static) or dynamically-learned from IGMPv3 (Dynamic). This variable a read-only field.
Activity	Displays the current activity of the selected (S,G) entry. True indicates that traffic is flowing to the switch, otherwise, it appears false. This variable a read-only field.
AdminState	Configures the administrative state for the selected static entry. This state determines whether the switch uses the static entries. Configure this field to enable (default) to use the entry or disable to save for future use.

Configuring SSM range and global parameters

Configure the SSM range parameter to extend the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without changing their group configurations.

Before you begin

- To change the RangeGroup configuration, you must first disable PIM.
- To change the RangeGroup configuration, you must delete all entries in the SSM channel table before you configure the new IP multicast group address.

About this task

The other global parameters enable the IGMPv3 dynamic learning feature and configure the administrative state for all the entries in the SSM channel table.

Important:

If you change the RangeGroup configuration, the switch reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), this procedure also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap router (BSR) is local.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.

2. Click **IGMP**.
3. Click the **Ssm Global** tab.
4. Configure the appropriate fields.
5. Click **Apply**.

Ssm Global field descriptions

Use the data in the following table to use the **SsmGlobal** tab.

Name	Description
DynamicLearning	Activates the dynamic learning of SSM channel (S,G) pairs from IGMPv3 reports. As new SSM channels are learned, they appear in the SSM channel table.
RangeGroup	Configures the IP multicast group address. The lowest group address is 224.0.0.0 and the highest is 239.255.255.255. The default is 232.0.0.0.
RangeMask	Configures the address mask of the multicast group. The default is 255.0.0.0.
SsmMapAdminAction	Configures the administrative state, which determines whether the switch uses the table entries: • enableAll—Globally activates all the static entries in the SSM channel table. This value does not affect the dynamically learned entries. • disableAll—Globally inactivates all the static entries in the SSM channel table. This value does not affect the dynamically learned entries.

Configuring multicast stream limitation on an interface

Configure multicast stream limitation to limit the number of concurrent multicast streams on the interface. By limiting the number of concurrent multicast streams, you can protect the bandwidth on a specific interface and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the interface drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a specific limit of multicast streams on an interface.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.

3. Click the **StreamLimit** tab.
4. To change the status of an interface, double-click on the **StreamLimitEnable** field for the interface, and then select **enable** or **disable** from the menu. If the interface is enabled, you can edit the **Maximum Number of Stream** field.
5. Click **Apply**.

StreamLimit field descriptions

Use the data in the following tab to use the **StreamLimit** tab.

Name	Description
Interface	Displays the slot and port number or VLAN ID for this interface.
StreamLimitEnable	Enables or disables stream limitation on this interface.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this interface. The range is from 0–65535, and the default is 4.
Current Number Of Stream	Displays the current number of streams received on this interface. This value is a read-only value.

Configuring multicast stream limitation on a VLAN

Configure multicast stream limitation to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, you can protect the bandwidth on a specific VLAN and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the VLAN drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a specific limit of multicast streams on an interface.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Click the **IGMP** tab.
7. For StreamLimitEnable, select **enable**.

8. Configure the maximum number of streams.
9. Click **Apply**.

Configuring multicast stream limitation on a port

Configure multicast stream limitation to limit the number of concurrent multicast streams on the port. Limit the number of streams to protect the bandwidth on a specific port and control access to multicast streams.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **IGMP** tab.
5. In the StreamLimitEnable field, select the **Enable** option button.
6. Configure the maximum number of streams.
7. Click **Apply**.

Configuring multicast stream limitation members

Configure multicast stream limitation members on ports of the specified interface to configure the maximum number of streams on the interface.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit Members** tab.
4. Click **Insert**.
5. Type the number of the VLAN to which you want to add a member or click **Vlan** to select an ID from the list.
6. Type the number of the slot and port, as slot/port, that you want to add as a member or click **Port**, and then select one from the graphic display.

Important:

You must select one of the ports in the VLAN that you selected in step 4.

7. Type a maximum number of streams or accept the default of 4.
8. Click **Insert**.

StreamLimit Members field descriptions

Use the data in the following table to use the **StreamLimit Members** tab.

Name	Description
IfIndex	Displays the ID of the VLAN.
Port	Lists each slot and port number for this interface with stream limitation enabled.
MaxStreams	Configures the maximum number of allowed streams for this specific port. The number of allowed streams cannot exceed the maximum number for the interface. The range is from 0–65535 and the default is 4.
NumStreams	Displays the current number of streams received on this interface. This value is a read-only value.

Deleting multicast stream limitation member

Delete a multicast stream limitation member from an interface to remove it from the configuration.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit Members** tab.
4. Click on the row that lists the member you want to delete.
5. Click **Delete**.

Configuring the IGMP interface

Configure the IGMP interface to change global IGMP values for the interface. Use the Interface tab to view or edit the IGMP interface table.

About this task

If an interface does not use an IP address, it does not appear in the IGMP table. If an interface uses an IP address, but PIM-SM is not enabled, the interface appears as notInService in the Status field.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**
2. Click **IGMP**.
3. Click the **Interface** tab.
4. Edit the appropriate information.
5. Click **Apply**.

Interface field descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
IfIndex	Shows the interface where IGMP is enabled.
QueryInterval	Configures the frequency (in seconds) at which the IGMP host query packets transmit on the interface. The range is from 1–65535 and the default is 125.
Status	Shows the IGMP row status. If an interface uses an IP address and PIM-SM is enabled, the status is active. Otherwise, it is notInService.
Version	Configures the version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
OperVersion	Shows the version of IGMP that currently runs on this interface.
Querier	Shows the address of the IGMP querier on the IP subnet to which this interface attaches.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255, and the default is 100 tenths of a second (equal to 10 seconds.)

Table continues…

Name	Description
	Important: You must configure this value lower than the QueryInterval.
WrongVersionQueries	Shows the number of queries received with an IGMP version that does not match the interface. You must configure all routers on a LAN to run the same version of IGMP. If the interface receives queries with the wrong version, this value indicates a version mismatch.
Joins	Shows the number of times this interface added a group membership, which is the same as the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	Tunes for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, increase the robustness value. The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.
LastMembQueryIntvl	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of second. Avaya recommends that you configure this parameter to values greater than 3. If you do not need a fast leave process, Avaya recommends values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
OtherQuerierPresent Timeout	Shows the length of time that must pass before a multicast router determines that no other querier exists. If the local router is the querier, the value is 0.
FlushAction	Configures the flush action to one of the following: • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.

Table continues…

Name	Description
	Important: To maximize network performance, Avaya recommends that you configure this parameter according to the version of IGMP currently in use. • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SsmSnoopEnable	Enables SSM snoop.
SnoopQuerierEnable	Enables IGMP Layer 2 Querier.
SnoopQuerierAddr	Enables the IGMP Layer 2 Querier address.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts for each channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.
McastMode	Indicates the protocol configured on the VLAN. • snoop — Indicates IGMP snooping is enabled on a VLAN. • snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP Multicast over Fabric Connect for a Layer 2 VSN.) • routed-spb — Indicates IP Multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts. • pim — Indicates PIM is enabled.

Configuring IGMP sender entries

Configure IGMP sender entries to identify a source that sends multicast data to a multicast group.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Sender** tab.
4. Change the appropriate options.
5. Click **Apply**.

Sender field descriptions

Use the data in the following table to use the **Sender** tab.

Name	Description
GrpAddr	Specifies the multicast group address of the multicast stream. Within the indicated valid range (224.0.0.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an invalid message.
IfIndex	Specifies the interface where you enabled the IGMP entry.
MemberAddr	Specifies the IP address of a host.
Action	Flushes an entry or a group.
TPort	Identifies the T port.
State	Indicates whether a sender exists because of an IGMP access filter. The options are filtered and not filtered.

Configuring fast leave mode

Configure fast leave mode to control all IGMP fast leave enabled interfaces.

Before you begin

- You must enable explicit-host-tracking before configuring fast-leave mode. To enable explicit-host-tracking, see Configuring IGMP parameters on a port on page 74 and Configuring IGMP parameters on a VLAN on page 76.

About this task

Fast leave relies on an alternative leave process where the switch stops sending traffic for the group immediately after it receives a leave message, without issuing a query to check if other group members exist on the network. Use this global parameter to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Global** tab.

4. Select the mode.
5. Click **Apply**.

Global field descriptions

Use the data in the following table to use the **Global** tab.

Name	Description
FastLeaveMode	Configures the mode to one of the following values: • multipleUser: Removes from the group only the IGMP member who sent the leave message. Traffic does not stop if other receivers exist on the interface port. This value is the default. • oneUser: Removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.
GenerateTrap	Generates a trap. The default is disable.
GenerateLog	Generates a log message. The default is disable.

Configuring multicast access control for an interface

Configure multicast access control for a selected IGMP interface or VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Access Control** tab.
4. Click **Insert**.
5. Type the number of the slot and port or VLAN ID that you want to add as a member or click the appropriate button, and then select one from the graphic display.
6. Click the ellipsis button (**...**) next to **PrefixListId**.
7. Select a prefix list ID.
8. Click **OK**.

9. Type the host address and host mask.
10. Select the action mode that you want for the specified host.
11. Click **Insert**.

Access Control field descriptions

Use the data in the following table to use the **Access Control** tab.

Name	Description
IfIndex	Specifies the interface where the IGMP entry is enabled.
PrefixListId	Specifies a numeric string that identifies the prefix list.
HostAddr	Specifies the IP address of the host.
HostMask	Specifies the subnet mask that determines the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
PrefixListName	Specifies the name of the prefix list.
ActionMode	Specifies the action for the host identified by HostAddr. The options include the following: • denied IP multicast transmitted traffic (deny-tx). • denied IP multicast received traffic (deny-rx). • denied both IP multicast transmitted and received traffic (deny-both). • allowed IP multicast transmitted traffic (allow-only-tx). • allowed IP multicast received traffic (allow-only-rx). • allowed both IP multicast transmitted and received traffic (allow-only-both).

Viewing IGMP cache information

View IGMP cache information to view the group for which members exist on a specific interface.

About this task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**

3. Click the **Cache** tab.

Cache field descriptions

Use the data in the following table to use the **Cache** tab.

Name	Description
Address	Shows the IP multicast group address for this entry that contains this information.
IfIndex	Shows the interface from which the corresponding multicast group address is heard.
LastReporter	Shows the IP address of the source of the last membership report received for this IP multicast group address on this interface. If no membership report is received, the object uses the value 0.0.0.0.
ExpiryTime	Shows the amount of time (in seconds) that remain before this entry ages out.
Version1HostTimer	Shows the time that remains until the local router assumes that no IGMPv1 members exist on the IP subnet attached to the interface. Upon hearing IGMPv1 membership report, this value resets to the group membership timer. When the time that remains is nonzero, the local router ignores IGMPv2 leave messages for this group that it receives on this interface.
Type	Shows the type of IGMP entry.
StaticPorts	Shows the static ports associated with the entry.

Viewing IGMPv3 cache

View the IGMPv3 specific data corresponding to each interface, port, and multicast group pair on a router.

About this task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **IGMPv3 Cache** tab to view the IGMPv3 cache information.

IGMPv3 Cache field descriptions

Use the data in the following table to use the **IGMPv3 Cache** tab.

Name	Description
GroupAddress	Specifies the Multicast group Address (Class D) that others want to join. A group address can be the same for many incoming ports.
IfIndex	An unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports from various sources.
InPort	An unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports from various sources.
ModeExpiryTimer	Represents the time remaining before the interface EXCLUDE state expires and the interface state transitions to INCLUDE mode. This value is applicable only to IGMPv3-compatible nodes.
Version1HostTimer	Specifies the time remaining until the local router assumes that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. This entry only applies to IGMPv1 hosts. Upon hearing any IGMPv1 report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface.
Version2HostTimer	Specifies the time remaining until the local router assumes that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. Assuming no IGMPv1 hosts have been detected, the local router does not ignore any IGMPv2 Leave messages for this group that it receives on this interface.
SourceFilterMode	Specifies the current group state, applicable to IGMPv3-compatible nodes. The value indicates whether the state is INCLUDE or EXCLUDE.

Viewing and editing multicast router discovery information

View multicast router discovery information to view the current configuration.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.

3. Click the **Multicast Router Discovery** tab.
4. To edit the current configuration, double-click the value, make the change, and then click **Apply**.

Multicast Router Discovery field descriptions

Use the data in the following table to use the **Multicast Router Discovery** tab.

Name	Description
Interface	Shows the interface where IGMP is enabled.
MrdiscEnable	Enables (true) or disables (false) the router interface to listen for multicast router discovery messages to determine where to send multicast source data and IGMPv2 reports. If you enable snoop, you automatically enable multicast router discovery.
DiscoveredRouterPorts	Lists ports that the Multicast Router Discovery (MRDISC) protocol discovers. **Important:** The Avaya Virtual Services Platform 9000 does not support the MRDISC protocol on brouter ports.
MaxAdvertiseInterval	Shows the maximum time allowed between sending router advertisements from the interface, in seconds. The range is from 2–180 seconds. The default is 20 seconds.
MinAdvertiseInterval	Shows the minimum time allowed between sending unsolicited router advertisements from the interface, in seconds. This value must be more than 3 seconds but no greater than the value assigned to the MaxAdvertiseInterval value.
MaxInitialAdvertiseInterval	Configures the maximum number (in seconds) of multicast advertisement intervals that you can configure on the switch.
MaxInitialAdvertisements	Configures the maximum number of initial multicast advertisements that you can configure on the switch.
NeighborDeadInterval	Shows the time interval (in seconds) before the router interface drops traffic after a user leaves the multicast group.

Viewing the IGMP router source list

View the source list entries corresponding to each interface and multicast group pair on a router.

About this task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Igmp Router Source List** tab to view the IGMPv3 cache information.

Igmp router source list field descriptions

Use the data in the following table to use the **Igmp Router Source List** tab.

Name	Description
GroupAddress	Specifies the IP multicast group address for which this entry contains information.
IfIndex	Specifies the interface for which this entry contains information for an IP multicast group address.
InPort	Specifies a unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports for this source.
HostAddress	Specifies the host address to which this entry corresponds.
MemberAddress	Specifies the IP Address of a member that has sent source specific report wishing to join this source.
Expire	This value indicates the relevance of the source list entry, where a non-zero value indicates this is an INCLUDE state value, and a zero value indicates this to be an EXCLUDE state value.
Mode	Specifies the current member state, applicable to IGMPv3-compatible nodes. The value indicates whether the state is INCLUDE or EXCLUDE.
MemberExpire	This value indicates the time until the member for this source expires.

Viewing IGMP snoop information

View information about IGMP snoop to see the current configuration.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.

3. Click the **Snoop** tab.

Snoop field descriptions

Use the data in the following table to use the **Snoop** tab.

Name	Description
Interface	Shows the VLAN ID for the VLAN.
SnoopEnable	Shows the status of IGMP snoop. IGMP snoop works only if a multicast router exists in the VLAN.
SsmSnoopEnable	Shows the status of SSM snoop.
ProxySnoopEnable	Indicates whether the IGMP report proxy feature is enabled. If you enable this feature, the switch forwards reports from hosts to the multicast router once for each group for each query interval, or after new group information is available. If you disable this feature, the switch forwards all reports from different hosts to multicast routers, and can forward more than one group report for the same multicast group for each query interval. The default is enabled.
FastLeaveEnable	Shows the status of fast leave for this port.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router. **Important:** Configure this variable only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.
SnoopActiveMRouterPorts	Shows the active multicast router ports. Active multicast router ports are ports that directly attach to a multicast router. These ports include the querier port and all ports in the forwarding state that you configure as well as those that were dynamically learned through receiving queries.
SnoopMRouterExpiration	Indicates the time that remains before the multicast router ages out. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The query maximum response interval (obtained from the queries received) is used as the timer resolution.

Viewing IGMP group information

View information about IGMP groups to see the current group operation on the switch.

About this task

Note:

The following procedure displays the dynamically learned IGMP groups. **IP** > **IGMP** > **Static** displays statically configured IGMP groups. This is in contrast to the ACLI command **`show ip igmp group`**, which displays both dynamically learned and statically configured IGMP groups, and the ACLI command **`show ip igmp static`**, which displays only the statically configured groups.

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Groups** tab.

Groups field descriptions

Use the data in the following table to use the **Groups** tab.

Name	Description
IpAddress	Shows the multicast group address (Class D). A group address can be the same for many incoming ports.
Members	Shows the IP address of the host that issues the membership report to this group.
InPort	Shows the port that receives the group membership report.
IfIndex	Shows a unique value that identifies a physical interface or a logical interface (VLAN) that receives the membership report.
Expiration	Shows the time left before the group report expires on this port. This variable is updated after the port receives a group report.

Chapter 10: Route management using ACLI

With multicast route commands, you can configure and view IP multicast routing parameters on the Avaya Virtual Services Platform 9000.

Configuring multicast stream limits

Limit the number of multicast streams to protect a Control Processor (CP) module from multicast data packet bursts generated by malicious applications, such as viruses that cause the CP module to reach 100 percent utilization or that prevent the CP module from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CP module through a port during a sampling interval, the port shuts down until you take appropriate action.

About this task

You can enable or disable the mroute stream limit for the entire device or for individual ports when the switch is operating. If you enable the mroute stream limit for the device and for an individual port, only the periodic check is performed for that port.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable stream limitation globally:

   ```
   ip mroute stream-limit
   ```

3. Log on to the GigabitEthernet Interface Configuration mode.

   ```
   interface gigabitethernet {slot/port[-slot/port][,...]}
   ```

4. Enable stream limits:

   ```
   ip mroute stream-limit
   ```

5. For Gigabit Ethernet interfaces, configure the maximum number of streams and the interval at which to sample:

   ```
   ip mroute max-allowed-streams <1-32768> max-allowed-streams-timer-check <1-3600>
   ```

6. Show the mroute stream limit configuration:

   ```
   show ip mroute interface gigabitethernet [{slot/port[-slot/port]
   [,...]}]
   ```

Example

```
VSP-switch:1(config)#ip mroute stream-limit
VSP-switch:1(config)#interface gigabitethernet 4/11
VSP-switch:1(config-if)#ip mroute stream-limit
VSP-switch:1(config-if)#ip mroute max-allowed streams 1000 max-allowed-streams-timer-
check 20
VSP-switch:1(config)#show ip mroute interface gigabitethernet 4/11

================================================================================
                              Port Mroute-Limit
================================================================================
PORT  MROUTE STR    MROUTE STR    ENABLE
      LIMIT         LIMIT TMR
--------------------------------------------------------------------------------
4/11  1000          20            true
```

Variable definitions

Use the data in the following table to use the **`interface`** command.

Variable	Value
<1–4084>	Specifies the VLAN ID to enter VLAN Interface Configuration mode.
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Use the data in the following table to use the **`ip mroute`** command.

Variable	Value
max-allowed-streams *<1–32768>*	Configures the maximum number of streams on the specified port. The port is shut down if the number of streams exceeds this limit. The value is a number between 1–32768. The default value is 1984 streams. To configure this option to the default value, use the default operator with the command.
max-allowed-streams-timer-check *<1–3600>*	Configures the sampling interval, which checks if the number of ingress multicast streams to the CP module is under a configured limit or if the port needs to shut down. The range is between 1–3600. The default value is 10 seconds. To configure this option to the default value, use the default operator with the command.

Job aid

The following message appears if the system shuts down the port due to excessive multicast streams:

```
Shutdown port <port> due to excessive multicast streams <# of streams
ingressed>; Configured limit max streams <configured limit> in
<configured sampling interval> sec. Please disable and re-enable the
port.
```

The following table shows the field descriptions for the **`show ip mroute interface`** command.

Table 8: show ip mroute interface field descriptions

Field	Description
PORT	Indicates the slot and port number.
MROUTE STR LIMIT	Indicates the maximum number of multicast streams that can enter the CP module through this port.
MROUTE STR LIMIT TIMER	Indicates the sampling period (in seconds) to check the number of multicast streams that enter the CP module through this port.
ENABLE	Indicates the status of the mroute stream limit on the port.

Configuring multicast static source groups

Configure static source group entries in the Protocol Independent Multicast (PIM) multicast routing table. The PIM cannot prune these entries from the distribution tree.

Before you begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - PIM-Sparse Mode (SM)
  - PIM-Source Specific Multicast (SSM)

About this task

Even if no receivers exist in the group, the multicast stream for a static source group entry remains active.

The maximum number of static source groups must not exceed 1024.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a static source group entry:

   ```
   ip mroute static-source-group <A.B.C.D> <A.B.C.D/X>
   ```

Example

Create a static source group for two multicast groups: 224.32.2.1 and 226.50.2.2. The static mroute for group 224.32.2.1 is for a source subnet 10.10.10.0/24. The static mroute for group 226.50.2.2 is for the host 20.20.20.100/32.

```
VSP-9012:1(config)# ip mroute static-source-group 224.32.2.1
10.10.10.0/24
VSP-9012:1(config)# ip mroute static-source-group 226.50.2.2
20.20.20.100/32
```

Variable definitions

Use the data in the following table to use the **`ip mroute static-source-group`** command.

Variable	Value
A.B.C.D	Specifies the IP address of the multicast group. Use the no operator to later remove this configuration.
A.B.C.D/X	Specifies the multicast source IP address and subnet mask for the static source group entry. You cannot create duplicate groups. How you configure the source address depends on the protocol and mode you use. Use the no operator to later remove this configuration.

Configuring IP multicast software forwarding

When you use the IP multicast software forwarding feature you can avoid initial data loss experienced by multicast applications; this is suitable for low bandwidth conditions.

When you configure the IP multicast software forwarding feature the system forwards the initial packets of an IP multicast data stream it receives and creates a corresponding hardware record for subsequent packets.

By default, multicast software forwarding is disabled.

About this task

IP multicast software forwarding is a global system configuration feature that applies to all IP multicast-enabled interfaces and protocols. If you enable IP multicast software forwarding, the hardware continues to forward IP multicast traffic. The software only forwards initial data traffic.

After a new data stream arrives, the first data packet is sent to the CP which programs the multicast route in hardware, and all packets that arrive subsequent to this programming are forwarded by hardware only.

If you enable software forwarding, all initial packets received before hardware programming is complete are sent to the CP for forwarding and packet suppression by the hardware is disabled.

If you do not enable software forwarding, only the first data packet is sent to the CP and subsequent packets are suppressed by the hardware so that the CP is not overwhelmed with traffic. During this time, packets suppressed by the hardware are dropped.

Important:

To avoid overloading the CP module, Avaya recommends that you do not use the IP multicast software forwarding feature for video multicast applications.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Enable software forwarding:

   ```
   multicast software-forwarding
   ```

3. Show the software forwarding configuration:

   ```
   show multicast software-forwarding
   ```

Example

```
VSP-9012:1#show multicast software-forwarding

================================================================================
                    Mcast Software Forwarding - GlobalRouter
================================================================================
McastSoftwareForwarding        :disabled
```

Configuring the resource usage counter for multicast streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing the switch.

About this task

After you configure the counter thresholds for ingress and egress records, if the record usage exceeds the threshold, you receive notification by a trap on the console, a logged message, or both.

If you do not configure the thresholds, ACLI displays only the ingress and egress records currently in use.

You can configure the resource usage counter on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the thresholds:

   ```
   ip mroute resource-usage egress-threshold <0–32767> ingress-threshold <0–32767>
   ```

3. Configure one of the following notification methods:
   - Configure a log-only notification method:

     ```
     ip mroute resource-usage log-msg
     ```

   - Configure a trap-only notification method:

     ```
     ip mroute resource-usage trap-msg
     ```

   - Configure both notification methods:

     ```
     ip mroute resource-usage log-msg trap-msg
     ```

Example

Configure the egress threshold to 200.

```
VSP-9012:1(config)# ip mroute resource-usage egress-threshold 200
```

Configure the ingress threshold to 100.

```
VSP-9012:1(config)# ip mroute resource-usage ingress-threshold 100
```

Enable the log message notification method.

```
VSP-9012:1(config)# ip mroute resource-usage log-msg
```

Variable definitions

Use the data in the following table to use the **`ip mroute resource-usage`** command.

Variable	Value
egress-threshold *<0–32767>*	Configures the egress record threshold (S,G). The system sends a notification message after the number of streams exceeds a threshold level. To configure this option to the default value, use the default operator with the command. The default is 0.
ingress-threshold *<0–32767>*	Configures the ingress record threshold. The system sends a notification message after the number of streams exceeds a threshold level.

Table continues…

Variable	Value
	To configure this option to the default value, use the default operator with the command. The default is 0.

Configuring prefix lists

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

About this task

Important:

When you configure a prefix list for a route policy, add the prefix as a.b.c.d/32. You must enter the full 32-bit mask to exact a full match of a specific IP address.

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a prefix list:

   ```
   ip prefix-list WORD<1-64> {A.B.C.D/X} [ge <0-32>] [le <0-32>]
   ```

3. **(Optional)** Rename an existing prefix list:

   ```
   ip prefix-list WORD<1-64> name WORD<1-64>
   ```

4. Display the prefix list:

   ```
   show ip prefix-list [prefix {A.B.C.D}] [vrf WORD<1-16>] [vrfids
   WORD<0-512>] [WORD <1-64>]
   ```

Example

Configure a prefix-list. Display the prefix list.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#ip prefix-list LIST1 47.17.121.50/255.255.255.0
VSP-9012:1(config)#show ip prefix-list LIST1
================================================================================
                          Prefix List - GlobalRouter
================================================================================

       PREFIX           MASKLEN FROM TO
```

```
--------------------------------------------------------------------------------

List 1    LIST1:
        47.17.121.50     24      24   24
1 Total Prefix List entries configured
--------------------------------------------------------------------------------
Name Appendix for Lists Converted from Old Config:
@A=conv addr list, @N=conv net list, @NR=conv net list modified as range
```

Variable definitions

Use the data in the following table to use the **ip prefix-list** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and the mask in one of the following formats: • a.b.c.d/x • a.b.c.d/x.x.x.x • default
ge <0–32>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
le <0–32>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
name WORD<1-64>	Renames the specified prefix list. The name length is 1–64 characters.
WORD<1-64>	Specifies the name for a new prefix list.

Use the data in the following table to use the **show ip prefix-list** command.

Variable	Value
{A.B.C.D}	Specifies the prefix to include in the command output.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

Use the following table to use the **show ip prefix-list** command output.

Variable	Value
PREFIX	Indicates the member of a specific prefix list.

Table continues…

Variable	Value
MASKLEN	Indicates the prefix mask length in bits.
FROM	Indicates the prefix mask starting point in bits.
TO	Indicates the prefix mask endpoint in bits.

Chapter 11: Route management using EDM

View or edit interface configuration information for Layer 3 IP multicast protocols on the switch.

Viewing multicast route information

View multicast route information for troubleshooting purposes.

This tab shows multicast routing information for IP datagrams from a particular source and addressed to a particular IP multicast group address.

About this task

You can view the multicast routes for a Layer 3 Virtual Services Network (VSN) the same way you view the Global Router except that you must first launch the appropriate VRF context. For more information about Layer 3 VSNs, see *Configuring Avaya Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Routes** tab.

Routes field descriptions

Use the data in the following table to use the **Routes** tab.

Name	Description
Group	Displays the IP multicast group address for this entry that contains multicast routing information.
Source	Displays the network address that, when combined with the corresponding route SourceMask value, identifies the source that contains multicast routing information.
SourceMask	Displays the network mask that, when combined with the corresponding route Source value, identifies the multicast source.

Table continues…

Name	Description
UpstreamNeighbor	Shows the address of the upstream neighbor from which the IP datagrams from these sources are received. The address is 0.0.0.0 if the network is local.
Interface	Displays the interface, slot and portnumber, or VLAN ID where IP datagrams sent by these multicast sourcesto this multicast address are received.
ExpiryTime	Displays the amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11)

Viewing multicast next-hop information

View all multicast next-hop information.

This tab shows information about the next hops used by outgoing interfaces to route IP multicast datagrams. Each entry is one in a list of next hops on outgoing interfaces for particular sources that send to a particular multicast group address.

About this task

You can view multicast next-hop information for a Layer 3 VSN the same way you view the Global Router except that you must first launch the appropriate VRF context. For more information about Layer 3 VSNs, see *Configuring Avaya Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Next Hops** tab.

Next Hops field descriptions

Use the data in the following table to use the **Next Hops** tab.

Name	Description
Group	Displays the IP multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
OutInterface	Displays the interface slot and portnumber or VLAN ID for the outgoing interface for this next hop.
Address	Displays the address of the next hop specific to this entry. For most interfaces, it is identical to the next-hop group. Non Broadcast Multiple Access (NBMA) interfaces, however, can use multiple next hop addresses out of a single outgoing interface.
State	Displays whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. A value of forwarding indicates the information is currently used; pruned indicates it is not used.
ExpiryTime	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
ClosestMemberHops	Displays the minimum number of hops between this router and members of the IP multicast group reached through the next hop on this outgoing interface. IP multicast datagrams for the group that use a time-to-live less than this number of hops are not forwarded to the next hop.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11)

Viewing multicast interface information

View multicast interface information to verify the multicast configuration.

This tab shows multicast routing information specific to interfaces.

About this task

You can view multicast interface information for a Layer 3 VSN the same way you view the Global Router except that you must first launch the appropriate VRF context. For more information about

Layer 3 VSNs, see *Configuring Avaya Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Interfaces** tab.

Interfaces field descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Interface	Displays the slot and port number or VLAN ID for this entry.
Ttl	Displays the datagram time-to-live (TTL) threshold for the interface. IP multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means that all multicast packets are forwarded out of the interface.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11)

Adding new static source groups

Add a new static source group to create an entry that the switch cannot prune from the distribution tree. An attempt to add a duplicate of an existing source-group entry results in an error message.

Before you begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - PIM-SM
  - PIM-SSM

About this task

Virtual Services Platform 9000 supports PIM only in the Global Router. You cannot configure static source groups for specific VRF contexts.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Static Source Group** tab.
4. Click **Insert**.
5. Complete the information in the dialog box.
6. Click **Insert**.

Editing static source groups

Configure static source-group entries in the PIM multicast routing table. PIM cannot prune these entries from the distribution tree. In other words, even if no receivers exist in the group, the multicast stream for a static source-group entry stays active.

Before you begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - PIM-Sparse Mode (SM)
  - PIM-Source Specific Multicast (SSM)

About this task

The maximum number of static source groups must not exceed 1024.

Virtual Services Platform 9000 supports PIM only in the Global Router. You cannot configure static source groups for specific VRF contexts.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Static Source Group** tab.
4. Edit the required information.
5. Click **Apply**.

Static Source Group field descriptions

Use the data in the following table to use the **Static Source Group** tab.

Name	Description
GroupAddress	Configures the multicast group IP address for this static source-group entry.
SourceSubnet	Configures the multicast source address for this static source-group entry. How you configure the source address depends on the protocol and mode you use.
SrcSubnetMask	Configures the subnet mask of the source for this static source-group entry.

Configuring IP multicast software forwarding

Configure IP multicast software forwarding to enable the system to initially forward IP multicast data until a hardware record is created. The system forwards the initial packets of a stream it receives and creates a corresponding hardware record for subsequent packets. The advantage of this feature is that it avoids initial data loss experienced by multicast applications and is most suited for low bandwidth.

About this task

The IP multicast software forwarding is a global system configuration feature that applies to all IP multicast-enabled interfaces and protocols. After you enable IP multicast software forwarding, the hardware still forwards IP multicast traffic. The software forwards only initial data traffic.

After a new data stream arrives, the first data packet is sent to the CP which programs the multicast route in hardware, and all packets that arrive subsequent to this programming are forwarded by hardware only. If you enable software forwarding, all initial packets received before hardware programming is complete are sent to the CP for forwarding. If you enable software forwarding, packet suppression by the hardware is disabled. If you do not enable software forwarding, only the first data packet is sent to the CP and subsequent packets are suppressed by the hardware so that the CP is not overwhelmed with traffic. During this time, packets suppressed by the hardware are dropped.

By default, the feature is disabled.

Important:

To avoid overloading the CPU, do not use the IP multicast software forwarding feature for video multicast applications.

If you configure multicast software forwarding from within a VRF context, the configuration applies to the Global Router and all VRF contexts. You cannot change the multicast software forwarding configuration for individual VRF contexts.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Globals** tab.
4. Select the **SWForwardingEnable** check box.

5. Click **Apply**.

Globals field descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
MRouteStatsEnabled	Enables the collection of multicast route statistics. The default is disabled.
SWForwardingEnable	Enables the system to initially forward IP multicast data until a hardware record is created. The default is disabled.
MulticastSquareSmltEnable	Enables square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. The default is disabled.

Configuring mroute stream limit

Limit the number of multicast streams to protect a CPU from multicast data packet bursts generated by malicious applications, such as viruses that cause the CPU to reach 100 percent utilization or that prevent the CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **General**.
4. Select the **Mroute Stream Limit** tab.
5. Select the **StreamLimitEnable** box.
6. Edit other fields as required.
7. Click **Apply**.

Mroute Stream Limit field descriptions

Use the data in the following table to use the **Mroute Stream Limit** tab.

Name	Description
StreamLimitEnable	Enables or disables mroute stream limit on the port.
StreamLimit	Specifies the maximum number of multicast streams allowed to enter the CPU through this port.
StreamTimerCheck	Specifies the sampling period, in seconds, to check the number of multicast streams that enter the CPU through this port.

Configuring resource usage counter for multicast streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing the switch. After you configure the counter thresholds for ingress and egress records, if the record usage goes beyond the threshold, you receive notification through a trap on the console, a logged message, or both.

About this task

Important:

If you do not configure the thresholds, EDM displays only the ingress and egress records that are currently in use.

You can configure the resource usage counter on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Select the **Resource Usage** tab.
4. Configure the ingress and egress thresholds.
5. Configure the notification methods.
6. Click **Apply**.

Resource Usage field descriptions

Use the data in the following table to use the **Resource Usage** tab.

Name	Description
Egress Records In-Use	Displays the number of egress records traversing the switch.
Ingress Records In-Use	Displays the number of ingress records (source or group) traversing the switch.

Table continues…

Name	Description
Egress Threshold	Configures the egress threshold level (0–32767).
Ingress Threshold	Configures the ingress threshold level (0–32767).
SendTrapOnly	Sends only trap notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type. You can configure only one notification type.
SendTrapAndLog	Sends both trap and log notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type.
LogMsgOnly	Sends only log notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type.

Configuring a prefix list

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

Before you begin

- Change the VRF instance as required to configure a prefix list on a specific VRF instance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Policy**.
3. Click the **Prefix List** tab.
4. Click **Insert**.
5. In the **Id** box, type an ID for the prefix list.
6. In the **Prefix** box, type an IP address for the route.
7. In the **PrefixMaskLength** box, type the length of the prefix mask.
8. Configure the remaining parameters as required.
9. Click **Insert**.

Prefix List field descriptions

Use the data in the following table to use the **Prefix List** tab.

Name	Description
Id	Configures the list identifier.
Prefix	Configures the IP address of the route.
PrefixMaskLen	Configures the specified length of the prefix mask. You must enter the full 32-bit mask to exact a full match of a specific IP address, for example, if you create a policy to match on the next hop.
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name length can use from 1 to 64 characters.
MaskLenFrom	Configures the lower bound of the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Configures the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Chapter 12: Multicast MAC filtering using ACLI

With multicast media access control (MAC) filtering, you can create a smaller flooding domain inside a VLAN. You can specify a multicast MAC address and a subset of ports for a VLAN. When clients send data to that designated MAC address, only that subset of ports receive the traffic.

Configuring Layer 2 multicast MAC filtering

Configure Layer 2 multicast MAC filtering to direct MAC multicast flooding to a specific set of ports.

About this task

Important:

Avaya recommends that you do not use MAC addresses beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive) with the MAC address parameter.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure Layer 2 multicast MAC filtering:

   ```
   vlan mac-address-static <1-4084> <0x00:0x00:0x00:0x00:0x00:0x00>
   {slot/port[-slot/port][,...]}
   ```

Example

Add a multicast MAC address 01:02:03:04:05:06 as a static MAC in VLAN 2. Add ports so that traffic destined for the MAC address forwards to ports 4/1 through 4/4, instead of flooding to all VLAN 2 ports.

```
VSP-9012:1(config)# vlan mac-address-static 2 01:02:03:04:05:06 4/1-4/4
```

Variable definitions

Use the data in the following table to use the **vlan mac-address-static** command.

Variable	Value
0x00:0x00:0x00:0x00:0x00:0x00	Specifies the MAC address in hexadecimal format.
1–4084	Specifies a VLAN from 1–4084.
{slot/port[-slot/port][,...]}	Specifies the port or ports that receive the multicast flooding. Type the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Configuring Layer 3 multicast MAC filtering

Configure Layer 3 multicast MAC filtering to route an IP frame to a unicast IP address and flood it with a destination multicast MAC address. You must manually define a static ARP entry that associates an IP address with a multicast MAC address, flooding ports, and a multilink trunk.

About this task

Important:

Avaya recommends that you do not use MAC addresses beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive) with the MAC address parameter.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Configure Layer 3 multicast MAC filtering:

   ```
   ip arp static-mcast {A.B.C.D} <0x00:0x00:0x00:0x00:0x00:0x00> vid
   <1-4084> [mlt WORD<1-16>][port {slot/port[-slot/port][,...]}]
   [WORD<1-16>]
   ```

Example

Add a multicast MAC address 01:01:01:01:01:02 as a static ARP entry in VLAN 2. Add ports and a multilink trunk group so that traffic destined for the MAC address forwards to ports 4/14 and 4/43, and MLT 1, instead of flooding to all VLAN 2 ports.

```
VSP-9012:1(config)# ip arp static-mcast 2.2.2.100 01:01:01:01:01:02 vid 2
port 4/14-4/43 1
```

Variable definitions

Use the data in the following table to use the **`ip arp static-mcast`** command.

Variable	Value
<0x00:0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the MAC address in hexadecimal format.
vid*<1–4084>*	Specifies the VLAN ID.
{A.B.C.D}	Specifies the IP address.
mlt *WORD<1–16>*	Specifies the multilink trunk ID.
port*{slot/port[-slot/port][,...]}*	Specifies the port that receives the multicast flooding. Type the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).
WORD<1–16>	Specifies the multilink trunk ID.

Chapter 13: Multicast MAC filtering using EDM

With multicast media access control (MAC) filtering, you can create a smaller flooding domain inside a VLAN. You can specify a multicast MAC address and a subset of ports for a VLAN. When clients send data to a designated MAC address, only that subset of ports receives the traffic.

Important:

To configure and use Multicast MAC filtering on a VRF instance, you must first select and launch the VRF context.

To select and launch the VRF context, see Selecting and launching a VRF context view on page 68.

Configuring Layer 2 multicast MAC filtering

Configure Layer 2 multicast MAC filtering to direct MAC multicast flooding to a specific set of ports.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. From the table, select a VLAN.
5. Click **Bridge**.
6. Click the **Multicast** tab.
7. Click **Insert**.
8. In the **Address** box, type the MAC address for the multicast flooding domain.
9. Click the ellipsis button (**...**) next to the **ForwardingPorts** box, and then choose from the list of ports that appear.
10. Click **Ok**.
11. Click the ellipsis button (**...**) next to the **MltIds** box, and then choose from the list of MLT IDs that appear.

12. Click **Ok**.
13. Click **Insert**.

Multicast field descriptions

Use the data in the following table to use the **Multicast** tab.

Name	Description
VlanId	Specifies the VLAN ID.
Address	Configures the MAC address for the multicast flooding domain.This variable does not accept MAC addresses beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive). MAC addresses in this range are reserved.
ForwardingPorts	Specifies the ports to include in the multicast flooding domain.
MltIds	Specifies the multilink trunks to include in the multicast flooding domain.
NumMltIds	Specifies the number of MLT IDs.

Configuring Layer 3 multicast MAC filtering

About this task

Configure Layer 3 multicast MAC filtering to route an IP frame to a unicast IP address and flood it with a destination multicast MAC address. You must manually define a static Address Resolution Protocol (ARP) entry that associates an IP address with a multicast MAC address, flooding ports, and a multilink trunk.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Multicast ARP** tab.
4. Click **Insert**.
5. Click the ellipsis button (**...**) next to the **VlanId** box, and then choose from the list that appears.
6. Click **Ok**.
7. In the **MacAddress** box, type the MAC address.
8. In the **IpAddress** box, type the IP address.
9. Click the ellipsis button (**...**) next to the **Ports** box, and then choose from the list of ports that appear.
10. Click **Ok**.

11. Click the ellipsis button (**...**) next to the **MltiIds** box, and then choose from the list of ports that appear.
12. Click **Ok**.
13. Click **Insert**.

Multicast ARP field descriptions

Use the data in the following table to use the **Multicast ARP** tab.

Name	Description
VlanId	Specifies the ID number of the VLAN for the multicast ARP.
MacAddress	Configures the MAC address assigned to the IP address in the multicast ARP entry. This field does not accept MAC addresses beginning with 01:00:5e (01:00:5e: 00:00:00 to 01:00:5e:ff:ff:ff inclusive).
IPAddress	Configures the IP address of the multicast ARP.
Ports	Specifies the ports that receive the multicast flooding.
MltIds	Specifies the multilink trunks to include in the multicast flooding domain.

Chapter 14: ACLI show command reference

This reference information provides show commands to view the operational status of multicast routing on the Avaya Virtual Services Platform 9000.

General show commands

This section explains the show commands for general multicast routing operations.

Layer 2 multicast MAC filters

Use the **show vlan static-mcastmac** command to display the Layer 2 multicast media access control (MAC) filters. If you specify a VLAN ID, the command displays information for the specified VLAN. Without the VLAN ID, the command displays information for all configured VLANs. The syntax for this command is as follows.

```
show vlan static-mcastmac [<1–4084>]
```

The following table shows the field descriptions for this command.

Table 9: show vlan static-mcastmac command

Field	Description
VLAN ID	Indicates the VLAN ID.
MAC ADDRESS	Indicates the MAC address.
PORT LIST	Indicates the list of ports.
MLT GROUPS	Indicates the MultiLink Trunking (MLT) groups.

Layer 3 multicast MAC ARP data

Use the **show ip arp static-mcastmac** command to display Layer 3 multicast MAC ARP data. You can specify optional information to narrow the results to a specific virtual router and forwarder (VRF) name or ID, or to a specific network and subnet. The valid syntax for this command is as follows.

```
show ip arp static-mcastmac -s {A.B.C.D/X}
```

show ip arp static-mcastmac vrf WORD<1-16>

show ip arp static-mcastmac vrfids WORD<0-512>

show ip arp static-mcastmac {A.B.C.D} -s {A.B.C.D/X}

show ip arp static-mcastmac {A.B.C.D}

show ip arp static-mcastmac

The following table shows the field descriptions for this command.

Table 10: show ip arp static-mcastmac command

Field	Description
IP_ADDRESS	Indicates the multicast IP address
MAC ADDRESS	Indicates the multicast MAC address.
VLAN	Indicates the VLAN ID.
PORT	Indicates the list of ports.
MLT ID	Indicates the multilink trunk ID.

Multicast route information

Use the **show ip mroute route** command to display information about the multicast routes on the switch. The syntax for this command is as follows.

show ip mroute route [vrf WORD <1–32>] [vrfids <0–255>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following section shows sample output for the **show ip mroute route** command.

In this table, every stream uses one (*,G) entry and x (S,G) entries, depending on how many servers forward traffic to the same group.

The 0.0.0.0 mask is always tied to a (*,G) entry.

Every time a new stream comes in, Protocol Independent Multicast (PIM) creates two entries in the table; one is a (*,G) entry that points toward the rendezvous point (RP) router, and the other is an (S,G) entry that points toward the source.

```
VSP-9012:1#show ip mroute route

================================================================================
                        Mroute Route - GlobalRouter
================================================================================
GROUP           SOURCE          SRCMASK         UPSTREAM_NBR    IF     EXPIR  PROT
--------------------------------------------------------------------------------
233.252.0.1       0.0.0.0         0.0.0.0         0.0.0.0        V3       30 spb-access
233.252.0.1       198.51.100.99   255.255.255.0   0.0.0.0        -        0  spb-network
233.252.0.2       0.0.0.0         0.0.0.0         0.0.0.0         V2      30  pimsm
```

```
233.252.0.2       198.51.100.99   255.255.255.0   0.0.0.0         V3      151 pimsm

Total 4
```

The following table shows the field descriptions for this command.

Table 11: show ip mroute route command

Field	Description
GROUP	Indicates the IP multicast group for this multicast route.
SOURCE	Indicates the network address that, when combined with the corresponding value of SRCMASK, identifies the sources for this multicast route.
SRCMASK	Indicates the network mask that, when combined with the corresponding value of SOURCE, identifies the sources for this multicast route.
UPSTREAM_NBR	Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the (S,G) source is local or if the RP for this the (*,G) group is an address on this router.
IF	Indicates the value of ifIndex for the interface that receives IP datagrams sent by these sources to this multicast address. A value of 0 in a (*,G) route indicates that datagrams are not subject to an incoming interface check, but datagrams can be received on any interface.
EXPIR	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
PROT	Indicates the multicast protocol through which the switch learned this route. The spb-access and spb-network indicate the stream learned when IP Multicast over Fabric Connect is configured on the VLAN. spb-access indicates that it was learned on the access. spb-network indicates it was learned over the SPBM cloud.

Multicast route next hop

Use the **show ip mroute next-hop** command to show information about the next hop for the multicast routes on the switch. The syntax for this command is as follows.

show ip mroute next-hop [vrf WORD <0-32>] [vrfids <0-255>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 12: show ip mroute next-hop command

Field	Description
INTERFACE	Indicates the interface identity.

Table continues…

Field	Description
GROUP	Indicates the IP multicast group for which this entry specifies a next-hop PIM neighbor toward receivers for a specific outgoing interface.
SOURCE	Indicates the network address, which when combined with the corresponding value of SRCMASK, identifies the sources for which this entry specifies a next-hop PIM neighbor toward receivers for a specific outgoing interface.
SRCMASK	Indicates the network mask, which when combined with the corresponding value of SOURCE, identifies the sources for which this entry specifies a next-hop PIM neighbor toward receivers for a specific outgoing interface.
ADDRESS	Indicates the address of the next hop specific to this entry. The next hop must be the address of a PIM neighbor. This table does not represent local receivers.
STATE	Indicates whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. The value forwarding indicates the information is currently used; the value pruned indicates it is not used.
EXPTIME	Indicates the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
CLOSEHOP	Indicates the minimum number of hops between this router and members of this IP multicast group reached through this next hop on this outgoing interface. IP multicast datagrams for the group that use a time-to-live less than this number of hops are forwarded to the next hop
PROTOCOL	Indicates the routing mechanism through which the switch learned this next hop.

Multicast routes on an interface

Use the **show ip mroute interface** command to display information about the multicast routes on the switch for a specific interface. The syntax for this command is as follows.

```
show ip mroute interface gigabitethernet [{slot/port[-slot/port][,...]}]
```

```
show ip mroute interface [vrf WORD <1-32>][vrfids WORD <0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command if you do not use optional command parameters.

Table 13: show ip mroute interface command without parameters

Field	Description
INTERFACE	Indicates the interface.

Table continues…

Field	Description
TTL	Indicates the datagram TTL threshold for the interface. IP multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means all multicast packets are forwarded out of the interface.
PROTOCOL	Indicates the routing protocol running on this interface.

The following table shows the field descriptions for this command if you use optional command parameters.

Table 14: show ip mroute interface command with parameters

Field	Description
PORT	Shows the slot and port location.
MROUTE STR LIMIT	Indicates the maximum number of multicast streams that can enter the SF module or CP module through this port.
MROUTE STR LIMIT TMR	Indicates the sampling period (in seconds) to check number of multicast streams that use ingressed the SF module or CP module through this port.
ENABLE	Indicates the status of the mroute stream limit on the port.

Multicast hardware resource usage

Use the **`show ip mroute hw-resource-usage`** command to display information about the hardware resource use of an IP multicast route.

The syntax for this command is as follows:

`show ip mroute hw-resource-usage [vrf WORD <1-32>] [vrfids WORD <0-255>]`

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

Field	Description
EGRESS REC IN-USE	Displays the number of egress records traversing the switch.
INGRESS REC IN-USE	Displays the number of ingress records (source or group) traversing the switch.
EGRESS THRESHOLD	Displays the configured egress threshold level (0–32767). A notification message is sent if this value is exceeded. The default is 0.

Table continues…

Field	Description
INGRESS THRESHOLD	Displays the configured ingress threshold level (0–32767). A notification message is sent if this value is exceeded. The default is 0.
LOG MSG ONLY	Displays whether only log notification messages are sent after the threshold level is exceeded. The default is false (disabled).
SEND TRAP ONLY	Displays whether only trap notification messages are sent after the threshold level is exceeded. The default is false (disabled).
SENT TRAP AND LOG	Displays whether both trap and log notification messages are sent after the threshold level is exceeded. The default is false (disabled).

Static source groups

Use the **show ip mroute static-source-group** command to display information about the static source groups. You can see all the valid entries that were created. If an entry is created with a x bit mask, it shows as a x bit in the output. The syntax for this command is as follows.

show ip mroute static-source-group [<A.B.C.D>][vrf WORD <1–32>][vrfids WORD <0–255>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 15: show ip mroute static-source-group command

Field	Description
Group Address	Indicates the IP multicast group address.
Source Address	Indicates the network address.
Subnet Mask	Indicates the network mask.

VLAN port data

Use the **show vlan members** command to display VLAN port data. The syntax for this command is as follows.

```
show vlan members [<1–4084>][null—vlan] [port {slot/port[-slot/port]
[,...]} ]
```

The following table shows the field descriptions for this command.

Table 16: show vlan members command

Field	Description
VLAN ID	Indicates the VLAN ID.
PORT MEMBER	Indicates the set of ports that are members (static or dynamic) of this VLAN.
ACTIVE MEMBER	Indicates the set of ports that are currently active in this VLAN. Active ports include all static and dynamic ports that meet the VLAN policy.
STATIC MEMBER	Indicates the set of ports that are static members of this VLAN. A static member of a VLAN is always active and never ages.
NOT_ALLOW MEMBER	Indicates the set of ports that cannot become members of this VLAN.
VLAN PORT NUM	Indicates the VLAN port number for the passive OSPF interface.

IGMP show commands

This section explains the show commands for the Internet Group Management Protocol (IGMP).

IGMP access

Use the **show ip igmp access** command to display information about the IGMP multicast access control groups. The syntax for this command is as follows.

```
show ip igmp access [vrf WORD <1–16>] [vrfids WORD <0–512>]
```

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 17: show ip igmp access field descriptions

Field	Description
INTERFACE	Identifies the interface where multicast access control is configured.
GRP PREFIX	Shows an alphanumeric string that identifies the name of the access policy.
HOSTADDR	Shows the IP address of the host.
HOSTMASK	Shows the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the host network.
ACCESSMODE	Specifies the action of the access policy. The actions are: • deny-tx—deny IP multicast transmitted traffic. • deny-rx—deny IP multicast received traffic. • deny-both—deny both IP multicast transmitted and received traffic. • allow-only-rx—allow IP multicast transmitted traffic. • allow-only-rx—allow IP multicast received traffic. • allow-only-both—allow both IP multicast transmitted and received traffic.

IGMP cache

Use the **show ip igmp cache** command to display information about the IGMP cache. The syntax for this command is as follows.

```
show ip igmp cache [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 18: show ip igmp cache command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INTERFACE	Indicates the physical interface or a logical interface (VLAN), which received group reports from various sources.
LASTREPORTER	Indicates the IP address of the source of the last membership report received for this IP multicast group address on this interface. If the interface does not receive a membership report, this object uses the value 0.0.0.0.

Table continues…

Field	Description
EXPIRATION	Indicates the minimum amount of time that remains before this entry ages out.
V1HOSTTIMER	Indicates the time that remains until the local router assumes that no IGMPv1 members exist on the IP subnet attached to this interface.
TYPE	Indicates whether the entry is learned dynamically or is added statically.
STATICPORTS	Indicates the list of statically-defined ports.

IGMP group

Use the **show ip igmp group** command to display information about the IGMP group. The syntax for this command is as follows.

```
show ip igmp group [count] [member-subnet {default|A.B.C.D/X}] [group
{A.B.C.D} <detail|tracked-members>][vrf WORD <1-16>] [vrfids <0-512>]
```

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

Note:

The ACLI command **show ip igmp group** displays both static and dynamically learned IGMP groups, and the ACLI command **show ip igmp static** command displays only the statically configured IGMP groups. In contrast, the EDM display command under **IP > IGMP > Groups** displays the dynamically learned groups, and the EDM command under **IP > IGMP > Static** displays the statically configured IGMP groups.

The following table shows the field descriptions for this command.

Table 19: show ip igmp group command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INPORT	Indicates the physical interface or a logical interface (VLAN), which received group reports from various sources.
MEMBER	Indicates the IP address of a source that sent a group report to join this group.
EXPIRATION	Indicates the time left before the group report expires on this port. The port updates this variable after it receives a group report.
TYPE	Indicates the group type.

Example

```
VSP-9012:1(config)#show ip igmp group
```

```
================================================================================
                                Igmp Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
232.2.1.1       V2-6/8          2.2.2.2         234        Dynamic
232.2.1.1       V2-10/48        2.2.2.25        232        Dynamic
232.2.1.1       V2-6/8          2.2.2.150       231        Dynamic
239.255.255.250 V2-6/8          2.2.2.157       240        Dynamic

4 out of 4 group Receivers displayed

Total number of unique groups 2
```

IGMP interface

Use the **show ip igmp interface** command to display information about the interfaces where IGMP is enabled. This syntax for this command is as follows.

show ip igmp interface [gigabitethernet {slot/port[-slot/port][,...]}| vlan <1-4084>][vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command if you do not use optional parameters.

Table 20: show ip igmp interface command without parameters

Field	Description
IF	Indicates the interface where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
STATUS	Indicates the activation of a row, which activates IGMP on the interface. The destruction of a row disables IGMP on the interface.
VERS.	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
OPER VERS	Indicates the operational version of IGMP.
QUERIER	Indicates the address of the IGMP querier on the IP subnet to which this interface attaches.

Table continues…

Field	Description
QUERY MAXRSPT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
WRONG QUERY	Indicates the number of queries received where the IGMP version does not match the interface version. You must configure all routers on a LAN to run the same version of IGMP. If queries are received with the wrong version, a configuration error occurs.
JOINS	Indicates the number of times this interface added a group membership.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
LASTMEM QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.
MODE	Indicates the protocol configured on the VLAN added. • snoop — Indicates IGMP snooping is enabled on a VLAN. • snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP Multicast over Fabric Connect for a Layer 2 VSN.) • routed-spb — Indicates IP Multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts. • pim — Indicates PIM is enabled.

The following table shows the field descriptions for this command if you use the interface parameters.

Table 21: show ip igmp interface command with interface parameters

Field	Description
VLAN ID or PORT NUM	Identifies the VLAN or port where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
QUERY MAX RESP	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.

Table continues…

Field	Description
LAST MEMB QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1
PROXY SNOOP ENABLE	Indicates if proxy snoop is enabled on the interface.
SNOOP ENABLE	Indicates if snoop is enabled on the interface.
SSM SNOOP ENABLE	Indicates if SSM snoop is enabled on the interface.
FAST LEAVE ENABLE	Indicates if fast leave mode is enabled on the interface.
FAST LEAVE PORTS (VLAN parameter only)	Indicates the set of ports that are enabled for fast leave
SNOOP QUERIER ENABLE (VLAN parameter only)	Indicates if the IGMP Layer 2 Querier feature is enabled.
SNOOP QUERIER ADDRESS (VLAN parameter only)	Indicates the IP address of the IGMP Layer 2 querier.
DYNAMIC DOWNGRADE VERSION	Indicates if the dynamic downgrade feature is enabled.
COMPATIBILITY MODE	Indicates if compatibility mode is enabled.
EXPLICIT HOST TRACKING	Indicates if explicit host tracking is enabled to track all the source and group members.

Example

```
VSP-9012:1(config)#show ip igmp interface
================================================================================
                          Igmp Interface - GlobalRouter
================================================================================
       QUERY               OPER            QUERY   WRONG            LASTMEM
IF     INTVL STATUS VERS. VERS QUERIER    MAXRSPT QUERY JOINS ROBUST QUERY MODE
--------------------------------------------------------------------------------
V100   125    activ  2     2   0.0.0.0    100     0     0     2      10    routed-spb

1 out of 1 entries displayed
```

IGMP multicast router discovery

Use the **show ip igmp mrdisc** command to display information about the IGMP multicast discovery routes. The syntax for this command is as follows.

show ip igmp mrdisc [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 22: show ip igmp mrdisc command

Field	Description
VLAN ID	Indicates the VLAN ID.
MRDISC	Indicates the status of multicast router discovery.
DISCOVERED RTR PORTS	Indicates the ports discovered.

IGMP multicast router discovery neighbors

Use the **show ip igmp mrdisc neighbors** command to display information about the IGMP multicast router discovery neighbors. The syntax for this command is as follows.

show ip igmp mrdisc neighbors [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 23: show ip igmp mrdisc-neighbors command

Field	Description
VLAN ID	Indicates the VLAN ID.
SRC_PORT	Indicates the source port.
IP Addr	Indicates the IP address.
Advert-int	Indicates the advertisement interval in seconds.
QUERY-int	Indicates the query interval in seconds.
Robust-val	Indicates the tuning for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.

IGMP router-alert

Use the **show ip igmp router-alert** command to display the status of IGMP router alert. The syntax for this command is as follows.

show ip igmp router-alert [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 24: show ip igmp router-alert command

Field	Description
IFINDEX	Indicates the interface index number.
ROUTER ALERT ENABLE	Indicates the status of the router alert check.

IGMP sender

Use the **show ip igmp sender** command to display information about the IGMP senders. The syntax for this command is as follows.

show ip igmp sender [count] [member-subnet {default|A.B.C.D/X}] [group {A.B.C.D}] [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 25: show ip igmp sender command

Field	Description
GRPADDR	Indicates the IP multicast address.
IFINDEX	Indicates the interface index number.
MEMBER	Indicates the IP address of the host.
PORT/MLT	Indicates the IGMP sender ports.
STATE	Indicates if a sender exists because of an IGMP access filter. Options include filtered and nonfiltered.

Example

Display information about IGMP senders:

```
VSP-9012#show ip igmp sender
================================================================================
                                Igmp Sender - GlobalRouter
================================================================================
                                          PORT/
GRPADDR         IFINDEX    MEMBER         MLT        STATE
--------------------------------------------------------------------------------
239.0.0.1       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED
239.0.0.2       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED
239.0.0.3       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED
239.0.0.4       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED
239.0.0.5       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED

5 out of 5 entries displayed
```

IGMP snoop

Use the **show ip igmp snooping** command to display the status of IGMP snoop. The syntax of this command is as follows.

show ip igmp snooping [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 26: show ip igmp snooping command

Field	Description
IFINDEX	Indicates the interface index number.
SNOOP ENABLE	Indicates the status of IGMP snoop.
PROXY SNOOP ENABLE	Indicates the status of IGMP proxy snoop.
SSM SNOOP ENABLE	Indicates the status of IGMP Source Specific Multicast (SSM) snoop.
STATIC MROUTER PORTS	Indicates the set of ports in this VLAN that provide connectivity to an IP multicast router.
ACTIVE MROUTER PORTS	Indicates the active ports.
MROUTER EXPIRATION TIME	Indicates the multicast querier router aging timeout in seconds.
SNOOP QUERIER ENABLE	Indicates if the IGMP Layer 2 Querier feature is enabled.
SNOOP QUERIER ADDRESS	Indicates the IP address of the IGMP Layer 2 querier.
DYNAMIC DOWNGRADE VERSION	Indicates if the Virtual Services Platform 9000 downgrades the version of IGMP to handle older query messages.
COMPATIBILITY MODE	Indicates if IGMPv3 is compatible with IGMPv2

IGMP static and blocked ports

Use the **show ip igmp static** command to display information about the static and blocked ports for the IGMP-enabled interfaces. The syntax for this command is as follows.

show ip igmp static [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 27: show ip igmp static command

Field	Description
GRPADDR	Indicates the IP multicast address. The group address holds the starting range for the address range.
TO-GRPADDR	Indicates the end of the range for the group address.
INTERFACE	Indicates the interface IP address.
STATICPORTS	Indicates the egressing ports.
BLOCKEDPORTS	Indicates the ports not allowed to join.

Multicast group trace for IGMP snoop

Use the **show ip igmp snoop-trace** command to view multicast group trace information for IGMP snoop. Multicast group trace tracks the data flow path of the multicast streams. This command provides information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port. The syntax for the command is as follows.

show ip igmp snoop-trace [source {A.B.C.D}] [group {A.B.C.D}] [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table provides the field descriptions for this command.

Table 28: show ip igmp snoop-trace command

Field	Description
GROUP ADDRESS	Indicates the IP multicast group address for which this entry contains information.
SOURCE ADDRESS	Indicates the source of the multicast traffic.
IN VLAN	Indicates the incoming VLAN ID.
IN PORT	Indicates the incoming port number.
OUT VLAN	Indicates the outgoing VLAN ID.
OUT PORT	Indicates the outgoing port number.
TYPE	Indicates where the stream is learned. ACCESS indicates the stream is learned on UNI ports. NETWORK indicates the stream is learned over the SPBM network.

```
VSP-9012:1# show ip igmp snoop-trace

================================================================================
                           Snoop Trace - GlobalRouter
================================================================================
GROUP              SOURCE             IN     IN         OUT    OUT        TYPE
```

```
ADDRESS            ADDRESS            VLAN   PORT       VLAN   PORT
-------------------------------------------------------------------------------
233.252.0.1        192.0.2.6          500    spb        500    9/5         NETWORK
233.252.0.10        192.0.2.7         500    spb        500    10/10       NETWORK
```

SSM map information

Use the **show ip igmp ssm-map** command to display the list of SSM maps. The syntax for this command is as follows.

show ip igmp ssm-map [vrf WORD <1-16>] [vrfids WORD <0-512>]

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 29: show ip igmp ssm-map command

Field	Description
GROUP	Indicates the IP multicast group address that uses the default range of 232/8.
SOURCE	Indicates the IP address of the source that sends traffic to the group source.
MODE	Indicates that the entry is a statically configured entry (static) or a dynamically learned entry from IGMPv3 (dynamic).
ACTIVE	Indicates the activity on the corresponding source and group. If the source is active and traffic is flowing to the switch, this status is active; otherwise, it is nonactive.
STATUS	Indicates the administrative state and whether to use the entry. If the status is enabled (default), the entry is used. If the status is disabled, the entry is not used but is saved for future use.

Example

```
VSP-9012:1(config)#show ip igmp ssm-map
================================================================================
                              Igmp Ssm Channel - GlobalRouter
================================================================================
GROUP           SOURCE          MODE       ACTIVE     STATUS
--------------------------------------------------------------------------------
232.1.1.1       122.122.122.200 dynamic    false      enabled
232.1.1.2       122.122.122.200 dynamic    false      enabled
232.1.1.3       122.122.122.200 dynamic    false      enabled
232.1.1.4       122.122.122.200 dynamic    false      enabled
232.1.1.5       122.122.122.200 dynamic    false      enabled
232.1.1.6       122.122.122.200 dynamic    false      enabled
232.1.1.7       122.122.122.200 dynamic    false      enabled
232.1.1.8       122.122.122.200 dynamic    false      enabled
232.1.1.9       122.122.122.200 dynamic    false      enabled
232.1.1.10      122.122.122.200 dynamic    false      enabled

10 out of 10 entries displayed
```

SSM group range and dynamic learning status

Use the **show ip igmp ssm** command to display the SSM group range and the status of dynamic learning. The syntax for this command is as follows.

```
show ip igmp ssm [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 30: show ip igmp ssm command

Field	Description
DYNAMIC LEARNING	Indicates whether dynamic learning is enabled at a global level.
SSM GROUP RANGE	Indicates the IP address range for the SSM group.

PIM show commands

This section explains the show commands for Protocol Independent Multicast (PIM).

PIM active RP

Use the **show ip pim active-rp** command to display information about the active rendezvous point (RP) for all groups or a specific group. If you do not specify an IP address, you receive information about the active RP for all the running multicast groups on the switch. The syntax for this command is as follows.

```
show ip pim active-rp [group {A.B.C.D}]
```

The following table shows the field descriptions for this command.

Table 31: show ip pim active-rp command

Field	Description
GRPADDR	Shows the IP address of the multicast group.
RP-ADDR	Shows the IP address of the RP router. This address must be one of the local PIM-SM enabled interfaces.
RP-PRIORITY	Shows the priority of the RP.

Example

Display information about the active rendezvous points:

```
VSP-9012#show ip pim active-rp
================================================================================
                          Pim Grp->RP Active RP Table - GlobalRouter
================================================================================
GRPADDR          RP-ADDR          RP-PRIORITY
--------------------------------------------------------------------------------
239.0.0.1         20.0.0.90         0
239.0.0.2         20.0.0.90         0
239.0.0.3         20.0.0.90         0
239.0.0.4         20.0.0.90         0
239.0.0.5         20.0.0.90         0
239.255.255.250   20.0.0.90         0
```

PIM bootstrap router

Use the **show ip pim bsr** command to display information about the bootstrap router (BSR) for this PIM-SM domain. The syntax for this command is as follows.

show ip pim bsr

The following table shows the field descriptions for this command.

Table 32: show ip pim bsr command

Field	Description
Current BSR address	Shows the IP address of the current BSR for the local PIM domain.
Current BSR priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
Current BSR HaskMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. The hash-mask allows a small number of consecutive groups (for example, 4) to always hash to the same RP.
Current BSR Fragment	Shows a randomly generated number that distinguishes fragments that belong to different bootstrap messages. Fragments that belong to the same bootstrap message carry the same fragment tag.
Pim Boostrap Timer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

PIM candidate rendezvous points

Use the **show ip pim rp-candidate** command to display information about the candidate rendezvous points for the PIM-SM domain.The syntax for this command is as follows.

show ip pim rp-candidate

The following table shows the field descriptions for this command.

Table 33: show ip pim rp-candidate command

Field	Description
GRPADDR	Displays the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Displays the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
RPADDR	Displays the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

PIM interface

Use the **show ip pim interface** command to display information about the PIM-SM interface configuration on the switch. The syntax of this command is as follows.

show ip pim interface [gigabitethernet {slot/port[-slot/port][,...]}|vlan <1-4084>]

The following table shows the field descriptions for this command if you do not use optional parameters.

Table 34: show ip pim interface command without parameters

Field	Description
IF	Indicates the slot and port number or VLAN ID of the interface where PIM is enabled.
ADDR	Shows the IP address of the PIM interface.
MASK	Shows the network mask for the IP address of the PIM interface.
MODE	Indicates the configured mode of this interface. The valid modes are SSM and Sparse.
DR	Shows the designated router (DR) for this interface.
HLINT	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.

Table continues…

Field	Description
JPINT	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default join and prune interval is 60 seconds.
CBSPR	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
OPSTAT	Indicates the status of PIM on this interface: up or down.
INTF TYPE	Indicates whether the PIM interface is active or passive.

The following table shows the field descriptions for this command if you use optional parameters.

Table 35: show ip pim interface command with parameters

Field	Description
VLAN-ID or PORT-NUM	Indicates the slot and port number or VLAN ID of the interface where PIM is enabled.
PIM ENABLE	Indicates the administrative status of PIM
MODE	Indicates the configured mode of this interface. The valid modes are SSM and Sparse.
HELLOINT	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default join and prune interval is 60 seconds.
CBSRPREF	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
INTF TYPE	Indicates whether the PIM interface is active or passive.

Example

```
VSP-9012:1(config)#show ip pim interface
```

```
================================================================================================================
                                  Pim Interface - GlobalRouter
================================================================================================================
IF        ADDR            MASK            MODE       DR       HLINT  JPINT  CBSPR          OPSTAT INTF TYPE
Port4/4   12.10.11.1      255.255.255.0   ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Port4/30  13.10.10.1      255.0.0.0       ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Clip1     3.3.3.3         255.0.0.0       ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Clip2     5.5.5.5         255.0.0.0       ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Vlan2     2.2.2.32        255.0.0.0       ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Vlan3     30.30.30.32     255.255.255.0   ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Vlan4     40.1.1.32       255.255.255.0   ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Vlan5     10.10.10.1      255.0.0.0       ssm        0.0.0.0  30     60     -1 (disabled)  down   active
Vlan200   1.1.1.2         255.255.255.0   ssm        0.0.0.0  30     60     -1 (disabled)  down   active

Total PIM Interfaces Displayed 9/9
----------------------------------------------------------------------------------------------------------------
```

PIM mode

Use the **show ip pim mode** command to show the PIM mode (SM or SSM). The syntax for this command is as follows.

```
show ip pim mode
```

The following table shows the field description for this command.

Table 36: show ip pim mode command

Field	Description
Mode	Indicates the PIM mode as SM or SSM.

PIM neighbor

Use the **show ip pim neighbor** command to display information about the neighboring routers configured with PIM-SM. The syntax for this command is as follows.

```
show ip pim neighbor
```

The following table shows the field descriptions for this command.

Table 37: show ip pim neighbor command

Field	Description
INTERFACE	Indicates the interface number.
ADDRESS	Indicates the IP address of the PIM neighbor.
UPTIME	Indicates the elapsed time since this PIM neighbor last became a neighbor of the local router.
EXPIRE	Indicates the time that remains before this PIM neighbor times out.

PIM route

Use the **show ip pim mroute** command to display information from the route table. The syntax for this command is as follows.

```
show ip pim mroute [group <A.B.C.D>] [source <A.B.C.D>] [terse]
```

Note:

In a PIM-SM or PIM-SSM Layer 3 MLT/SMLT multicast environment, when an SMLT link down or SMLT link up event occurs, or when an individual port in an (S)MLT goes down or comes back up, traffic can be re-hashed (switched over) either to another port in the (S)MLT or to any of the IST's MLT ports. This is valid, as the nature of the (S)MLT environment is that traffic can ingress on any one of these ports and be successfully forwarded to receivers. However, the Incoming Port record in the following table may not accurately reflect which port the data is

arriving on at any given time. This does not cause traffic loss. Checking traffic statistics on the ports of the (S)MLT/ IST can be used to determine the ingress port.

The following table shows the field descriptions for this command.

Table 38: show ip pim mroute command

Field	Description
Src	Displays the IP address of the source that sends the multicast stream. A nonzero value indicates that a source sends multicast traffic. 0.0.0.0 indicates that this entry is created in response to a receiver that wants to receive this traffic.
Grp	Displays the IP multicast group address.
RP	Displays the IP address of the RP router.
Upstream	Displays the IP address of the next hop that a multicast packet takes when received on the correct port as listed on the incoming interface.
Flags	Displays the flags configured based on the condition of the receivers, the RP, and the senders. Use the legend at the bottom of the output to explain the flag values.
Incoming Port	Lists the port through which a multicast packet can ingress. If the port is a member of a Multi-Link Trunk (MLT), the packets can ingress on any port of the MLT.
Outgoing Ports	Lists all ports through which traffic that enters on incoming ports exit.
Joined Ports	Lists all ports that received PIM join messages.
Pruned Ports	List all ports that received PIM prune messages.
Leaf Ports	Lists multicast receivers that directly connect to the router.
Asserted Ports	Lists all ports that received assert messages. The router uses assert messages to help determine the best path to the source.
Prune Pending Ports	Lists all ports currently in the prune-pending state.
Assert Winner Ifs	Lists interfaces elected the assert winner. The winner continues to forward multicast traffic to the LAN.
Assert Loser Ifs	Lists interfaces not elected as the assert winner. The loser interface is pruned.
Timers	Displays the up time and expiration time for the entry in the routing table.
AssertVifTimer	Displays the time after which the assert winner state refreshes.

Example

```
VSP-9012:1(config)#show ip pim mroute
```

```
================================================================================
                              Pim Multicast Route - GlobalRouter
================================================================================
Src: 10.1.1.3      Grp: 232.2.1.1    RP: 0.0.0.0     Upstream: 70.70.70.4
Flags: SPT CACHE SG
Incoming  Port: Vlan70-MLT-4(6/24),
Outgoing  Ports: Vlan2-6/8,10/48,
Joined    Ports: Vlan2-6/8,
Pruned    Ports:
Leaf      Ports: Vlan2-10/48,
Asserted  Ports:
Prune Pending Ports:
Assert Winner Ifs:
Assert Loser Ifs:
TIMERS:
Entry   JP   RS  Assert
207     9    0       0
VLAN-Id:     2     3     4    70
Join-P:   191     0     0     0
Assert:     0     0     0     0
--------------------------------------------------------------------------------
Src: 10.1.1.4      Grp: 232.2.1.1    RP: 0.0.0.0     Upstream: 70.70.70.4
Flags:
SPT CACHE SG
Incoming  Port: Vlan70-MLT-4(6/24),
Outgoing Ports: Vlan2-6/8,10/48,
Joined   Ports: Vlan2-6/8,
Pruned   Ports:
Leaf     Ports: Vlan2-10/48,
Asserted Ports:
Prune Pending Ports:
Assert Winner Ifs:
Assert Loser Ifs:
TIMERS:
Entry   JP   RS  Assert
230     19    0       0
VLAN-Id:     2     3     4    70
Join-P:   203     0     0     0
Assert:     0     0     0     0
--------------------------------------------------------------------------------
Total Num of Entries Displayed 2/2
```

PIM virtual neighbor

Use the **show ip pim virtual-neighbor** command to display the virtual neighbor. The syntax for this command is as follows.

show ip pim virtual-neighbor

The following table shows the field descriptions for this command.

Table 39: show ip virtual-neighbor command

Field	Description
INTERFACE	Indicates the interface.
ADDRESS	Indicates the IP address of the virtual neighbor.

Rendezvous points (for groups)

Use the **show ip pim rp-hash** command to display information about the RPs selected for a multicast group. The syntax for this command is as follows.

```
show ip pim rp-hash
```

The following table shows the field descriptions for this command.

Table 40: show ip pim rp-hash command

Field	Description
GRPADDRESS	Shows the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Shows the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
ADDRESS	Shows the IP address of the C-RP router.
HOLDTIME	Shows the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
EXPTIME	Shows the time that remains before this C-RP router times out.

Static RP table

Use the **show ip pim static-rp** command to display the static RP table. The syntax for this command is as follows.

```
show ip pim static-rp
```

The following table shows the field descriptions for this command.

Table 41: show ip pim static-rp command

Field	Description
GRPADDR	Indicates the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a static RP.

Table continues…

Field	Description
GRPMASK	Indicates the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a static RP.
RPADDR	Indicates the IP address of the static RP. This address must be one of the local PIM-SM enabled interfaces.
STATUS	Indicates the status of static RP.

Example

Display the static RP table:

```
VSP-9012#show ip pim static-rp
```

```
================================================================================
                             Pim Static RP Table - GlobalRouter
================================================================================
GRPADDR          GRPMASK          RPADDR           STATUS
--------------------------------------------------------------------------------
239.0.0.0        255.0.0.0        20.0.0.90       valid
```

Glossary

bootstrap router (BSR)	A dynamically elected Protocol Independent Multicast (PIM) router that collects information about potential Rendezvous Point routers and distributes the information to all PIM routers in the domain.
Bridge Protocol Data Unit (BPDU)	A data frame used to exchange information among the bridges in local or wide area networks for network topology maintenance.
candidate bootstrap router (C-BSR)	Provides backup protection in case the primary rendezvous point (RP) or boostrap router (BSR) fails. Protocol Independent Multicast (PIM) uses the BSR and C-BSR.
Circuitless IP (CLIP)	A CLIP is often called a loopback and is a virtual interface that does not map to any physical interface.
designated router (DR)	A single router elected as the designated router for the network. In a broadcast or nonbroadcast multiple access (NBMA) network running the Open Shortest Path First (OSPF) protocol, a DR ensures all network routers synchronize with each other and advertises the network to the rest of the Autonomous System (AS). In a multicast network running Protocol Independent Multicast (PIM), the DR acts as a representative router for directly connected hosts. The DR sends control messages to the rendezvous point (RP) router, sends register messages to the RP on behalf of directly connected sources, and maintains RP router status information for the group.
distribution tree	A set of multicast routers and subnetworks that allow the group members to receive traffic from a source.
Internet Assigned Numbers Authority (IANA)	The central registry for various assigned numbers, for example, Internet protocol parameters (such as port, protocol, and enterprise numbers), options, codes, and types.
Internet Group Management Protocol (IGMP)	IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.
interswitch trunking (IST)	A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one

	interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.
IP Multicast over Fabric Connect	With IP Multicast over Fabric Connect, Avaya introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. These extensions, combined with the use of IGMP snooping and querier functions at the edge of the SPBM cloud, efficiently transport IP multicast data by using sub-trees of the VSN shortest path tree per IP multicast group.
last member query interval (LMQI)	The time between when the last Internet Group Management Protocol (IGMP) member leaves the group and the stream stops.
latency	The time between when a node sends a message and receipt of the message by another node; also referred to as propagation delay.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 2 Virtual Services Network	The Layer 2 Virtual Services Network (L2 VSN) feature provides IP connectivity over SPBM for VLANs. Backbone Edge Bridges (BEBs) handle Layer 2 virtualization. At the BEBs you map the end-user VLAN to a Service Instance Identifier (I-SID). BEBs that have the same I-SID configured can participate in the same Layer 2 Virtual Services Network (VSN).
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
Layer 3 Virtual Services Network	The Layer 3 Virtual Services Network (L3 VSN) feature provides IP connectivity over SPBM for VRFs. Backbone Edge Bridges (BEBs) handle Layer 3 virtualized. At the BEBs through local provisioning, you map the end-user IP enabled VLAN or VLANs to a Virtualized Routing and Forwarding (VRF) instance. Then you map the VRF to a Service Instance Identifier (I-SID). VRFs that have the same I-SID configured can participate in the same Layer 3 Virtual Service Network (VSN).
multicast router discovery (MRDISC)	Provides the automatic discovery of multicast-capable routers. By listening to multicast router discovery messages, Layer 2 devices can determine where to send multicast source data and Internet Group Management Protocol (IGMP) host membership reports.
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
nonbroadcast multiaccess (NBMA)	Interconnects multiple devices over a broadcast network through point-to-point links. NBMA reduces the number of IP addresses required for point-to-point connections.

packet loss	Expressed as a percentage of packets dropped over a specified interval. Keep packet loss to a minimum to deliver effective IP telephony and IP video services.
Protocol Independent Multicast, Source Specific (PIM-SSM)	PIM-SSM is a multicast routing protocol for IP networks. PIM-SSM uses only shortest-path trees to provide multicast services based on subscription to a particular (source, group) channel. PIM-SSM eliminates the need for starting with a shared tree by immediately joining a source through the shortest path tree. This method enables PIM-SSM to avoid using a rendezvous point (RP) and RP-based shared tree, which can be a potential bottleneck.
Protocol Independent Multicast, Sparse Mode (PIM-SM)	PIM-SM is a multicast routing protocol for IP networks. PIM-SM provides multicast routing for multicast groups that can span wide-area and inter-domain networks, where receivers are not densely populated. PIM-SM sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic. PIM-SM adds a Rendezvous Point router to avoid multicast-data flooding. Use PIM-SM when receivers for multicast data are sparsely distributed throughout the network.
quality of service (QoS)	QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can configure a higher priority for IP deskphones, which need a fixed bit rate and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.
rendezvous point (RP)	The root of the shared tree. One RP exists for each multicast group. The RP gathers information about available multicast services through the reception of control messages and the distribution of multicast group information. Protocol Independent Multicast (PIM) uses RPs.
reverse path forwarding (RPF)	Prevents a packet from forging its source IP address. Typically, the system examines and validates the source address of each packet.
route table manager (RTM)	Determines the best route to a destination based on reachability, route preference, and cost.
routing policy	A form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.
Shortest Path Bridging (SPB)	Shortest Path Bridging is a control Link State Protocol that provides a loop-free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.

Shortest Path Bridging MAC (SPBM)	Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.
shortest path tree (SPT)	Creates a direct route between the receiver and the source for group members in a Protocol Independent Multicast-Sparse Mode (PIM-SM) domain.
SMLT aggregation switch	One of two IST peer switches that form a split link aggregation group. It connects to multiple wiring closet switches, edge switches, or customer premise equipment (CPE) devices.
Switch Fabric (SF) module	The Switch Fabric module connects to all I/O and Control Processor modules. You can install 6 SF modules in Virtual Services Platform 9000, using 5 SF modules plus 1 as a hot backup. The SF modules comprise a data path and a control path, and provide a back end switching solution in the midplane chassis.
time-to-live (TTL)	The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.
trunk	A logical group of ports that behaves like a single large port.
trunk port	A port that connects to the service provider network such as the MPLS environment.
virtual router forwarding (VRF)	Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF as a separate physical router.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.