Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000

Release 4.0
NN46250-510
Issue 03.01
December 2014

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.

License 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 provides instructions to configure Avaya VENA Fabric Connect on the Avaya Virtual Services Platform 9000. Fabric Connect includes Shortest Path Bridging (SPB, the MAC-in-MAC variant of IEEE 802.1aq), Intermediate System to Intermediate System (IS-IS), and Connectivity Fault Management (CFM).

Using the document

The document is organized into three main sections:

1. Infrastructure configuration — You must first configure your base SPB and IS-IS architecture described in the infrastructure configuration chapter. This allows SPB to operate on the switch. The chapter includes initial steps to configure the minimum SPB and IS-IS parameters to enable Fabric Connect on your network, at Configuring minimum SPBM and IS-IS parameters on page 41.
2. Services configuration — After you have completed the infrastructure configuration, you configure the appropriate services for your network to run on top of your base architecture. Services can include: Layer 2 VSNs, IP Shortcut Routing, Layer 3 VSNs, and Inter-VSN Routing.
3. Operations and management — Finally, Virtual Services Platform 9000 provides tools to monitor and troubleshoot your Fabric Connect network.

The document also includes configuration examples at the end of each chapter to show basic configurations for use of SPBM.

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 detail what is new in *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510 for Release 4.0.

Features

See the following sections for information about feature changes.

9012QQ-2 I/O module

Release 4.0.1 introduces a second generation 9012QQ-2 Input/Output (I/O) module.

The 9012QQ-2 module is a 12-port 40 gigabits per second (Gbps) module that supports QSFP+ fiber transceivers, and also directly-attached cables (DACs).

For more information about the 9012QQ-2 I/O module, see

- SPBM restrictions and limitations on page 30.

For more information about the 9012QQ-2 module specifications, see *Installing Modules in Avaya Virtual Services Platform 9000,* NN46250-301.

9048XS-2 module

Release 4.0 introduces a second generation 9048XS-2 Input/Output (I/O) module.

The 9048XS-2 module is a 48 port 10 Gigabit per second (Gbps) module that supports both 10GBASE-R small form factor pluggable plus (SFP+) and 1000BASE-X SFP transceivers. You can use second generation I/O modules in first generation mode or second generation mode. For more information, see:

- SPBM restrictions and limitations on page 30.

CFM Global tab

Release 4.0 adds additional options for the **Global** tab. For more information, see CFM Global field descriptions on page 350 and CFM Global field descriptions on page 351.

ECMP within ISIS routes

Release 4.0 updates tie-breaking rules related to Equal Cost Multipath (ECMP). Multiple BEBs can announce the same route, either because the Layer 2 LAN connects to multiple BEBs for redundancy, or because segments of the LAN are Layer 2 bridged. If the device has to tie-break

between the multiple sources, the device uses the following precedence rules to tie-break. In the following order, the device prefers:

1. Local routes over Inter-VSN routes.
2. Routes with the lowest route preference.

   You can change this with route-map within the IS-IS accept policy.
3. Routes with the lowest SPBM cost.
4. Routes with lowest prefix cost.

   You can change this with route-map on the remote advertising node with the **`redistribute`** command, or with route-map on the local node with the IS-IS accept policy.
5. Routes from the VSN with the lower Layer 3 VSN I-SID.

   The device considers the Global Routing Table (GRT) to have an I-SID equal to zero.

For more information, see ECMP within ISIS routes on page 164.

Intermediate-System-to-Intermediate-System (IS-IS) accept policies

Release 4.0 adds Intermediate-System-to-Intermediate-System (IS-IS) accept policies. You can use IS-IS accept policies with Layer 3 VSNs or IP Shortcuts to filter incoming IS-IS routes over the SPBM cloud. IS-IS accept policies enable the device to determine whether to add an incoming route to the routing table.

You can create an IS-IS accept policy for the Global Routing Table (GRT) or a Virtual Routing and Forwarding (VRF) instance. You can create an IS-IS accept policy on a switch that operates at a global default level or for a specific advertising BEB. You can also use the filter mechanism for IS-IS accept policies to redistribute routes between different VRFs, or between a VRF and the GRT.

IS-IS accept policies can use route policies to determine what traffic to accept into the routing table.

IS-IS policies can also use either a service instance identifier (I-SID) or an I-SID list to filter incoming traffic. For inter-VRF route redistribution, an I-SID value of 0 represents the GRT. For inter-VRF route redistribution between VRFs, the I-SID is the source VRF (or remote VRF).

For information on configuring IS-IS accept policies, see:

- IS-IS accept policies on page 172 for conceptual information.
- IS-IS redistribution policies on page 222 for conceptual information.
- Configuring IS-IS accept policies on page 181 for ACLI configuration.
- Configuring inter-VRF accept policies on VRFs on page 184 for ACLI configuration.
- Viewing IS-IS accept policy information on page 186 for ACLI information.
- Applying IS-IS accept policies globally on page 208 for EDM configuration.
- Configuring an IS-IS accept policy for a specific advertising BEB on page 209 for EDM configuration.
- Configuring an IS-IS accept policy to apply for a specific I-SID on page 210 for EDM configuration.
- Configuring an I-SID list for an IS-IS accept policy on page 212 for EDM configuration.
- Configuring an IS-IS accept policy for a specific I-SID list on page 213 for EDM configuration.

- Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list on page 214 for EDM configuration.

Layer 2 ping

Syntax corrected in the procedure for the **`l2 ping`** command. For more information, see Triggering a Layer 2 ping on page 334.

Maintenance Domain name and Maintenance Association name

Release 4.0 updates the variable range to <1–22> for Maintenance Domain (MD) name and Maintenance Association (MA) name for Connectivity Fault Management (CFM). For more information, see:

- Configuring CFM MD on page 326.
- Configuring CFM MA on page 327.
- Configuring CFM MEP on page 330.
- Triggering a loopback test (LBM) on page 332.
- Triggering linktrace (LTM) on page 333.

Router node name

The Virtual Services Platform 9000 does not support the routernodename option for C-VLANs with the **`l2 ping`** and **`l2 traceroute`** commands. For more information, see:

- Variable definitions on page 336.
- Variable definitions on page 339.

SPBM drop count

Virtual Services Platform 9000 does not support the **`show isis drop-stats port unknown-unicast-sa`** drop count parameter for second generation cards in this release. The device always displays the ACLI output for second generation modules as 0 for this counter. For more information, see Displaying SPBM packet drop statistics by port on page 80.

Update to show isis spbm multicast-fib command

Release 4.0 updates the **`show isis spbm multicast-fib`** command to add the field `INCOMING INTERFACE` to the show output. For more information, see

- SPBM FIB on page 28.
- Displaying the IP unicast FIB, multicast FIB, unicast FIB, and unicast tree on page 68.
- Verifying Layer 2 VSN operation on page 158.

New parameter for show isis lsdb command

Release 4.0 updates the **`show isis lsdb`** command with the new **`ip-unicast [i-sid <0-16777215>][lspid <xxxx.xxxx.xxxx.xx-xx>][sysid <xxxx.xxxx.xxxx>]`** parameters. For more information, see Displaying IS-IS LSDB and adjacencies on page 72.

TLVs

Release 4.0 adds additional type-length-value (TLVs) for use with Intermediate-System-to-Intermediate-System (IS-IS). Review the following table for more information on the new TLVs.

TLV	Description
3	End system neighbors — The end system neighbors TLV lists adjacent level 1 routers and end systems.
5	Prefix neighbors — The prefix neighbors TLV specifies reachable address prefix neighbors.
12	Optional checksum — The optional checksum can be included in all Complete Sequence Numbers Protocol Data (CSNP), Partial Sequence Numbers Protocol Data (PSNP) PSNP, and IS-IS Hello Packets (IIH) packets.
128	IP addresses — Specifies the IP addresses known by a router.
137	Dynamic host name — Specifies the symbolic name of the router where the Link-State packet (LSP) originates. This TLV is optional.

For more information, see Standard TLVs on page 22.

Update to show ip route command output

Release 4.0 updates the output for the **`show ip route [vrf WORD<1-16>]`** command. The NH VRF column now displays as NH VRF/ISID. The column can either be the next hop VRF or the destination I-SID. The NH VRF/ISID column displays the I-SID in the following examples:

- Only for inter-Virtual Services Network (VSN) routes redistributed using IS-IS accept policies.
- Only if the VRF does not exist locally.
- Only if the I-SID for which the routes are redistributed does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays.

If the I-SID is 0, which represents the GlobalRouter, the column displays as GlobalRouter.

The existing IS-IS routes in Shortest Path Bridging (SPB) Layer 3 VSN continue to display as the VRF name of the IP VSN.

For more information on IS-IS accept policies, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

For more information on the **`show ip route [vrf WORD<1-16>]`** command, see:

- Viewing IS-IS accept policy information on page 186.
- IP Shortcuts configuration example on page 217.
- Verifying Layer 3 VSN operation on page 284.
- Verifying Inter-VSN Routing operation on page 307.

Update to show isis spbm ip-unicast-fib command

Release 4.0 adds the `IP ROUTE PREFERENCE` column and **Dest ISID** to the **`show isis spbm ip-unicast-fib`** command output and the **DestIsid** and **Preference** column to the **IP Unicast FIB** tab.

The **`show isis spbm ip-unicast-fib`** command and **IP Unicast FIB** tab display all of the IS-IS routes in the IS-IS LSDB. The new columns for ACLI and EDM display the IP route preference.

Routes within the same VSN are added to the LSDB with a default preference of 7. Inter-VSN routes are added to the LSDB with a route preference of 200. IS-IS accept policies allow you to change the route preference for incoming routes. If the same route is learned from multiple sources with different route preferences, then the routes are not considered equal cost multipath (ECMP) routes. The route with the lowest route preference is the preferred route.

For more information, see

- Displaying the IP unicast FIB, multicast FIB, unicast FIB, and unicast tree on page 68.
- Displaying the IP unicast FIB on page 98.
- Viewing IS-IS accept policy information on page 186.
- IP Shortcuts configuration example on page 217.
- Verifying Layer 3 VSN operation on page 284.

Update to show vlan i-sid

Release 4.0 updates the output for the **`show vlan i-sid <1–4084>`** command. For more information, see:

- Enabling IP multicast over SPBM globally on page 49.
- Configuring SPBM Layer 2 VSN on page 130.

VLANs without member ports

VLANs without member ports on page 19 is added to the document to explain how the VSP 9000 designates a VLAN without member ports as operationally up.

Update to VRF name variable

Release 4.0 updates the VRF name variable to **`vrf WORD<1–16>`** in ACLI. EDM VRF name options are also updated throughout the document to a range of 1–16.

Other changes

See the following section for information about changes that are not feature-related.

Purpose

Release 4.0 updates the purpose section of the document to include Avaya VENA Fabric Connect. For more information, see Purpose on page 8.

Document title

Release 4.0 updates the document title to *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000*, NN46250–510, from *Avaya Virtual Services Platform 9000 Configuration – Shortest Path Bridging MAC (SPBM)*, NN46250–510.

Chapter 3: SPBM and IS-IS infrastructure configuration

This chapter provides concepts and procedures to configure the basic infrastructure for Shortest Path Bridging MAC (SPBM).

SPBM and IS-IS infrastructure fundamentals

Shortest Path Bridging MAC (SPBM) is a next generation virtualization technology that revolutionizes the design, deployment, and operations of carriers and service providers, along with enterprise campus core networks and enterprise data centers. SPBM provides massive scalability while at the same time reducing the complexity of the network.

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 that provides all virtualization services in an integrated model. In addition, by relying on endpoint service provisioning only, the idea of building your network once and not touching it again becomes a true reality. This technology provides all the features and benefits required by carrier-grade, enterprise and service provider deployments without the complexity of alternative technologies, for example, Multiprotocol Label Switching (MPLS).

SPBM simplifies deployments by eliminating the need to configure multiple points throughout the network. When you add new connectivity services to an SPBM network you do not need intrusive core provisioning. The simple endpoint provisioning is done where the application meets the network, with all points in between automatically provisioned through the robust link-state protocol, Intermediate-System-to-Intermediate-System (IS-IS).

Most Ethernet based networks use 802.1Q tagged interfaces between the routing switches. SPBM uses two Backbone VLANs (B-VLANs) that are used as the transport instance. A B-VLAN is not a traditional VLAN in the sense that it does not flood unknown, broadcast or multicast traffic, but only forwards based on IS-IS provisioned backbone MAC (B-MAC) tables. After you configure the B-VLANs and the IS-IS protocol is operational, you can map the services to service instances.

SPBM uses IS-IS to discover and advertise the network topology, which enables it to compute the shortest path to all nodes in the SPBM network. SPBM uses IS-IS shortest path trees to populate forwarding tables for the individual B-MAC addresses of each participating node.

To forward customer traffic across the core network backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a B-MAC destination address (BMAC-DA) and a B-MAC source address (BMAC-SA). Encapsulating customer

MAC addresses in B-MAC addresses improves network scalability (no end-user C-MAC learning is required in the core) and also significantly improves network robustness (loops have no effect on the backbone infrastructure.)

The SPBM B-MAC header includes a Service Instance Identifier (I-SID) with a length of 32 bits with a 24-bit ID. I-SIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. You can use I-SIDs in a Virtual Services Network (VSN) for VLANs or VRFs across the MAC-in-MAC backbone:

- **Unicast**
  - For a Layer 2 VSN, the device associates the I-SID with a customer VLAN, which the device then virtualizes across the backbone. Layer 2 VSNs associate one VLAN per I-SID.
  - With Layer 3 VSN, the device associates the I-SID with a customer VRF, which the device virtualizes across the backbone. Layer 3 VSNs associate one VRF per I-SID.
  - With Inter-VSN routing, Layer 3 devices, routers, or hosts connect to the SPBM cloud using the SPBM Layer 2 VSN service. The Backbone Core Bridge can transmit traffic between different VLANs with different I-SIDs.
  - With IP shortcuts, no I-SID is required, forwarding for the Global Routing Table (GRT) is done using IS-IS based shortest path BMAC reachability.
- **Multicast**
  - With Layer 2 VSN with IP multicast over SPBM, the BEB associates a data I-SID with the multicast stream and the scope I-SID is based on the Layer 2 VSN I-SID.
  - With Layer 3 VSN with IP multicast over SPBM, the BEB associates a data I-SID with the multicast stream and the scope I-SID is based on the Layer 3 VSN I-SID.
  - With IP Shortcuts with IP multicast over SPBM, the BEB associates a data I-SID with the multicast stream, but there is no I-SID for the scope, which is the Global Routing Table (GRT).

    Note:

    Inter-VSN routing for IP multicast over SPBM is not supported.

Virtual Services Platform 9000 supports the IEEE 802.1aq standard of SPBM, which allows for larger Layer 2 topologies and permits faster convergence.

MAC-in-MAC encapsulation

To forward customer traffic across the core network backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses.

The originating node creates a MAC header that is used for delivery from end to end. As the MAC header stays the same across the network, there is no need to swap a label or do a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end.

The encapsulation of customer MAC addresses in backbone MAC addresses greatly improves network scalability, as no end-user MAC learning is required in the backbone, and also significantly

improves network robustness, as customer-introduced network loops have no effect on the backbone infrastructure.

I-SID

SPBM introduces a service instance identifier called I-SID. SPBM uses I-SIDs to separate services from the infrastructure. After you create an SPBM infrastructure, you can add additional services (such as VLAN extensions or VRF extensions) by provisioning the endpoints only. The SPBM endpoints are BEBs, which mark the boundary between the core MAC-in-MAC SPBM domain and the edge customer 802.1Q domain. I-SIDs are provisioned on the BEBs to be associated with a particular service instance. In the SPBM core, the bridges are BCBs. BCBs forward encapsulated traffic based on the BMAC-DA.

The SPBM B-MAC header includes an I-SID. The length of the I–SID is 32 bits with a 24–bit ID. I-SIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. These I-SIDs are used in a VSN for VLANs or VRFs across the MAC-in-MAC backbone:

- For a Layer 2 VSN, the I-SID is associated with a customer VLAN, which is then virtualized across the backbone. Layer 2 VSNs offer an any-any LAN service type. Layer 2 VSNs associate one VLAN per I-SID.
- For a Layer 2 VSN with IP multicast over SPBM, the BEB associates a data I-SID with the multicast stream and a scope I-SID that defines the scope as Layer 2 VSN. A multicast stream with a scope of Layer 2 VSN can only transmit a multicast stream for the same Layer 2 VSN.
- For a Layer 3 VSN, the I-SID is associated with a customer VRF, which is also virtualized across the backbone. Layer 3 VSNs are always full-mesh topologies. Layer 3 VSNs associate one VRF per I-SID.
- For a Layer 3 VSN with IP multicast over SPBM, the BEB associates a data I-SID with the multicast stream and a scope I-SID that defines the scope as Layer 3 VSN. A multicast stream with a scope of Layer 3 VSN can only transmit a multicast stream for the same Layer 3 VSN.
- For IP Shortcuts with IP multicast over SPBM, the BEB associates a data I-SID with the multicast stream and defines the scope as Layer 3 GRT. A multicast stream with a scope of Layer 3 GRT can only transmit a multicast stream for a Layer 3 GRT.

Note:

I-SID configuration is required only for virtual services such as Layer 2 VSN and Layer 3 VSN. With IP Shortcuts with unicast, no I-SID is required, forwarding for the Global Routing table is done using IS-IS based shortest path B-MAC reachability.

BCBs and BEBs

The boundary between the core MAC-in-MAC SPBM domain and the edge customer 802.1Q domain is handled by Backbone Edge Bridges (BEBs). I-SIDs are provisioned on the BEBs to be associated with a particular service instance.

In the SPBM core, the bridges are referred to as Backbone Core Bridges (BCBs). BCBs forward encapsulated traffic based on the BMAC-DA.

Important:

SPBM separates the payload from the transport over the SPBM infrastructure. Configure all virtualization services on the BEBs at the edge of the network. There is no provisioning required on the core SPBM switches. This provides a robust carrier grade architecture where configuration on the core switches never needs to be touched when adding new services.

A BEB performs the same functionality as a BCB, but it also terminates one or more Virtual Service Networks (VSNs). A BCB does not terminate any VSNs and is unaware of the VSN traffic it transports. A BCB simply knows how to reach any other BEB in the SPBM backbone.

VLANs without member ports

If a VLAN is attached to an I-SID there must be another instance of that same I-SID in the SPBM network.

- If another instance of that I-SID exists, the device designates that VLAN as operationally up regardless of whether it has a member port or not.

  When the VLAN is operationally up, the IP address of the VLAN will be in the routing table.

- If no matching instance of the I-SID exists in the SPBM network, then that VLAN has no reachable members and does not act as an NNI interface.

  The VLAN does not act as a UNI interface because it does not have a member port.

  Therefore, the device does not designate the VLAN as operationally up because the VLAN does not act as a UNI or an NNI interface.

If the device acts as a BCB with two VLANs configured and two I-SIDs, there must be a UNI side with the corresponding I-SID existing in the network.

If the device acts as both BEB and BCB, then there must be a member port in that VLAN to push out the UNI traffic.

Basic SPBM network topology

The following figure shows a basic SPBM network topology, specifically a Layer 2 VSN. Switches A and D are the Backbone Edge Bridges (BEB) that provide the boundary between the customer VLANs (C-VLAN) and the Backbone. Switches B and C are the Backbone Core Bridges (BCB) that form the core of the SPBM network.

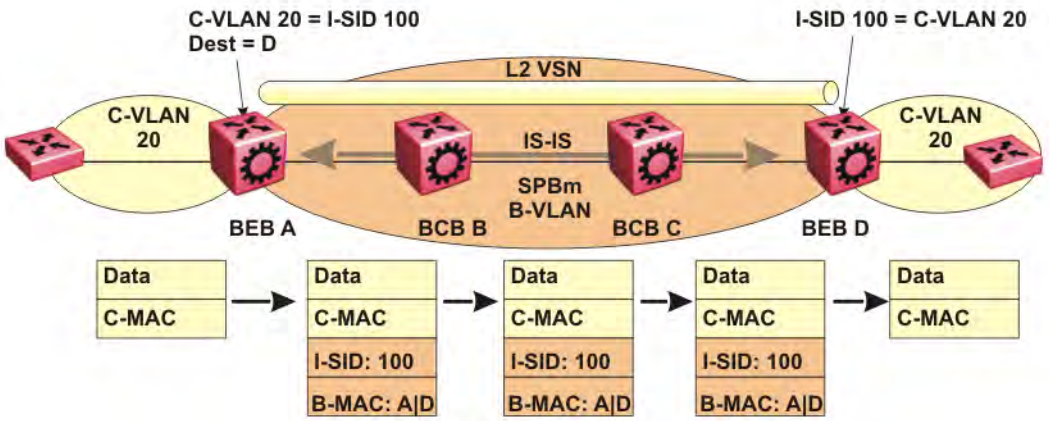

Figure 1: SPBM L2 VSN

SPBM uses IS-IS in the core so that all BEBs and BCBs learn the IS-IS System-ID (B-MAC) of every other switch in the network. For example, BEB-A uses IS-IS to build an SPBM unicast forwarding table containing the B-MAC of switches BCB-B, BCB-C, and BEB-D.

The BEBs provide the boundary between the SPBM domain and the virtualized services domain. For a Layer 2 VSN service, the BEBs map a C-VLAN to an I-SID based on local service provisioning. Any BEB in the network that has the same I-SID configured can participate in the same Layer 2 VSN.

In this example, BEB A and BEB D are provisioned to associate C-VLAN 20 with I-SID 100. When BEB A receives traffic from C-VLAN 20 that must be forwarded to the far-end location, it performs a lookup and determines that C-VLAN 20 is associated with I-SID 100 and that BEB D is the destination for I-SID 100. BEB A then encapsulates the data and C-MAC header into a new B-MAC header, using its own nodal B-MAC: A as the source address and B-MAC: D as the destination address. BEB A then forwards the encapsulated traffic to BCB B.

To forward traffic in the core toward the destination node D, BCB B and BCB C perform Ethernet switching using the B-MAC information only.

At BEB D, the node strips off the B-MAC encapsulation, and performs a lookup to determine the destination for traffic with I-SID 100. BEB D identifies the destination on the C-VLAN header as C-VLAN 20 and forwards the packet to the appropriate destination VLAN and port.

IS-IS

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 (IS-IS). IS-IS provides virtualization services, both Layer 2 and Layer 3, using a pure Ethernet technology base. SPBM also uses IS-IS to discover and advertise the network topology, which enables it to compute the shortest path to all nodes in the SPBM network.

IS-IS is a link-state, interior gateway protocol that was developed for the International Organization for Standardization (ISO). ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System-to-Intermediate System (IS-IS).

To provide a loop-free network and to learn and distribute network information, SPBM uses the Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol. IS-IS is designed to find the shortest path from any one destination to any other in a dynamic fashion. IS-IS creates any-to-any connectivity in a network in an optimized, loop-free manner, without the long convergence delay experienced with the Spanning Tree Protocol. IS-IS does not block ports from use, but rather employs a specific path. As such, all links are available for use.

IS-IS dynamically learns the topology of a network and constructs unicast and multicast mesh connectivity. Each node in the network calculates a shortest-path tree to every other network node based on System-IDs (B-MAC addresses).

Unlike in an IP Open Shortest Path First (OSPF) environment, the SPBM use of IS-IS does not require transport of any IP addressing for topology calculations. In the SPBM environment for Layer 2 VSNs, IS-IS carries only pure Layer 2 information with no requirement for an underlying IP control plane or forwarding path. IS-IS runs directly over Layer 2.

Note:

SPBM carries Layer 3 information for Layer 3 VSNs.

In SPBM networks, IS-IS performs the following functions:

- Discovers the network topology
- Builds shortest path trees between the network nodes:
  - Forwards unicast traffic
  - Determines the forwarding table for multicast traffic
- Communicates network information in the control plane:
  - Service Instance Identifier (I-SID) information

SPBM can distribute I-SID service information to all SPBM nodes, as the I-SIDs are created. SPBM includes I-SID information in the IS-IS Link State protocol data units (PDUs). When a new service instance is provisioned on a node, its membership is flooded throughout the topology using an IS-IS advertisement.

Standard TLVs

IS-IS uses Type-Length-Value (TLV) encoding. SPBM employs IS-IS as the interior gateway protocol and implements additional TLVs to support additional functionality. Virtual Services Platform 9000 also supports Sub-TLVs. TLVs exist inside IS-IS packets and Sub-TLVs exist as additional information in TLVs.

Virtual Services Platform 9000 supports standard 802.1 aq TLVs. The IEEE ratified the 802.1aq standard that defines SPBM and the Type-Length-Value (TLV) encoding that IS-IS uses to support SPBM services. Avaya is in full compliance with the IEEE 802.1aq standard. The following table lists a the TLVs that VSP 9000 supports.

Figure 2: Standard TLVs

TLV	Description	Usage
1	Area addresses — The Area Addresses TLV contains the area addresses to which the IS-IS is connected.	IS-IS area
22	Extended IS reachability — The Extended IS Reachability TLV contains information about adjacent neighbors.	IS-IS adjacencies Sub-TLV 29: SPBM link metric is carried within this TLV.
129	Protocols supported — The Protocol supported TLV carries the Network Layer Protocol Identifiers (NLPID) for the Network Layer protocols where the IS-IS can be used.	SPBM in addition to existing NLPID (IPV4 0xCC, IPV6 0x*E..), IEEE 802.1aq defined SPBM NLPID as 0xC1.
135	Extended IP reachability — The Extended IP Reachability TLV 135 is used to distribution IP reachability between IS-IS peers.	SPBM uses this existing IS-IS TLV to carry IP Shortcut routes in the Global Routing Table (GRT).
143	Multi-topology port aware capability (MT-Port-Capability) TLV This TLV carries the SPB instance ID in a multiple SPB instances environment. This TLV is carried within IS-IS Hello Packets (IIH).	This TLV carries the following SPBM Sub TLV: Sub-TLV 6: SPB B-VID Sub TLV indicates the mapping between a VLAN and its equal cost tree (ECT) algorithm. To form an adjacency, both nodes must have a matching primary (BVLAN, ECT) pair, and secondary (BVLAN, ECT) pair, the number of B-VLANs must be equal, B-VLAN values must match, ECT values for the B-VLANs must match. Used in IS-IS Hellos only.

TLV	Description	Usage
144	Multi-topology Capability (MT-Capability) TLV. This TLV carries the SPB instance ID in a multiple SPB instance environment. This TLV is carried within LSPs. In multicast over SPBM, TLV 144 on the BEB bridge, where the sender is located, has the transmit (Tx) bit set. On the BEB bridge, where the receiver is located the receive (Rx) bit is set.	TLV 144 is the service identifier TLV. TLV 144 advertizes B-MAC and I-SID information. This TLV carries the following Sub TLVs: Sub-TLV 1: SPB instance Sub TLV contains a unique SPSourceID (nickname) to identify the SPBM node within this SPB topology. Sub-TLV 3: SPB Service ID (I-SID) is stored in TLV 144 sub-TLV 3. Sub-TLV 3 carries service group membership (I-SIDs) for a particular SPBM B-VLAN.
184	SPBM IP VPN reachability — IS-IS TLV 184 is used to advertise SPBM L3 VSN route information across the SPBM cloud.	IP reachability for Layer 3 VSNs
185	IPVPN multicast TLV with IPMC sub TLV — The IPVPN multicast TLV contains information about the scope I-SID.	TLV 185 on the BEB bridge, where the source is located, displays the multicast source and group addresses and has the transmit (Tx) bit set. Each multicast group has its own data I-SID that maps to the source and group addresses. As part of the IPVPN TLV, sub-TLVs define IPv4 unicast, IPv6 unicast and IPv4 multicast information. Layer 2 VSN IP multicast over SPBM and Layer 3 VSN IP multicast over SPBM (using VRF) use TLV 185.
186	IP multicast TLV (GRT) — TLV 186 on the BEB bridge, where the source is located, displays the multicast source and group addresses and has the transmit (Tx) bit set. Each multicast group has its own data I-SID that maps to the source and group addresses.	IP Shortcuts with IP multicast over SPBM use TLV 186. All multicast streams are constrained within the level in which they originate, which is called the scope level.

IS-IS hierarchies

IS-IS is a dynamic routing protocol that operates within an autonomous system (or domain). IS-IS provides support for hierarchical routing, which enables you to partition large routing domains into smaller areas. When used separately from SPBM, IS-IS uses a two-level hierarchy, dividing the domain into multiple Level 1 areas and one Level 2 area. When used separately from SPBM, the Level 2 area serves as backbone of the domain, connecting to all the Level 1 areas. SPBM currently uses only Level 1 areas.

Important:

The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled in the current release.

IS-IS PDUs

Intermediate System to Intermediate System Hello (IIH) packets discover IS-IS neighbors and establish and maintain IS-IS adjacencies. An IIH is sent in every Hello-interval to maintain the established adjacency. If a node has not heard IIHs from its neighbor for (hello-interval x hello-multiple) seconds, the node tears down the adjacency. In the current release, IIH carries TLV 143 and SPB-B-VLAN Sub-TLV (among other sub-TLVs). For two nodes to form an adjacency the B-VLAN pairs for primary B-VLAN and secondary B-VLAN must match.

Link State Packets (LSP) advertise link state information. The system uses the link state information to compute the shortest path. LSP also advertises MT-capability TLV 144 and SPB instance Sub-TLV, and SPB I-SIDs Sub-TLV.

Complete Sequence Number Packets (CSNP) contain the most recent sequence numbers of all LSPs in the database. CSNP notifies neighbors about the local LSDB. After a neighbor receives a CSNP, it compares the LSPs in the CSNP with the LSP in the local LSDB. If the neighbor is missing LSPs, it sends a Partial Sequence Number Packets (PSNP) to request the missing LSPs. This process synchronizes the LSDBs among neighbors. A synchronized LSDB among all nodes in the network is crucial to producing a loop-free shortest path.

IS-IS configuration parameters

IS-IS system identifiers

The IS-IS system identifiers consist of three parts:

- System ID — The system ID is any 6 bytes that are unique in a given area or level. The system ID defaults to the baseMacAddress of the chassis but you can configure a non-default value. The system ID must use a unicast MAC address; do not use a multicast MAC address. A MAC address that has the low order bit 1 set in the highest byte is a multicast MAC address. For example, the following are multicast MAC addresses: x1xx.xxxx.xxxx, x3xx.xxxx.xxxx, x5xx.xxxx.xxxx, x7xx.xxxx.xxxx, x9xx.xxxx.xxxx, xBxx.xxxx.xxxx, xDxx.xxxx.xxxx, and xFxx.xxxx.xxxx.

- Manual area — The manual area or area ID is up to 13 bytes long. The first byte of the area number (for example, 49) is the Authority and Format Indicator (AFI). The next bytes are the assigned domain (area) identifier, which is up to 12 bytes (for example, 49.0102.0304.0506.0708.0910.1112). IS-IS supports a maximum of three manual areas, but the current VSP 9000 release only supports one manual area.
- NSEL — The last byte (00) is the n-selector. In the Avaya Virtual Services 9000 implementation, this part is automatically attached. There is no user input accepted.

The Network Entity Title (NET) is the combination of all three global parameters.

All routers have at least one manual area. Typically, a Level 1 router does not participate in more than one area.

The following are the requirements for system IDs:

- All IS-IS enabled routers must have one manual area and a unique system ID.
- All routers in the same area must have the same area ID.
- All routers must have system IDs of the same length (6 bytes).
- All IS-IS enabled routers must have a unique nickname.

PSNP interval

You can change the PSNP interval rate. A longer interval reduces overhead, while a shorter interval speeds up convergence.

CSNP periodic and interval rate

You can configure the CSNP periodic and interval rate. A longer interval reduces overhead, while a shorter interval speeds up convergence.

Parameters for the link state packet (LSP)

LSPs contain vital information about the state of adjacencies, which must be exchanged with neighboring IS-IS systems. Routers periodically flood LSPs throughout an area to maintain synchronization. You can configure the LSP to reduce overhead or speed up convergence.

The following list describes IS-IS parameters related to LSPs:

- The **`max-lsp-gen-interval`** is the time interval at which the generated LSP is refreshed. The default is 900 seconds with a range of 30 to 900.
- The **`retransmit-lspint`** is the minimum amount of time between retransmission of an LSP. When transmitting or flooding an LSP an acknowledgement (ACK) is expected. If the ack is not received within **`retransmit-lspint`**, the LSP is re-transmitted. The default is 5 seconds with a range of 1 to 300.

Point-to-point mode

All SPBM links are point-to-point links. Virtual Services Platform 9000 does not support broadcast links.

IS-IS interface authentication

Configure IS-IS interface authentication to improve security and to guarantee that only trusted routers are included in the IS-IS network. Interface level authentication only checks the IIH PDUs. If the authentication type or key in a received IIH does not match the locally-configured type and key, the IIH is rejected. By default, authentication is disabled.

You can use either one of the following authentication methods:

- Simple password authentication — Uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.
- MD5 authentication — Creates a Message Digest (MD5) key.

Password considerations

The passwords for all authentications are saved as cleartext in the configuration file on the Avaya Virtual Services Platform 9000. The passwords for simple and HMAC-MD5 are displayed in cleartext through ACLI. The HMAC-MD5 packet is encrypted when transmitted over the network.

To reset the authentication password type, you must set the type to none.

The current release supports only interface level authentication. The current release does not support area level or domain level authentication.

Hellos

To update the identities of neighboring routers, you can configure the:

- Interface Hello interval
- Interface Hello multiplier

Interface Hello interval

IS-IS uses Hello packets to initialize and maintain adjacencies between neighboring routers.

You can configure the interface level Hello interval to change how often Hello packets are sent out from an interface level.

Hello multiplier

You can configure the Hello multiplier to specify how many Hellos the Avaya Virtual Services Platform 9000 must miss before it considers the adjacency with a neighboring switch down. By default, the hold (wait) time is the Hello interval multiplied by the Hello multiplier. By default, if the Hello interval is 9 and the Hello multiplier is 3, the hold time is 27. If the Hello multiplier is increased to 10, the hold time is increased to 90.

Link metric

You can configure the link metric to overwrite the default metric value. By configuring the metric, you can specify a preferred path. Low cost reflects high-speed media, and high cost reflects slower media. For the wide metric, the value ranges from 1 to 16,777,215.

In this release, only the wide metric is supported.

The total cost of a path equals the sum of the cost of each link.

The default value for wide metrics is 10.

Disabling IS-IS

You can disable IS-IS globally or at the interface level. If IS-IS is globally disabled, then all IS-IS functions stop. If IS-IS is enabled at the global level and disabled at one of the interface levels, then IS-IS continues on all other interfaces.

Overload bit

If the overload bit parameter is configured, the Avaya Virtual Services Platform 9000 sets the overload bit in its LSP. The **`overload`** parameter works in conjunction with the **`overload-on-`**

`startup` parameter. When the **`overload-on-startup`** timer expires, the SPBM node clears the overload bit and re-advertises its LSP.

When an LSP with an overload bit is received, the Avaya Virtual Services Platform 9000 ignores the LSP in its SPF calculation so that the transit traffic will not go through the overloaded node. The overloaded node can still receive traffic that is destined for the node itself. If you set the overload bit, the device does not consider the node for use as a transit node in IS-IS computations. By default, overload is set to false.

SPBM B-VLAN

Each SPBM network instance is associated with at least one backbone VLAN (B-VLAN) in the core SPBM network.

Note:

SPB internally uses spanning tree group (STG) 63 or Multiple Spanning Tree Instance (MSTI) 62. STG 63 or MSTI 62 cannot be used by another VLAN or MSTI. For non-SPB customer networks, if you use STG 63 or MSTI 62 in the configuration, you must delete STG 63 or MSTI 62 before you can configure SPBM.

This VLAN is used for both control plane traffic and dataplane traffic.

Note:

Avaya recommends to always configure two B-VLANs in the core to allow load distribution over both B-VLANs.

SPBM alters the behavior of the VLAN. When a B-VLAN is associated with an SPBM network the following VLAN attributes and behaviors are modified for the B-VLAN:

- Flooding is disabled
- Broadcasting is disabled
- Source address learning is disabled
- Unknown MAC discard is enabled

Ports cannot be added to a B-VLAN manually, IS-IS takes care of adding ports to the B-VLAN. Essentially the B-MAC addresses are programmed into the B-VLAN Forwarding Information Bases (FIBs) by IS-IS instead of the traditional VLANs flooding and learning approach. Modification of the VLAN behavior is necessary to ensure proper control over the SPBM traffic.

Pre-populated FIB

An Ethernet network usually learns MAC addresses as frames are sent through the switch. This process is called reverse learning and is accomplished through broadcast.

SPBM does not allow any broadcast flooding of traffic on the B-VLAN in order to prevent looping accomplished through flooding packets with unknown destinations (although multicast traffic is supported). As such, MAC addresses must be distributed within SPBM. This is accomplished by carrying the necessary B-MAC addresses inside the IS-IS link state database. To that end, SPBM

supports an IS-IS TLV that advertises the I-SID and B-MAC information across the network. This functionality enables the powerful end-point-provisioning of SPBM.

These Backbone MAC addresses are populated into the SPBM VLAN Forwarding Information Base (FIB) to maximize efficiency and to allow Reverse Path Forwarding Check (RPFC) to operate properly.

RPFC

A loop prevention mechanism is required at Layer 2 to stop wayward traffic from crippling the network. Reverse Path Forwarding Check (RPFC) is the chosen method of suppressing loop traffic with SPBM. RPFC was originally designed for IP traffic at Layer 3 where it checks the source address of the packet against the routing entry in the routing table. The source address must match the route for the port it came in on otherwise the packet is illegitimate and therefore dropped.

With SPBM, the node matches the source B-MAC address against the ingress port to establish validity. If the frame is not supposed to come in that port, it is immediately suppressed imposing a guaranteed loop control. If there is no VLAN FDB entry to the source MAC address with the outgoing port as the ingress port, the frame will be dropped.

SPBM FIB

This section describes the SPBM unicast and multicast FIBs.

Unicast FIB

The unicast computation runs a single Dijkstra (unlike all pair Dijkstras for multicast). SPBM produces only one Shortest Path First (SPF) tree and the tree is rooted on the computing node.

The unicast computation generates an entry for each node in the network. The Destination Address (DA) for that entry is the system-id of the node. In addition, if a node advertises MAC addresses other than the system-id, each MAC address has an entry in the unicast FIB table, and the shortest path to that MAC should be exactly the same as the path to the node.

Unicast FIB entries are installed to the vlan-fdb table.

The following text shows an example of the unicast FIB.

```
VSP-9012:1# show isis spbm unicast-fib
================================================================================
                          SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION        BVLAN  SYSID           HOST-NAME         OUTGOING     COST
ADDRESS                                   NAME              INTERFACE
--------------------------------------------------------------------------------
00:80:2d:35:93:df  10     0080.2d35.93df  86-10             MLT-32       0
00:80:2d:35:93:df  11     0080.2d35.93df  86-10             MLT-32       0
00:80:86:10:86:20  11     0080.2d35.93df  86-10             MLT-32       0
00:e0:7b:84:57:df  10     00e0.7b84.57df  86-30             2/12         0
00:e0:7b:84:57:df  11     00e0.7b84.57df  86-30             2/12         0
```

Multicast FIB

SPBM runs all pair Dijkstras to produce the multicast FIB. The computing node loops through each node to run Dijkstra using that node as the root, and then prunes paths to only keep the shortest paths. The computing node then computes the intersection of the set of I-SIDs for which the root node transmits, with the set of I-SIDs for which the path endpoints receive.

The multicast addresses are built out of two pieces: the instance-ID (nickname) and the I-SID ID converted to hexadecimal format to form the multicast MAC address.

```
|-----------3 bytes ------------|---------------------------------|
       nickname & 3                    hexadecimal I-SID
```

For example, if the nickname is 0.00.10 and the I-SID is 100 (0x64), the multicast address is 03:00:10:00:00:64.

The following text shows an example of the multicast FIB.

```
Switch:1#show isis spbm multicast-fib

================================================================================

              SPBM MULTICAST FIB ENTRY INFO
================================================================================
MCAST DA          ISID  BVLAN  SYSID          HOST-NAME OUTGOING-INTERFACES INCOMING
                                                                    INTERFACE
--------------------------------------------------------------------------------
03:00:61:00:00:64 100    10    0080.2dc1.37ce 9000-1       4/7             5/7
03:00:61:00:00:c8 200    10    0080.2dc1.37ce 9000-1       4/2,4/1         5/2

--------------------------------------------------------------------------------
 Total number of SPBM MULTICAST FIB entries 2
--------------------------------------------------------------------------------
```

SPBM packet drop statistics

This enhanced packet drop feature keeps statistics on a per-port basis. The SPBM Drop Stats by Port tables display information on the source (SA) or destination (DA) MAC address of the last packet dropped.

This feature collects and displays frame drops on ingress at SPBM NNI interfaces in the following categories:

- Last drop

  Displays information for the last packets dropped.

- Reverse Path Forwarding Check (RPFC) multicast SA drops

  Displays the total number of SPBM RPFC multicast packets dropped. These drops occur when packets with a specific source MAC address ingress on a port different than what IS-IS expected.

- RPFC unicast SA drops

  Displays the total number of SPBM RPFC unicast packets dropped. These drops occur when packets with a specific source MAC address ingress on a port different than what IS-IS expected.

- Unknown unicast DA drops

Displays the total number of SPBM unknown unicast DA packets dropped. These drops occur when the unicast destination MAC address of the packet is not known.

- Unknown unicast SA drops

  Displays the total number of SPBM unknown unicast SA packets dropped. These drops occur when the unicast source MAC address of the packet is not known. VSP 9000 does not support the unknown-unicast-sa drop count parameter for second generation modules in this release. The device always displays the ACLI output for second generation modules as 0 for this counter.

- Unknown multicast DA drops

  Displays the total number of SPBM unknown multicast DA packets dropped. These drops occur when the multicast destination MAC address of the packet is not known.

SPBM restrictions and limitations

This section describes the restrictions and limitations associated with SPBM.

RSTP and MSTP

The following list identifies restrictions and limitations associated with RSTP and MSTP:

- RSTP mode does not support SPBM.
- A C-VLAN-level loop across SPBM NNI ports cannot be detected and needs to be resolved at the provisional level.
- SPBM NNI ports are not part of the Layer 2 VSN C-VLAN, and BPDUs are not transmitted over the SPBM tunnel. SPBM can only guarantee loop-free topologies consisting of the NNI ports. Avaya recommends that you always use Simple Loop Prevention Protocol (SLPP) in an SMLT environment.

  Note:

  Avaya recommends that you deploy SLPP on C-VLANs to detect loops created by customers in their access networks. However, SLPP is not required on B-VLANs, and it is not supported. The B-VLAN active topology is controlled by IS-IS that has loop mitigation and prevention capabilities built into the protocol.

- SPB internally uses Multiple Spanning Tree Instance (MSTI) 62. MSTI 62 cannot be used by another VLAN or MSTI. For non-SPB customer networks, if you use MSTI 62 in the configuration, you must delete MSTI 62 before you can configure SPBM.
- You must configure SPBM B-VLANs on all devices in the same MSTP region. MSTP requires this configuration to generate the correct digest.

SPBM IS-IS

The following list identifies restrictions and limitations associated with SPBM IS-IS:

- The current release does not support IP over IS-IS as defined by RFC 1195. IS-IS protocol is only to facilitate SPBM.

- The current release uses level 1 IS-IS. The current release does not support level 2 IS-IS. The ACLI command **show isis int-l2-contl-pkts** is not supported in the current release because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.
- The IS-IS standard defines wide (32bit) metrics and narrow (8 bits) metrics. The current release supports the wide metric.
- SPBM supports full High Availability (HA). The SPBM and IS-IS configuration and dynamic information (such as adjacencies and LSPs) are all synchronized to the standby CPU to ensure seamless switchover. Because the HA framework cannot guarantee seamless switchover, there is a 6 to 7 seconds gap between the master CPU going down and the backup CPU coming up. To avoid IS-IS adjacencies bouncing during the switchover, the recommended hello interval is 9 seconds and the hello multiple is 3.

Pay special attention to the expected scaling of routes in the network when you select configuration values for the **isis l1-hello-interval** and **isis l1-hello-multiplier** commands on IS-IS interfaces. The default values for these commands work well for most networks, including those using moderately-scaled routes. In highly-scaled networks, you may need to configure higher values for these commands. For example, if the total number of non IS-IS routes on a given BEB exceeds 25,000 in combination with approximately 60,000 IS-IS routes that the BEB receives from other BEBs in the network, Avaya recommends that you configure a value of 12 for **isis l1-hello-multiplier**, rather than use the default value of 3.

SPBM NNI SMLT

For NNI-facing SMLT, the current release supports only one link between each IS-IS adjacency pair.

The following modules support NNI ports:

- 9024XL I/O module
- 9048XS-2 I/O module
- 9012QQ-2 I/O module

Internet Protocol Flow Information eXport

The switch does not support Internet Protocol Flow Information eXport (IPFIX) on Intermediate-System-to-Intermediate-System (IS-IS) interfaces.

VLACP

VLACP is generally used when a repeater or switch exists between connected Virtual Services Platform 9000 switches to detect when a connection is down even when the link LED is lit. If you configure VLACP on an SPBM link that is also an IST link, during a connection fail over (where the link LED stays lit) the IS-IS hellos time out first (after 27 seconds, using default values) and take down the IS-IS adjacency. IS-IS then calculates the new shortest path and fails over the SPBM traffic. 90 seconds after the connection failure (using default values), VLACP goes down but the IST link was already taken down by IS-IS.

In this scenario, there is no data traffic impact because IS-IS can find another path in the SPBM network before VLACP goes down.

SNMP traps

On each SPBM peer, if you configure the SPBM B-VLANs to use different VLAN IDs, for example, VLAN 10 and 20 on one switch, and VLAN 30 and 40 on the second, the system does not generate a trap message to alert of the mismatch because the two switches cannot receive control packets from one another. Configure the SPBM B-VLANs to use matching VLAN IDs.

Other

The following list identifies other restrictions and limitations:

- The current release does not support I-SID filters.
- You cannot enable C-VLAN and B-VLAN on the same port, except the IST ports.
- The current release supports NNI only on 10 GbE interface modules.

Legacy IS-IS

An SPBM node can form an adjacency with a legacy IS-IS router. This adjacency means you can introduce SPBM into existing networks and provide easy migration.

IP multicast over SPBM

Avaya leads the industry with a new approach to transporting IP multicast using SPBM. IP multicast over SPBM greatly simplifies multicast deployment, with no need for any multicast routing protocols such as Protocol Independent Multicast-Sparse Mode (PIM-SM) or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM). A BEB can forward a multicast stream anywhere in an SPBM network where IS-IS advertises the stream to the rest of the fabric.

The advantage of this solution over traditional approaches is the simplicity in provisioning and deploying IP multicast bridging and routing. Also, due to the fact that only one control plane protocol (IS-IS) exists, convergence times in the event of a network failure, are typically sub second.

You can compare the quick convergence times for IP multicast over SPBM to Interior Gateway Protocols like Open Shortest Path First (OSPF) combined with PIM-SM or PIM-SSM. OSPF combined with PIM-SM or PIM-SSM can have recovery times that are sub optimal with convergence times that take tens of seconds. PIM experiences longer convergence times, in part, because unicast IP routing protocols must converge before PIM can converge. PIM also maintains the network state for every multicast group and uses a mechanism based on each hop to update the network about state changes, which affects scalability.

IP multicast over SPBM is extremely scalable because you only apply the multicast bridging and routing functionality at the SPBM fabric edge, with the streams mapped to SPBM multicast trees in the fabric.

With IP multicast over SPBM, Avaya introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. IP multicast over SPBM uses these extensions, along with the Internet Group Management Protocol (IGMP) Snooping and Querier functions at the edge of the SPBM cloud, to create sub-trees of the VSN SPB for each multicast group to transport IP multicast data.

With IP multicast over SPBM, the switch supports the following:

- Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network (Layer 2 VSN with IP multicast over SPBM). Example application: Multicast in data centers.
- IP multicast routing support for IP Shortcuts using SPBM in the core and IGMP on the access (IP Shortcuts with IP multicast over SPBM). Example applications: Video surveillance, TV/Video/Ticker/Image distribution, VX-LAN.
- Layer 3 Virtual Services Network with VRF based IP multicast routing support over SPBM in the core and IGMP on the access (Layer 3 VSN with IP multicast over SPBM). Example

applications: Video surveillance, TV/Video/Ticker/Image Distribution, VX-LAN, Multi-tenant IP multicast.

How IP multicast over SPBM works

The BEBs act as the boundary between the multicast domain (currently only IGMP dynamic or static) and the SPBM domain. Multicast senders (sources) and receivers connect directly or indirectly (using Layer 2 switches) to the BEBs. You can enable SPBM multicast services at the Layer 2 VSN level or the Layer 3 VSN level (including the GRT).

The following figure shows how multicast senders and receivers connect to the SPBM cloud using BEBs.

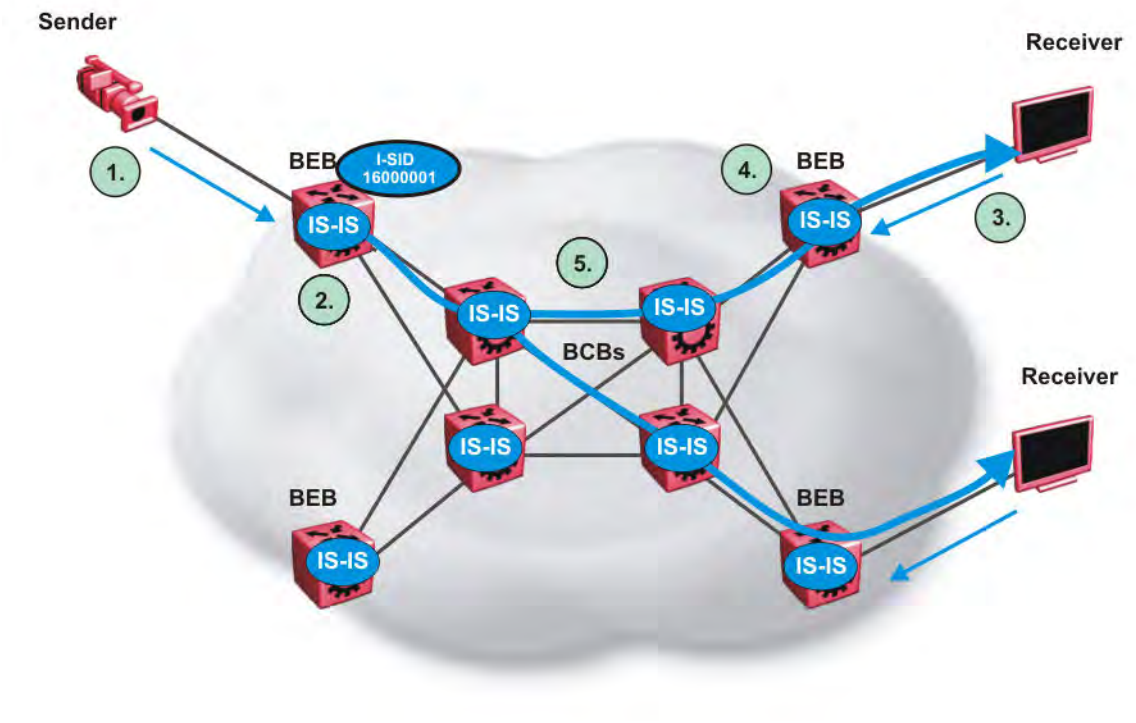

Figure 3: IP multicast over SPBM streams

The following list describes how multicast senders and receivers connect to the SPBM cloud using BEBs in the preceding diagram:

1. The sender transmits multicast traffic with group IP address 233.252.0.1.
2. After the BEB receives the IP multicast stream from the sender, the BEB allocates data I-SID 16000001 for the S,G multicast stream. The BEB sends an LSP with the TLV 185 (for Layer 2 VSN multicast and Layer 3 VSN multicast) or TLV 186 (for IP Shortcuts multicast) with the transmit bit set. The BEB also sends an IS-IS service identifier and unicast address sub-TLV (where the unicast address has the multicast bit set and the I-SID is the Data I-SID).
3. The receiver sends a join request to Group 233.252.0.1.

4. The BEB (acting as the IGMP Querier) queries the IS-IS database to find all senders for group 233.252.0.1. If the group exists, the BEB sends an LSP with the IS-IS service identifier and unicast address sub-TLV (where the unicast address has the multicast bit set and the nickname is the stream transmitter BEB and the I-SID is the data I-SID).
5. The multicast tree is calculated for the data I-SID and the data starts flowing from the sender.

Scope level

IP multicast over SPBM constrains all multicast streams within the level in which they originate, which is called the scope level. In other words, if a sender transmits a multicast stream to a BEB on a C-VLAN (a VLAN that is mapped to an I-SID, for instance, a L2 VSN) with IP multicast over SPBM enabled, only receivers that are part of the same Layer 2 VSN can receive that stream. Similarly, if a sender transmits a multicast stream to a BEB on a VLAN that is part of the GRT or a Layer 3 VSN with IP multicast over SPBM enabled, only receivers that are part of the same Layer 3 instance (GRT or L3 VSN) can receive that stream.

Note:

In the context of SPBM multicast, scope is either the Global Routing Table or the I-SID value of the Layer 2 or Layer 3 VSN associated with the local VLAN on which the IP multicast data was received.

Data I-SID

After the BEB receives the IP multicast stream from the sender, a BEB allocates a data Service Identifier (I-SID) in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the S,G, V tuple, which is the source IP address, the group IP address, and the local VLAN the multicast stream is received on.

The BEB propagates this information through the SPBM cloud by using IS-IS TLV updates in LSPs, which results in the creation of a multicast tree for that stream. All BEBs now know what data I-SID to use for that stream and its scope. The data I-SID is a child of the scope or VSN I-SID. If no receiver requests the IP multicast stream, the ingress BEB does not forward the multicast stream.

IGMP

After a BEB receives an IGMP join message from a receiver, a BEB queries the IS-IS database to check if a sender exists for the requested stream within the scope of the receiver. If the requested stream does not exist, the IGMP information is kept, but no further action is taken. If the requested stream exists, the BEB sends an IS-IS TLV update to its neighbors to inform them of the presence of a receiver, and this information is propagated through the SPBM cloud.

IS-IS acts dynamically using the TLV information it receives from BEBs that connect to the sender and the receivers to create a multicast tree between them. IS-IS creates very efficient multicast trees for the data I-SID allocated at the sender edge of the SPBM cloud to transport data between the sender and the receivers. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID. After IS-IS creates the multicast tree, the sender transports data to the receiver across the SPBM cloud using the data I-SID.

The trigger to send IS-IS updates to announce a multicast stream into the SPBM cloud is the multicast traffic arriving at the BEB. Because the BEB only interacts with IGMP and not PIM in this release, all multicast traffic must be drawn towards the BEB for the stream to be announced, which SPBM accomplishes by making the BEB an IGMP Querier. In a VLAN, the IGMP Querier sends out periodic IGMP queries.

Note:

The BEB must be the only IGMP Querier in the VLAN. If the BEB receives an IGMP query from any other device, it causes unexpected behavior, including traffic loss.

BEB as IGMP Querier

The BEB acts as the IGMP Querier and creates tables for links that need IP multicast streams. IGMP and IGMP Snooping cannot work without an IGMP Querier that sends out periodic IGMP queries.

In the current release, the BEB only interacts with IGMP messages and not PIM. All multicast traffic must enter the BEB for the data stream to be announced.

The BEB must be the only IGMP Querier in the VLAN. If the BEB receives an IGMP query from any other device, unexpected behavior results, including traffic loss.

The IGMP query message is an IP packet and requires a source IP address. However, Layer 2 IGMP Snooping with SPBM by default turns on the service without the configuration of an IP address on the VLAN. By default, the BEB sends an IGMP query message with an IP source address of 0.0.0.0. If there are interoperability issues with third party vendors as a result of the 0.0.0.0 IP address, then you can configure the querier address under IGMP, without having to configure an IP address for the Layer 2 VSN VLAN.

IGMP Snooping, operating on the Layer 2 VSN, listens to conversations between hosts and routers, and maintains a table for links that need IP multicast streams.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

For more concept and configuration information on IGMP, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504.

Network Load Balancing

SPBM supports Network Load Balancing (NLB) unicast. SPBM does not support NLB multicast or NLB multicast with IGMP.

NLB is a clustering technology available with Microsoft Windows 2000, Microsoft Windows 2003, Microsoft Windows 2008, and Microsoft Windows 2012 server family of operating systems. You can use NLB to share the workload among multiple clustering servers. NLB uses a distributed algorithm to load balance TCP/IP network traffic across a number of hosts, enhancing the scalability and availability of mission critical, IP based services, such as Web, VPN, streaming media, and firewalls. NLB also provides high availability by detecting host failures and automatically redistributing traffic to remaining operational hosts.

For more information on NLB, see *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500.

Switch clustering at the edge of the SPBM network

Typical customer deployments require redundancy all the way to the access side of the network. IP multicast over SPBM supports Avaya switch clustering, Split Multilink Trunking (SMLT) technology, at the edge of the SPBM fabric, providing redundancy to the access Layer 2 switch where you can attach multicast senders and receivers. Typical SPBM fabric deployments use two or more B-VLANs for Equal Cost Multipath (ECMP) and resiliency. For simplicity in understanding how the SPBM network works, assume that there are two B-VLANs (primary and secondary).

The following figure shows how multicast senders and receivers connect to the SPBM cloud using BEBs.

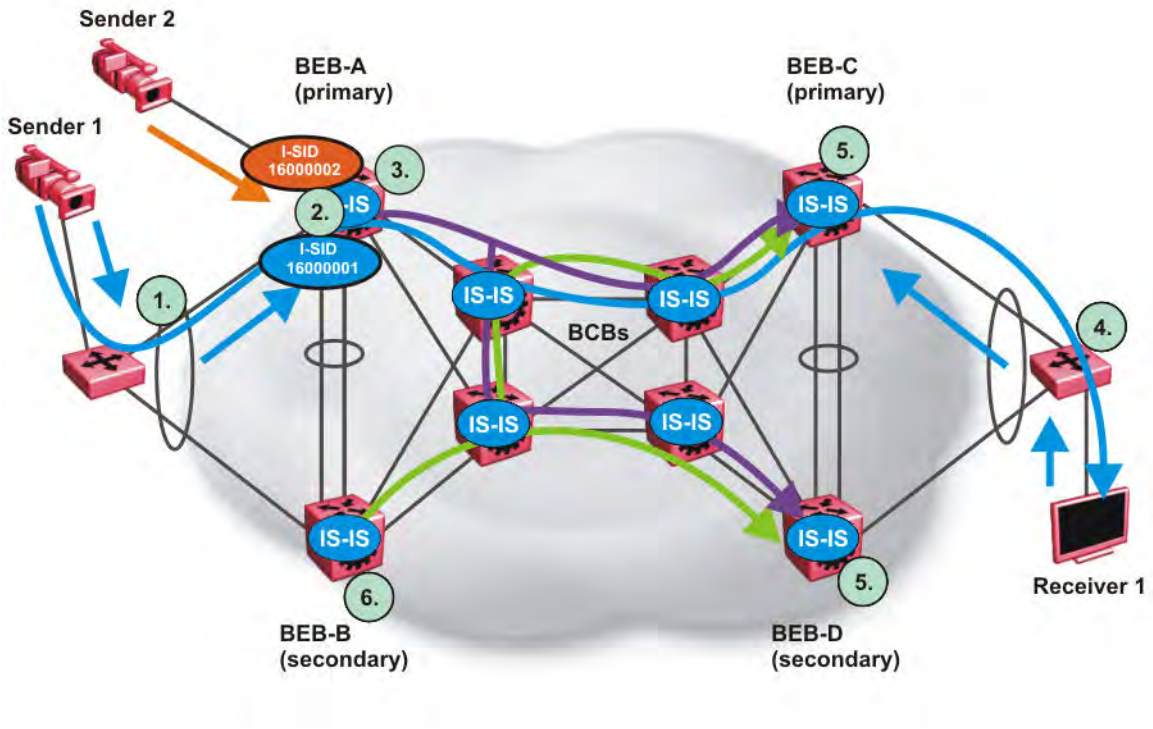

Figure 4: IP multicast over SPBM streams in an SMLT configuration

The following list describes the preceding diagram:

1. The edge switch hashes the sender multicast data to a specific MLT link.
2. A multicast stream received at the edge of the SPBM fabric is mapped to a dedicated multicast data I-SID.
3. For the non-SMLT attached sender 2, the stream is hashed to the primary or secondary B-VLAN based on whether the data I-SID is even or odd numbered. For the SMLT attached to sender 1, IS-IS advertises the stream to the rest of the fabric on the primary B-VLAN and synchronizes information to the IST peer.
4. The edge switch hashes the receiver IGMP join to a specific MLT link.
5. Both BEBs on both B-VIDs advertise the IGMP join.

6. The multicast tree is built for (S1,G1), which is rooted in the primary sender BEB. The multicast tree is built for (S1,G1), which is rooted in the secondary sender BEB.

IGMP Snooping is widely used on Layer 2 access switches to prune multicast traffic. In IP multicast over SPBM, BEBs are the IGMP Queriers, therefore access switches forward multicast data from the senders as well as IGMP control messages from receivers to the BEBs.

Multicast sender

When a sender transmits multicast data to the Layer 2 access switch that has an MLT to the switch cluster, it is hashed towards one or the other BEBs in the switch cluster. The receiving BEB allocates a data I-SID and sends a TLV update on either the primary B-VLAN or the secondary B-VLAN, depending on whether the BEB is the primary or secondary switch. The primary switch uses the primary B-VLAN, whereas, the secondary switch uses the secondary B-VLAN. This information is propagated through the SPBM fabric so all BEBs are aware of this stream availability.

The sender information is also synchronized over the IST to the peer switch. Then the peer switch allocates a data I-SID for the multicast stream and sends a TLV update on the appropriate B-VLAN to announce the availability of the stream. The data I-SIDs allocated by the primary and secondary switch cluster peers may be the same or different, as they are allocated independently by each switch.

Note:

If a sender attaches to only one BEB in a switch cluster, the sender information is not synchronized over the IST because it is not SMLT attached. The sender information is advertised, and data is sent on either the primary or secondary B-VLAN. The odd-numbered data I-SIDs use the primary B-VLAN, and the even-numbered data I-SIDs use the secondary B-VLAN. The same hashing rules apply to the forwarding of multicast data.

Multicast receiver

When a receiver sends an IGMP join message to the Layer 2 access switch that has an MLT to the switch cluster, it is hashed towards one or the other BEBs in the switch cluster. The receiving BEB queries the IS-IS Link State Database (LSDB) to check if a sender exists for the requested stream within the scope of the receiver.

If the requested stream does not exist, the BEB keeps the IGMP information but no further action is taken. If the requested stream exists, the BEB sends an IS-IS Link State Packet (LSP), with TLV update information, for both primary and secondary B-VLANs to its neighbors to inform them of the presence of a receiver. The BEB propagates this information through LSPs through the SPBM cloud. The receiver information is also synchronized over the IST to the peer switch. The peer switch then queries its IS-IS Link State Database (LSDB) and, if the requested stream exists, it sends an IS-IS LSP, with a TLV update, for both primary and secondary B-VLANs to its neighbors to inform them of the presence of the receiver.

IS-IS uses these TLV updates in LSPs to create multicast shortest path first trees in the SPBM fabric. IS-IS creates a shortest path first tree for the primary and secondary B-VLANs, but only one of the B-VLANs transports multicast data with the other in active standby in case of failures at the SPBM edge. After IS-IS creates the trees, multicast data flows between senders and receivers.

IP multicast over SPBM and SMLT

The following section summarizes the IP multicast over SPBM actions in an SMLT environment. The BEBs on the sender side behave as follows:

- Primary SMLT peer BEB always advertises the streams it receives, and sends data for them on the primary B-VLAN.
- Secondary SMLT peer BEB always advertises the streams it receives, and sends data for them on the secondary B-VLAN.
- Non-SMLT BEBs or SMLT BEBs with single attached senders advertise streams, and send data on the primary or secondary B-VLAN based on hash criteria (odd-numbered data I-SIDs use primary B-VLAN; even-numbered data I-SIDs use secondary B-VLAN).

The BEBs on the receiver side behave as follows:

- The primary SMLT peer BEB that receives multicast data on the primary B-VLAN sends it to both SMLT and non-SMLT SPBM access (UNI) links.
- The primary SMLT peer BEB that receives multicast data on the secondary B-VLAN sends it to non-SMLT SPBM access (UNI) links only.
- The secondary SMLT peer BEB that receives multicast data on primary B-VLAN sends it to non-SMLT SPBM access (UNI) links only.
- The secondary SMLT peer BEB that receives multicast data on secondary B-VLAN sends data to both SMLT and non-SMLT SPBM access (UNI) links.
- The non-SMLT BEB that receives multicast data on primary or secondary B-VLAN sends data to all SPBM access (UNI) links.

Layer 2 Querier behavior for a switch cluster

In VSP 9000 for C-VLANs in an SMLT environment, the IST ports are part of the VLAN.

In ERS 8800, VSP 8000, and VSP 4000, for C-VLANs in an SMLT environment, the vIST ports are not part of the VLAN.

IGMP on a C-VLAN behaves as follows to account for the fact that IST peers do not see the membership queries of each other:

- The IST peer with the higher IP address sends the queries out all SMLT and non-SMLT ports on SPBM access links.
- The IST peer with the lower IP address only sends out queries on its non-SMLT ports. This includes SMLT ports whose remote ports are down (SMLT state of ‘norm’).
- With the existence of an IST peer with a higher IP address and an IST peer with a lower IP address, it means two queriers exist within the C-VLAN. Having two queriers poses no problems in this SPB environment, as all SMLT access devices see the IST peer with the higher IP address as the querier, and non-SMLT access devices see the directly connected IST peer as the querier. Non-SMLT access devices that connect on either side of the IST peers can talk to each other using the SPBM cloud.

Considerations when you connect an IP multicast over SPBM network to a PIM network

The current implementation of IP multicast over SPBM does not integrate PIM functionality. Apply the following considerations when you connect to a PIM network:

- You must configure static IGMP receivers on the BEB access interface that faces the PIM network when the sender is on the SPBM access network and the receiver is on the PIM network.

  Note:

  The PIM router must have a configuration option to accept streams with non-local sources or the router drops the packets. The switch does not currently support a configuration option to accept streams with non-local sources.

  You must configure static IGMP receivers on the PIM interface that face the SPBM multicast network when the sender is on the PIM network and the receiver is on the SPBM access network.

  Note:

  For security reasons and to limit unnecessary multicast streams from being injected into the SPBM domain, you should configure ACLs on the BEB facing the PIM network.

IP multicast over SPBM restrictions

Review the following restrictions for the IP multicast over SPBM feature.

IGMP

The BEB must be the only IGMP querier in the network. If the BEB receives an IGMP query from any other device, it causes unpredictable behavior, including traffic loss.

SPBM supports IGMP Snooping on a C-VLAN, but it does not support PIM on a C-VLAN. If you enable IGMP Snooping on a C-VLAN, then its operating mode is Layer 2 VSN with IP multicast over SPBM.

SPBM supports Network Load Balancing (NLB) unicast. SPBM does not support NLB multicast or NLB multicast with IGMP.

You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is either the same as the IGMP version configured on the IGMP Snooping VLAN, or that compatibility mode is enabled.

SSM

If you delete any **ssm-map** in a static range group, the switch deletes the entire static range group. For example, create an ssm-map for 232.122.122.122 to 232.122.122.122.128 and after that configure this same range in a static group. If you delete any ssm-map between 232.122.122.122. to 232.122.122.128, the switch deletes the entire static range group.

PIM

There can be no interaction with PIM and multicast routers on the access.

The BEB only interacts with IGMP messages and not PIM, so all multicast traffic must be drawn towards the BEB, which acts as the IGMP querier, for the stream to be announced.

IP multicast over SPBM does not integrate PIM functionality so the following considerations apply when connecting to a PIM network:

- You must configure static IGMP receivers on the BEB access interface facing the PIM network when the sender is on the SPBM access network and the receiver is on the PIM network. Static IGMP receivers make the PIM router accept streams and avoid a Reverse Path Forwarding (RPF) check that can change the source of the stream.
- You must configure static IGMP receivers on the PIM interface facing the SPBM multicast network when the sender is on the PIM network and the receiver is on the SPBM access network.
- You must configure Access Control Lists (ACLs) on the BEB facing the PIM network for security.

Data I-SID

The BEB matches a single multicast stream to a particular data I-SID. As a result there is a one-to-one mapping between the S,G to data I-SID for each BEB.

IP address

IP multicast over SPBM only supports IPv4 multicast traffic in this release.

Supported services

The switch does not support IP multicast over SPBM routing on inter-VSN routing interfaces.

The switch supports the following modes of IP multicast over SPBM:

- Layer 2 VSN multicast service — Multicast traffic remains within the same Layer 2 VSN across the SPBM cloud.
- Layer 3 VSN multicast service — Multicast traffic remains within the same Layer 3 VSN across the SPBM cloud.
- IP Shortcuts multicast service — Multicast traffic can cross VLAN boundaries but remains confined to the subset of VLANs with the Global Routing Table that have IP multicast over SPBM enabled.

SPBM and IS-IS infrastructure configuration using ACLI

This section provides procedures to configure SPBM and IS-IS using Avaya Command Line Interface (ACLI).

Configuring minimum SPBM and IS-IS parameters

Use the following procedure to configure the minimum required SPBM and IS-IS parameters to allow SPBM to operate on the switch.

Note:

Virtual Services Platform 9000 does not support IPFIX on IS-IS interfaces.

Note:

SPB internally uses spanning tree group (STG) 63 or Multiple Spanning Tree Instance (MSTI) 62. STG 63 or MSTI 62 cannot be used by another VLAN or MSTI. For non-SPB customer networks, if you use STG 63 or MSTI 62 in the configuration, you must delete STG 63 or MSTI 62 before you can configure SPBM.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable SPBM globally:

   ```
   spbm
   ```

3. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

4. Create the SPBM instance (in this release, only one SPBM instance is supported):

   ```
   spbm <1-100>
   ```

5. Associate the SPBM B-VLAN to the SPBM instance:

   ```
   spbm <1-100> b-vid {<vlan-id [-vlan-id][,...]} [primary <1-4084>]
   ```

   Note:

   You must associate the SPBM B-VLAN to the SPBM instance before you create the SPBM B-VLAN.

6. Exit to Global Configuration mode:

```
exit
```

7. Create the SPBM backbone VLAN (B-VLAN):

   ```
   vlan create <2-4084> type spbm-bvlan
   ```

8. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

9. Configure the system nickname (2.5 bytes in the format <x.xx.xx>):

   ```
   spbm <1-100> nick-name <x.xx.xx>
   ```

   Note:

   Although it is not strictly required for SPBM operation, Avaya recommends that you change the IS-IS system ID from the default B-MAC value to a recognizable address to easily identify a switch (Log on to IS-IS Router configuration mode and use the `system-id <xxxx.xxxx.xxxx>` command) . This helps to recognize source and destination addresses for troubleshooting purposes.

10. Configure an IS-IS manual area (1-13 bytes in the format <xx.xxxx.xxxx...xxxx>. In this release, only one manual area is supported.):

    ```
    manual-area <xx.xxxx.xxxx...xxxx>
    ```

11. Exit to Global Configuration mode:

    ```
    exit
    ```

12. Enter Interface Configuration mode:

    ```
    enable
    configure terminal
    interface GigabitEthernet {slot/port[-slot/port][,...]}
    ```
    or `interface mlt <1-512>`

13. Create an IS-IS circuit and interface on the selected ports or MLTs:

    ```
    isis
    ```

14. Enable the SPBM instance on the IS-IS interfaces:

    ```
    isis spbm <1-100>
    ```

15. Enable the IS-IS circuit/interface on the selected ports or MLTs:

    ```
    isis enable
    ```

16. Exit to Global Configuration mode:

    ```
    exit
    ```

17. Enable IS-IS globally:

    ```
    router isis enable
    ```

18. Display the SPBM configurations:

    ```
    show isis spbm
    ```

19. Display the global IS-IS configuration:

    ```
    show isis
    ```

20. Display the interface IS-IS configuration:

    ```
    show isis interface
    ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# spbm
VSP-9012:1(config)# router isis
VSP-9012:1(config-isis)# spbm 1
VSP-9012:1(config-isis)# spbm 1 b-vid 1000,2000 primary 1000
VSP-9012:1(config-isis)# exit
VSP-9012:1(config)# vlan create 1000 type spbm-bvlan
VSP-9012:1(config)# vlan create 2000 type spbm-bvlan
VSP-9012:1(config)# router isis
VSP-9012:1(config-isis)# spbm 1 nick-name 1.11.16
VSP-9012:1(config-isis)# manual-area 00.2000.0000.01
VSP-9012:1(config-isis)# exit
VSP-9012:1(config)# interface GigabitEthernet 3/21
VSP-9012:1(config-if)# isis
VSP-9012:1(config-if)# isis spbm 1
VSP-9012:1(config-if)# isis enable
VSP-9012:1(config-if)# exit
VSP-9012:1(config)# router isis enable
VSP-9012:1(config)# show isis spbm
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM       B-VID      PRIMARY    NICK     LSDB     IP       MULTICAST
INSTANCE              VLAN       NAME     TRAP
```

```
-------------------------------------------------------------------------------
1         1000,2000   1000      1.11.16  disable  disable  disable

===============================================================================
                             ISIS SPBM SMLT Info
===============================================================================
SPBM       SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
-------------------------------------------------------------------------------
1          secondary           00:14:c7:e1:33:e0     0018.b0bb.b3df


-------------------------------------------------------------------------------
 Total Num of SPBM instances: 1
-------------------------------------------------------------------------------
```

```
VSP-9012:1(config)# show isis
```

```
===============================================================================
                              ISIS General Info
===============================================================================
                            AdminState : enabled
                            RouterType : Level 1
                             System ID :0014.c7e1.33df
                  Max LSP Gen Interval : 900
                                Metric : wide
                   Overload-on-startup : 20
                              Overload : false
                         Csnp Interval : 10
                         PSNP Interval : 2
                     Rxmt LSP Interval : 5
                             spf-delay : 100
                           Router Name :VSP-9012
                     ip source-address :
                     Num of Interfaces : 2
                Num of Area Addresses : 1
```

```
VSP-9012:1(config)# show isis interface
```

```
===============================================================================
                              ISIS Interfaces
===============================================================================
IFIDX      TYPE    LEVEL        OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
-------------------------------------------------------------------------------
Mlt2       pt-pt   Level 1          UP         UP        1        1      10
Port3/21   pt-pt   Level 1          UP         UP        1        1      10
```

Variable definitions

Use the data in the following table to use the **isis** command.

Variable	Value
enable	Enables or disables the IS-IS circuit/interface on the specified port or MLT. The default is disabled. Use the no option to disable IS-IS on the specified interface.
spbm <1–100>	Enable the SPBM instance on the IS-IS interfaces.

Use the data in the following table to use the **manual-area** command.

Variable	Value
<xx.xxx.xxx...xxx>	Specifies the IS-IS manual-area (1–13 bytes in the format <xx.xxx.xxx...xxx>. In this release, only one manual area is supported. For IS-IS to operate, you must configure at least one area. Use the no option to delete the manual area.

Use the data in the following table to use the **spbm** command.

Variable	Value
<1–100>	Creates the SPBM instance. In this release, only one SPBM instance is supported.
b-vid *{<vlan-id [-vlan-id] [,...]}*	Sets the ISIS SPBM instance data VLANs. Use the no option to remove the specified B-VLAN from the SPBM instance.
nick-name <*x.xx.xx*>	Specifies a nickname for the SPBM instance globally. The value is 2.5 bytes in the format <x.xx.xx>. Use the no or default options to delete the configured nickname.
primary *<1–4084>*	Sets the IS-IS instance primary data VLAN.

Use the data in the following table to use the **vlan create** command.

Variable	Value
<2–4084>	Specifies the VLAN ID. Creates an SPBM Backbone VLAN (B-VLAN). You can optionally specifiy a name for the SPBM B-VLAN.
type {ipsubnet-mstprstp|port-mstprstp|protocol-mstprstp|spbm-bvlan|srcmac-mstprstp}	Specifies the type of VLAN created. • ipsubnet-mstprstp — Create a VLAN by IP subnet. • port-mstprstp — Create a VLAN by port. • protocol-mstprstp — Create a VLAN by protocol. • spbm-bvlan — Create an SPBM-BVLAN. • srcmac-mstprstp — Create a VLAN by source MAC address.

Job aid

Important:

After you have configured the SPBM nickname and enabled IS-IS, if you require a change of the system ID, you must also change the nickname. However, for naming convention purposes or

configuration purposes, you may not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps:

1. Disable IS-IS.
2. Change the system ID.
3. Change the nickname to a temporary one.
4. Enable IS-IS.
5. Disable IS-IS.
6. Change the nickname to the original nickname.
7. Enable IS-IS.

Displaying global SPBM parameters

Use the following procedure to verify the proper global SPBM configuration.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the SPBM configuration:

   ```
   show isis spbm
   ```

3. You can also use the following command to identify SPBM VLANs. For spbm-bvlan, the attribute `TYPE` displays `spbm-bvlan` instead of `byport`.

   ```
   show vlan basic
   ```

Example

```
VSP-9012:# enable
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM       B-VID      PRIMARY    NICK     LSDB     IP       MULTICAST
INSTANCE              VLAN       NAME     TRAP
--------------------------------------------------------------------------------
1         1000,2000   1000       1.11.16  disable  disable  disable

================================================================================
                              ISIS SPBM SMLT Info
================================================================================
SPBM       SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
--------------------------------------------------------------------------------
1          secondary           00:14:c7:e1:33:e0     0018.b0bb.b3df


--------------------------------------------------------------------------------
 Total Num of SPBM instances: 1
--------------------------------------------------------------------------------
```

```
VSP-9012:1# show vlan basic
================================================================================
```

```
                                    Vlan Basic
================================================================================
VLAN                                INST
ID    NAME              TYPE          ID  PROTOCOLID SUBNETADDR      SUBNETMASK
--------------------------------------------------------------------------------
1     Default           byPort        0   none       N/A             N/A
10    VLAN-10           spbm-bvlan    62  none       N/A             N/A
20    VLAN-20           spbm-bvlan    62  none       N/A             N/A
100   VLAN-100          byPort        0   none       N/A             N/A
1000  VLAN-1000         byPort        25  none       N/A             N/A

All 5 out of 5 Total Num of Vlans displayed
```

Job aid

The following table describes the fields in the output for the **`show isis spbm`** command.

Parameter	Description
SPBM INSTANCE	Indicates the SPBM instance identifier. You can only create one SPBM instance.
B-VID	Indicates the SPBM B-VLAN associated with the SPBM instance.
PRIMARY VLAN	Indicates the primary SPBM B-VLAN.
NICK NAME	Indicates the SPBM node nickname. The nickname is used to calculate the I-SID multicast MAC address.
LSDB TRAP	Indicates the status of the IS-IS SPBM LSDB update trap on this SPBM instance. The default is disable.
IP	Indicates the status of SPBM IP shortcuts on this SPBM instance. The default is disable.
MULTICAST	Indicates if SPBM multicast is enabled. The default is disabled.
SPBM INSTANCE	Indicates the SPBM instance identifier. You can only create one SPBM instance.
SMLT-SPLIT-BEB	Specifies whether the switch is the primary or secondary IST peer.
SMLT-VIRTUAL-BMAC	Specifies a virtual MAC address that can be used by both peers.
SMLT-PEER-SYSTEM-ID	Specifies the IST peer system ID.

Displaying global IS-IS parameters

Use the following procedure to display the global IS-IS parameters.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display IS-IS configuration information:

   ```
   show isis
   ```

3. Display the IS-IS system-id:

   ```
   show isis system-id
   ```

4. Display IS-IS net info:

   ```
   show isis net
   ```

Example

```
VSP-9012:1#show isis
================================================================================
                              ISIS General Info
================================================================================
                           AdminState : enabled
                           RouterType : Level 1
                            System ID : 0014.c7e1.33df
                  Max LSP Gen Interval : 900
                             Metric : wide
                   Overload-on-startup : 20
                               Overload : false
                         Csnp Interval : 10
                         PSNP Interval : 2
                          Rxmt LSP Interval : 5
                             spf-delay : 100
                           Router Name : VSP-9012
                         ip source-adddress :
                       Num of Interfaces : 2
                        Num of Area Addresses : 1
```

```
VSP-9012:1#show isis system-id
================================================================================
                                ISIS System-Id
================================================================================
SYSTEM-ID
--------------------------------------------------------------------------------
0014.c7e1.33df
```

```
VSP-9012:1#show isis net
================================================================================
                                ISIS Net Info
================================================================================
NET
--------------------------------------------------------------------------------
c0.2000.0000.0000.14c7.e133.df00
```

Job aid

The following sections describe the fields in the outputs for the global IS-IS show commands.

show isis

The following table describes the fields in the output for the **show isis** command.

Parameter	Description
AdminState	Indicates the administrative state of the router.
RouterType	Indicates the router Level: l1, l2, or l1/2.
System ID	Indicates the system ID.
Max LSP Gen Interval	Indicates the maximum time between LSP updates in seconds.
Metric	Indicates if the metric is narrow or wide.
Overload-on-startup	Indicates the overload-on-startup value.
Overload	Indicates if there is an overload condition.

Parameter	Description
Csnp Interval	Indicates the interval between CSNP updates in seconds.
PSNP Interval	Indicates the interval between PSNP updates in seconds.
Rxmt LSP Interval	Indicates the received LSP time interval.
spf-delay	Indicates the Shortest Path First delay in milliseconds.
Router Name	Indicates the IS-IS name of the router.
ip source-address	Indicates the IP source address used for SPBM IP shortcuts.
Num of Interfaces	Indicates the number of interfaces on the router.
Num of Area Addresses	Indicates the number of area addresses on the router.

show isis system-id

The following table describes the fields in the output for the **`show isis system-id`** command.

Parameter	Description
SYSTEM-ID	Shows the system ID. Output from this show command is from the global IS-IS configuration of the system ID. There is one system ID configured. The system ID is 6 bytes in length.

show isis net

The following table describes the fields in the output for the **`show isis net`** command.

Parameter	Description
NET	Shows the NET address. Output from this command is from the global IS-IS configuration of the manual area and the configuration of the system ID. There is only one manual areas defined and only one system ID. The manual area is from 1-13 bytes in length. The system ID is 6 bytes in length.

Enabling IP multicast over SPBM globally

Use this procedure to enable IP multicast over SPBM globally on the Backbone Edge Bridges (BEBs) that directly or indirectly (using Layer 2 switches) connect to IP multicast senders or receivers. By default, IP multicast over SPBM is disabled. There is no need to enable IP multicast over SPBM on the Backbone Core Bridges (BCBs).

You must configure IP multicast over SPBM at the global level, and then enable it on the service option or options you choose.

Note:

SPBM multicast uses I-SIDs starting at 16,000,000 and above. If Layer 2 or Layer 3 I-SIDs are in this range, the system displays an error message and the switch does not enable IP multicast over SPBM.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Verify no I-SIDs exist in the default reserved range:
   a. For Layer 2 use the following command:

      ```
      show vlan i-sid
      ```

   b. For Layer 3 use the following command:

      ```
      show ip ipvpn vrf WORD<1-16>
      ```

3. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

4. Enable SPBM multicast globally:

   ```
   spbm <1-100> multicast enable
   ```

   Note:

   In this release, the switch only supports one SPBM instance.

5. **(Optional)** Disable SPBM multicast globally:

   ```
   no spbm <1-100> multicast enable
   default spbm <1-100> multicast enable
   ```

Example

Enable IP multicast over SPBM globally:

```
Switch:1(config)#show vlan i-sid

================================================================================
                  Vlan I-SID
================================================================================
VLAN_ID    I-SID
--------------------------------------------------------------------------------
1
50          200
51
52
53
54
55
56
57
```

```
9 out of 9 Total Num of Vlans displayed
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 multicast enable
```

Variable definitions

Use the data in the following table to use the **spbm** command.

Variable	Value
<1–100>	Enables IP multicast over SPBM globally. The default is disabled. Specifies the SPBM instance. **Note:** In this release, the switch only supports one instance.

Displaying IP multicast over SPBM information

Use this procedure to display IP multicast over SPBM summary information.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the status of the global SPBM multicast configuration:

   ```
   show isis spbm multicast
   ```

3. Display IP multicast over SPBM summary information for each S, G, V tuple:

   ```
   show isis spb-mcast-summary [host-name WORD<0-255>][lspid
   <xxxx.xxxx.xxxx.xx-xx>]
   ```

4. Display information about the multicast routes on the switch:

   ```
   show ip mroute route
   ```

Example

Display IP multicast over SPBM global configuration information:

```
Switch:1>enable
Switch:1#show isis spbm multicast

                            multicast : enable
                    fwd-cache-timeout(seconds) : 210

Switch:1#show isis spb-mcast-summary

================================================================================
                              SPB multicast - Summary
================================================================================
SCOPE    SOURCE           GROUP            DATA          LSP  HOST
```

```
I-SID    ADDRESS          ADDRESS          I-SID   BVID  FRAG NAME
-------------------------------------------------------------------------
GRT      192.0.2.102     233.252.0.1        16000001  63  0x0  DIST5A

Switch: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

Total 4
```

Variable definitions

Use the data in the following table to use the **`show isis spb-mcast-summary`** command.

Variable	Value
host-name *WORD<0–255>*	Displays the IP multicast over SPBM summary information for a specific host-name.
lspid *<xxx.xxx.xxx.xx-xx>*	Displays the IP multicast over SPBM summary information for the specified LSP ID that you enter in xxx.xxx.xxx.xx-xx — 8 byte format.

Job aid

The following table describes the fields in the output for the **`show isis spbm multicast`** command.

Parameter	Description
multicast	Specifies if multicast is enabled.
fwd-cache-timeout (seconds)	Specifies the forward cache timeout value in seconds.

The following table describes the fields in the output for the **`show isis spb-mcast-summary`** command.

Parameter	Description
SCOPE I-SID	Indicates the I-SID that specifies the multicast streams when the scope is either the Layer 3 VSN or the Layer 2 VSN or any combination.
SOURCE ADDRESS	Indicates the IP multicast source address that maps to the I-SID.
GROUP ADDRESS	Indicates the IP multicast group address that maps to the I-SID.
DATA I-SID	Indicates the data I-SID for the IP multicast route, which includes the source IP address, group IP address, and the local VLAN that the stream is

Parameter	Description
	received on (S,G,V tuple). SPBM uses the data I-SID to create the multicast tree.
BVID	Indicates the ID of the SPBM backbone VLAN (B-VLAN) on which the multicast stream forwards in the SPBM cloud.
LSP FRAG	Indicates the fragment number of the LSP ID.
HOST-NAME	Indicates the host name of the router.

The following table describes the fields in the output for the **`show ip mroute route`** command.

Parameter	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. The field displays the value of 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 values indicate the stream learned when SPBM multicast is configured on the VLAN. The spb-access value indicates that it was learned on the access. The spb-network value indicates it was learned over the SPBM cloud.

Displaying IS-IS areas

Use the following procedure to display IS-IS areas.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Use the following procedure to display IS-IS areas.

   ```
   show isis manual-area
   ```

Example

```
VSP-9012:1#show isis manual-area
================================================================================
                            ISIS Manual Area Address
================================================================================

AREA ADDRESS
--------------------------------------------------------------------------------
c0.2000.0000.00
```

Job aid

The following table describes the fields in the output for the **`show isis manual-area`** command.

Parameter	Description
AREA ADDRESS	Shows the manual areas defined. There can only be one area. Use the same manual area across the entire SPBM cloud. The manual area can be from 1-13 bytes in length.

Configuring SMLT parameters for SPBM

Use the following procedure to configure the required Split MultiLink Trunking (SMLT) parameters to allow SPBM to interoperate with SMLT on the switch.

Note:

If you want to modify SMLTs that contain NNI ports, Avaya recommends you do the modification during maintenance windows. Otherwise, if you create or delete SMLTs that contain NNI ports on Virtual Services Platform 9000 running MSTP, IS-IS adjacencies that connect to those ports can bounce even if the SMLT is not used.

Note:

- The assignment of primary and secondary roles to the IST peers is automatic. The switch with the lower system ID (between the two IST peers) is primary, and the switch with the higher system ID is secondary.

- SMLT virtual B-MAC is an optional configuration. If SMLT virtual B-MAC is not configured, the system derives SMLT virtual B-MAC from the configured SMLT peer system ID and the nodal MAC of the device (IS-IS system ID). The system compares the nodal MAC of the device with the SMLT peer system ID configured and takes the small one, plus 0x01, as the SMLT virtual B-MAC.
- Different from Ethernet Routing Switch 8800 SPBM implementation, on Virtual Services Platform 9000 when you map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID), you must include the IST MLT peer switches in the VLAN, which means all IST member ports must be members of this C-VLAN. In ERS 8800, you cannot configure IST MLT to be part of the C-VLAN. For more information about this and other differences between the two products, see *Platform Migration Reference for Avaya Virtual Services Platform 9000,* NN46250-107.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable IS-IS on the switch:

   ```
   no router isis enable
   ```

3. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

4. Specify the system ID of the IST peer, so that if it goes down, the local peer can take over forwarding for the failed peer:

   ```
   spbm <1-100> smlt-peer—system—id <xxxx.xxxx.xxxx>
   ```

5. Configure the virtual B-MAC, which is shared and advertised by both peers:

   ```
   spbm <1-100> smlt-virtual-bmac <0x00:0x00:0x00:0x00:0x00:0x00>
   ```

6. Exit to Global Configuration mode:

   ```
   exit
   ```

7. Enable IS-IS on the switch:

   ```
   router isis enable
   ```

8. Display the SPBM SMLT configuration:

   ```
   show isis spbm
   ```

Example

```
VSP:9012:1> enable
```

```
VSP:9012:1# configure terminal
```

Disable IS-IS on the switch:

```
VSP:9012:1(config)# no router isis enable
```

Enter the IS-IS Router Configuration mode:

```
VSP:9012:1(config)# router isis
VSP:9012:1(config-isis)# spbm 1 smlt-peer-system-id 0018.b0bb.b3df
VSP:9012:1(config-isis)# spbm 1 smlt-virtual-bmac 00:14:c7:e1:33:e0
VSP:9012:1(config-isis)# router isis enable
VSP:9012:1(config-isis)# show isis spbm
```

```
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM       B-VID      PRIMARY    NICK     LSDB     IP       MULTICAST
INSTANCE              VLAN       NAME     TRAP
--------------------------------------------------------------------------------
1         1000,2000   1000      1.11.16  disable  disable  disable

================================================================================
                              ISIS SPBM SMLT Info
================================================================================
SPBM       SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
--------------------------------------------------------------------------------
1          secondary           00:14:c7:e1:33:e0     0018.b0bb.b3df


--------------------------------------------------------------------------------
 Total Num of SPBM instances: 1
--------------------------------------------------------------------------------
```

Variable definitions

Use the data in the following table to use the **spbm** command.

Variable	Value
smlt-peer-system-id <*xxxx.xxxx.xxxx*>	Specifies the IS-IS SPBM peer system ID. SMLT peer system ID is part of the required configuration. You must configure the SMLT peer system ID as the nodal MAC of the peer device. In the IS-IS network, the nodal MAC of devices should be eight apart from each other.
smlt-virtual-bmac <*0x00:0x00:0x00:0x00:0x00:0x00:0x00*>	Specifies a virtual MAC address that can be used by both peers. SMLT virtual B-MAC is an optional configuration. **Note:** • If SMLT virtual B-MAC is not configured, the system derives SMLT virtual B-MAC from the configured SMLT peer system ID and the nodal MAC of the device (IS-IS system ID). The system compares the nodal MAC of the

Variable	Value
	device with the SMLT peer system ID configured and takes the small one, plus 0x01, as the SMLT virtual B-MAC. • The system also derives SMLT split BEB from the SMLT peer system ID and nodal MAC of the device. The device with the lower system ID is primary, the device with the higher system ID is secondary.

Configuring optional SPBM parameters

Use the following procedure to configure optional SPBM parameters.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the SPBM ethertype:

   ```
   spbm ethertype {0x8100 | 0x88a8}
   ```

3. Configure the optional link-state database (LSDB) trap global parameter. To configure this parameter, you must globally disable IS-IS on the switch:

   ```
   no router isis enable
   ```

4. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

5. Enable a trap when the SPBM LSDB changes:

   ```
   spbm <1-100> lsdb-trap enable
   ```

6. Enable IS-IS on the switch:

   ```
   router isis enable
   ```

7. Exit to Global Configuration mode:

   ```
   exit
   ```

8. Configure the optional SPBM interface parameters. To configure these parameters, you must disable IS-IS on the interface.
9. Enter Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[-slot/port][,...]}
```
or `interface mlt <1-512>`

10. Disable IS-IS on the interface:

    ```
    no isis enable
    ```

11. Configure SPBM instance interface-type on IS-IS interface. SPBM supports only pt-pt:

    ```
    isis spbm <1-100> interface-type {broadcast|pt-pt}
    ```

12. Configure the SPBM instance level 1 metric on the IS-IS interface:

    ```
    isis spbm <1-100> l1-metric <1-16777215>
    ```

13. Enable IS-IS on the switch:

    ```
    isis enable
    ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# spbm ethertype 0x8100
VSP-9012:1(config-isis)# no router isis enable
VSP-9012:1(config)# router isis
VSP-9012:1(config-isis)# spbm 1 lsdb-trap enable
VSP-9012:1(config-isis)# router isis enable
VSP-9012:1(config-isis)# exit
VSP-9012:1(config)# interface gigabitethernet 3/7
VSP-9012:1(config-if)# no isis enable
VSP-9012:1(config-if)# isis spbm 1 interface-type pt-pt
VSP-9012:1(config-if)# isis spbm 1 l1-metric 500
VSP-9012:1(config-if)# isis enable
```

Variable definitions

Use the data in the following table to use the **spbm** command.

Variable	Value
ethertype *{0x8100 | 0x88a8}*	Configures the SPBM ethertype. The default value is 0x8100.
<1–100> lsdb-trap enable	Configures whether to enable or disable a trap when the SPBM LSDB changes.

Variable	Value
	The default is disabled. Use the no or default options to disable LSDB traps.

Use the data in the following table to use the **isis spbm** command.

Variable	Value
<1–100> interface-type *{broadcast|pt-pt}*	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT. SPBM only supports the point-to-point (pt-pt) interface type. The default is pt-pt. Use the no or default options to set this parameter to the default value of pt-pt.
<1–100> l1–metric *<1–16777215>*	Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10. Use the no or default options to set this parameter to the default.

Configuring optional IS-IS global parameters

Use the following procedure to configure optional IS-IS global parameters.

Procedure

1. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

2. Configure optional IS-IS global parameters:
   a. Specify the Complete Sequence Number Packet (CSNP) interval in seconds:

      ```
      csnp-interval <1-600>
      ```

   b. Configure the router type globally:

      ```
      is-type {l1|l12}
      ```

   c. Configure the maximum level, in seconds, between generated LSPs by this Intermediate System:

      ```
      max-lsp-gen-interval <30-900>
      ```

   d. Configure the IS-IS metric type:

      ```
      metric {wide}
      ```

   e. Set or clear the overload condition:

```
overload
```

f. Configure the overload-on-startup value in seconds:

```
overload-on-startup <15-3600>
```

g. Configure the Partial Sequence Number Packet (PSNP) in seconds:

```
psnp-interval <1-120>
```

h. Configure the minimum time between retransmission of an LSP:

```
retransmit-lsp-interval <1-300>
```

i. Configure the SPF delay in milliseconds:

```
spf-delay <0-5000>
```

j. Configure the name for the system:

```
sys-name WORD<0-255>
```

k. Configure the IS-IS system ID for the switch:

```
system-id <xxxx.xxxx.xxxx>
```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# router isis
VSP-9012:(config-isis)# csnp-interval 10
VSP-9012:(config-isis)# is-type l1
VSP-9012:(config-isis)# max-lsp-gen-interval 800
VSP-9012:(config-isis)# metric wide
VSP-9012:(config-isis)# overload
VSP-9012:(config-isis)# overload-on-startup 30
VSP-9012:(config-isis)# psnp-interval 10
VSP-9012:(config-isis)# retransmit-lsp-interval 10
VSP-9012:(config-isis)# default sys-name
VSP-9012:(config-isis)# spf-delay 200
```

Variable definitions

Use the data in the following table to use the **csnp-interval** command.

Variable	Value
<1–600>	Specifies the CSNP interval in seconds. This is a system level parameter that applies for level 1 CSNP

Variable	Value
	generation on all interfaces. A longer interval reduces overhead, while a shorter interval speeds up convergence. The default value is 10. Use the no or default options to set this parameter to the default value of 10.

Use the data in the following table to configure the **is-type** command.

Variable	Value
{l1|l12}	Sets the router type globally: • l1: Level-1 router type • l12: Not valid in the current release. The default value is l1. Use the no or default options to set this parameter to the default value of l1.

Use the data in the following table to configure the **max-lsp-gen-interval** command.

Variable	Value
<30–900>	Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate System. The default value is 900 seconds. Use the no or default options to set this parameter to the default value of 900.

Use the data in the following table to configure the **metric** command.

Variable	Value
{narrow|wide}	Specifies the IS-IS metric type. Only wide is supported in this release. The default value is wide. Use the no or default options to set this parameter to the default value of wide.

Use the data in the following table to configure the **overload** command.

Variable	Value
overload	Sets or clears the overload condition. The default value is disabled. Use the no or default options to set this parameter to the default value of disabled.

Use the data in the following table to configure the **overload-on-startup** command.

Variable	Value
<15–3600>	Specifies the IS-IS overload-on-startup value in seconds. The overload-on-startup value is used as a timer to control when to send out LSPs with the overload bit cleared after IS-IS startup. The default value is 20. Use the no or default options to set this parameter to the default value of 20.

Use the data in the following table to configure the **psnp-interval** command.

Variable	Value
<1–120>	Specifies the PSNP interval in seconds. This is a system level parameter that applies for level 1 PSNP generation on all interfaces. A longer interval reduces overhead, while a shorter interval speeds up convergence. The default value is 2. Use the no or default options to set this parameter to the default value of 2.

Use the data in the following table to configure the **retransmit-lsp-interval** command.

Variable	Value
<1–300>	Specifies the minimum time between retransmission of an LSP. This defines how fast the switch resends the same LSP. This is a system level parameter that applies for Level1 retransmission of LSPs. The default value is 5 seconds. Use the no or default options to set this parameter to the default value of 5.

Use the data in the following table to configure the **spf-delay** command.

Variable	Value
<0–5000>	Configures the delay, in milliseconds, to pace successive Shortest Path First (SPF) runs. The timer prevents more than two SPF runs from being scheduled back-to-back. The mechanism for pacing SPF allows two back-to-back SPF runs. The default value is 100 milliseconds. Use the no or default options to set this parameter to the default value of 100 milliseconds.

Use the data in the following table to configure the **sys-name** command.

Variable	Value
WORD<0–255>	Specifies a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763.

Variable	Value
	By default, the system name comes from the host name configured at the system level. Use the no or default options to set this parameter to the default value (host name). **Note:** In this release, no consistency checks appear when you edit sys-name on Virtual Services Platform 9000.

Use the data in the following table to configure the **system-id** command.

Variable	Value
<xxxx.xxxx.xxxx>	Specifies the IS-IS system ID for the switch. Use the no or default options to set this parameter to the default value (node BMAC).

Job aid

Important:

After you have configured the SPBM nickname and enabled IS-IS, if you require a change of the system ID, you must also change the nickname. However, for naming convention purposes or configuration purposes, you may not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps:

1. Disable IS-IS.
2. Change the system ID.
3. Change the nickname to a temporary one.
4. Enable IS-IS.
5. Disable IS-IS.
6. Change the nickname to the original nickname.
7. Enable IS-IS.

Configuring optional IS-IS interface parameters

Use the following procedure to configure optional IS-IS interface parameters.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   ```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[-slot/port][,...]}
```
or
```
interface mlt <1-512>
```

2. Configure optional IS-IS interface parameters:
   a. Specify the authentication type used for IS-IS hello packets on the interface:
      ```
      isis hello-auth type {none|simple|hmac-md5}
      ```
   b. If you select `simple` as the `hello-auth type`, you must also specify a key value but the key-id is optional:
      ```
      isis hello-auth type simple key WORD<1-16> [key-id <1-255>]
      ```
   c. If you select `hmac-md5`, you must also specify a key value but the key-id is optional:
      ```
      isis hello-auth type hmac-md5 key WORD<1-16> [key-id <1-255>]]
      ```
   d. Configure the level 1 IS-IS designated router priority:
      ```
      isis [l1-dr-priority <0-127>]
      ```
      Note:

      This parameter is not used for SPBM because SPBM only runs on point-to-point interfaces. This parameter is for designated router election on a broadcast LAN segment, which is not supported.
   e. Configure the level 1 hello interval:
      ```
      isis [l1-hello-interval <1-600>]
      ```
   f. Configure the level 1 hello multiplier:
      ```
      isis [l1-hello-multiplier <1-600>]
      ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# interface gigabitethernet 3/7
VSP-9012:1(config-if)# isis
VSP-9012:1(config-if)# isis hello-auth type hmac-md5 key test
VSP-9012:1(config-if)# isis l1-dr-priority 100
VSP-9012:1(config-if)# isis l1-hello-interval 20
VSP-9012:1(config-if)# isis l1-hello-multiplier 10
```

Variable definitions

Use the data in the following table to configure the **`isis`** command.

Variable	Value
hello-auth type *{none|simple|hmac-md5}] [key*[key *WORD<1–16>*] [key-id *<1–255>*]	Specifies the authentication type used for IS-IS hello packets on the interface. type can be one of the following: • none • simple: If selected, you must also specify a key value but the key id is optional. Simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5: If selected, you must also specify a key value but the key-id is optional. MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. There is an optional key ID. The default is none. Use the no or default options to set the hello-auth type to none.
l1-dr-priority *<0–127>*	Configures the level 1 IS-IS designated router priority to the specified value. The default value is 64. Use the no or default options to set this parameter to the default value of 64. **Note:** This parameter is not used for SPBM because SPBM only runs on point-to-point interfaces. This parameter is for designated router election on a broadcast LAN segment, which is not supported.
l1-hello-interval *<1–600>*	Configures the level 1 hello interval. The default value is 9 seconds. Use the no or default options to set this parameter to the default value of 9 seconds.
l1-hello-multiplier *<1–600>*	Configures the level 1 hello multiplier. The default value is 3 seconds. Use the no or default options to set this parameter to the default value of 3 seconds.

Displaying IS-IS interface parameters

Use the following procedure to display the IS-IS interface parameters.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display IS-IS interface configuration and status parameters (including adjacencies):

```
show isis interface [l1|l2|l12]
```

3. Display IS-IS interface authentication configuration:

   ```
   show isis int-auth
   ```

4. Display IS-IS interface timers:

   ```
   show isis int-timers
   ```

5. Display IS-IS circuit level parameters:

   ```
   show isis int-ckt-level
   ```

Example

```
VSP-9012:1#show isis interface
================================================================================
                                ISIS Interfaces
================================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ   SPBM-L1-METRIC
--------------------------------------------------------------------------------
Mlt2       pt-pt   Level 1      UP      UP         1        1     10
Port3/21   pt-pt   Level 1      UP      UP         1        1     10
```

```
VSP-9012:1#show isis int-auth
================================================================================
                              ISIS Interface Auth
================================================================================
IFIDX        AUTH-TYPE       AUTH-KEYID      AUTH-KEY
--------------------------------------------------------------------------------
Mlt2         none            0
Port3/21     none            0
```

```
VSP-9012:1#show isis int-timers
================================================================================
                              ISIS Interface Timers
================================================================================
IFIDX           LEVEL           HELLO               HELLO           HELLO
                                    INTERVAL        MULTIPLIER      DR
--------------------------------------------------------------------------------
Mlt2          Level 1             9               3               3
Port3/21        Level 1             9               3               3
```

```
VSP-9012:1#show isis int-ckt-level
================================================================================
                        ISIS Circuit level parameters
================================================================================
IFIDX           LEVEL                DIS                  CKTID
--------------------------------------------------------------------------------
Mlt2            Level 1                                     1
Port3/21        Level 1                                     2
```

Variable definitions

Use the data in the following table to use the IS-IS interface show command.

Variable	Value
[l1, l2, l12]	Displays the interface information for the specified level: l1, l2, or l12.

Job aid

The following sections describe the fields in the outputs for the IS-IS interface show commands.

show isis interface

The following table describes the fields in the output for the **`show isis interface`** command.

Parameter	Description
IFIDX	Indicates the interface index for the Ethernet or MLT interface.
TYPE	Indicates the type of interface configured (in this release, only pt-pt is supported).
LEVEL	Indicates the level of the IS-IS interface (Level 1 [default] or Level 2).
OP-STATE	Shows the physical connection state of the interface.
ADM-STATE	Shows the configured state of the interface.
ADJ	Shows how many adjacencies are learned through the interface.
UP-ADJ	Shows how many adjacencies are active through the interface.
SPBM-L1–METRIC	Indicates the SPBM instance Level 1 metric on the IS-IS interface.

show isis int-auth

The following table describes the fields in the output for the **`show isis int-auth`** command.

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.
AUTH-TYPE	Shows the type of authentication configured for the interface. Types include: • none for no authentication. • simple for a simple password. • hmac-md5 for MD5 encryption.
AUTH-KEYID	Shows the authentication password configured for the interface.
AUTH-KEY	Shows the HMAC-MD5 key needed for encryption. This is used only for HMAC-MD5.

show isis int-timers

The following table describes the fields in the output for the **`show isis int-auth`** command.

Parameter	Description
IFIDX	Indicates the interface index for the Ethernet or MLT interface.
LEVEL	Indicates the IS-IS interface level.
HELLO INTERVAL	Indicates the interval at which a Hello packet is sent to the IS-IS network.
HELLO MULTIPLIER	Indicates the multiplier that is used in conjunction with the Hello Interval.
HELLO DR	Indicates the interval at which a Hello packet is sent to the IS-IS network if the router is a designated router (DIS).

show isis int-ckt-level

The following table describes the fields in the output for the **show isis int-ckt-level** command.

Parameter	Description
IFIDX	Shows the interface index for the ethernet or MLT interface.
LEVEL	Shows the level of the IS-IS interface (Level 1 [default] or Level 2).
DIS	Shows the Designated Intermediate System (DIS) of the circuit.
CKT ID	Displays the CKT ID.

Displaying the IP unicast FIB, multicast FIB, unicast FIB, and unicast tree

Use the following procedure to display SPBM IP unicast Forwarding Information Base (FIB), SPBM multicast FIB, unicast FIB, and the unicast tree.

About this task

In SPBM, Backbone MAC (B-MAC) addresses are carried within the IS-IS link-state database (LSDB). To do this, SPBM supports an IS-IS Type-Length-Value (TLV) that advertises the Service Instance Identifier (I-SID) and B-MAC information across the network. Each node has a System ID, which also serves as B-MAC of the switch. These B-MAC addresses are populated into the SPBM Forwarding Information Base (FIB).

When the network topology is discovered and stored in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

I-SIDs are only used for virtual services (Layer 2 VSNs and Layer 3 VSNs). If you only enable IP Shortcuts on the Backbone Edge Bridges, I-SIDs are never exchanged in the network as IP Shortcuts allows for Global Routing Table (GRT) IP networks to be transported across IS-IS.

The **show isis spbm ip-unicast-fib** command displays all of the IS-IS routes in the IS-IS LSDB. The IP ROUTE PREFERENCE column in the **show isis spbm ip-unicast-fib** command displays the IP route preference.

Routes within the same VSN are added to the LSDB with a default preference of 7. Inter-VSN routes are added to the LSDB with a route preference of 200. IS-IS accept policies allow you to change the route preference for incoming routes. If the same route is learned from multiple sources with different route preferences, then the routes are not considered equal cost multipath (ECMP) routes. The route with the lowest route preference is the preferred route.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the SPBM IP unicast FIB:

```
show isis spbm ip-unicast-fib [all] [id <1-6777215] [spbm-nh-as-mac]
```

3. Display the SPBM multicast FIB:

   ```
   show isis spbm multicast-fib [i-sid <1-16777215>] [nick-name
   <x.xx.xx>] [summary] [vlan <2-4084>]
   ```

4. Display the SPBM unicast FIB:

   ```
   show isis spbm unicast-fib [b-mac <0x00:0x00:0x00:0x00:0x00:0x00>]
   [summary] [vlan <2-4084>]
   ```

5. Display the SPBM unicast tree:

   ```
   show isis spbm unicast-tree <2-4084> [destination <xxxx.xxxx.xxxx>]
   ```

Example

```
Switch:#enable
Switch:1#show isis spbm ip-unicast-fib

================================================================================
                          SPBM IP-UNICAST FIB ENTRY INFO
================================================================================
     VRF  DEST                                    OUTGOING  SPBM  PREFIX  IP ROUTE
VRF  ISID ISID  Destination     NH BEB   VLAN INTERFACE COST  COST    PREFERENCE
--------------------------------------------------------------------------------
GRT  -    -     1.1.1.13/32     VSP13    1000 10/7      10    1        7
GRT  -    -     1.1.1.13/32     VSP13    1001 10/7      10    1        7
GRT  -    7     5.7.1.0/24      VSP5     7    Local     0     1        200
GRT  -    11    11.1.1.0/24     VSP13    1000 10/7      10    1        200
GRT  -    11    11.1.1.0/24     VSP13    1001 10/7      10    1        200
GRT  -    11    11.11.11.11/32  VSP13    1000 10/7      10    1        200
GRT  -    11    11.11.11.11/32  VSP13    1001 10/7      10    1        200
GRT  -    7     13.7.1.0/24     VSP13    1000 10/7      10    1        200
GRT  -    7     13.7.1.0/24     VSP13    1001 10/7      10    1        200
GRT  -    -     34.34.34.34/32  VSP13    1000 10/7      10    1        7
GRT  -    -     34.34.34.34/32  VSP13    1001 10/7      10    1        7

--------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 11
--------------------------------------------------------------------------------
```

```
Switch:1#show isis spbm multicast-fib

================================================================================

              SPBM MULTICAST FIB ENTRY INFO
================================================================================
MCAST DA          ISID  BVLAN  SYSID          HOST-NAME OUTGOING-INTERFACES INCOMING
                                                                    INTERFACE
--------------------------------------------------------------------------------
03:00:61:00:00:64 100    10    0080.2dc1.37ce 9000-1        4/7             5/7
03:00:61:00:00:c8 200    10    0080.2dc1.37ce 9000-1        4/2,4/1         5/2

--------------------------------------------------------------------------------
 Total number of SPBM MULTICAST FIB entries 2
--------------------------------------------------------------------------------
```

```
Switch:1#show isis spbm unicast-fib
================================================================================
                          SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION          BVLAN SYSID          HOST-NAME    OUTGOING       COST
ADDRESS                                                INTERFACE
--------------------------------------------------------------------------------
00:16:ca:23:73:df   1000  0016.ca23.73df  SPBM-1      3/21           10
```

```
00:16:ca:23:73:df   2000  0016.ca23.73df  SPBM-1      3/21            10
00:18:b0:bb:b3:df   1000  0018.b0bb.b3df  SPBM-2      MLT-2           10
00:14:c7:e1:33:e0   1000  0018.b0bb.b3df  SPBM-2      MLT-2           10
00:18:b0:bb:b3:df   2000  0018.b0bb.b3df  SPBM-2      MLT-2           10

--------------------------------------------------------------------------------
 Total number of SPBM UNICAST FIB entries 5
--------------------------------------------------------------------------------
```

```
VSP-9012:1#show isis spbm unicast-tree 1000
Node:0018.b0bb.b3df.00 (VSP-9010) -> ROOT
Node:0016.ca23.73df.00 (VSP-9012) -> ROOT
```

Variable definitions

Use the data in the following table to use the **show isis spbm ip-unicast-fib** command.

Variable	Value
all	Displays entries for the Global Routing Table (GRT) and all Virtual Routing and Forwarding (VRF) instances. **Note:** If you use the command **show isis spbm ip-unicast-fib** the device displays only GRT entries. The command shows IP routes from remote Backbone Edge Bridges (BEBs).
id *<1–6777215>*	Displays IS-IS SPBM IP unicast Forwarding Information Base (FIB) information by Service Instance Identifier (I-SID) ID.
spbm-nh-as-mac	Displays the next hop B-MAC of the IP unicast FIB entry.

Use the data in the following table to use the **show isis spbm multicast-fib** command.

Variable	Value
i-sid *<1–16777215>*	Displays the FIB for the specified I-SID.
nick-name *<x.xx.xx>*	Displays the FIB for the specified nickname.
summary	Displays a summary of the FIB.
vlan *<2–4084>*	Displays the FIB for the specified SPBM VLAN.

Use the data in the following table to use the **show isis spbm unicast-fib** command.

Variable	Value
b-mac *<0x00:0x00:0x00:0x00:0x00:0x00:0x00>*	Displays the FIB for the specified BMAC.
summary	Displays a summary of the FIB.
vlan *<2–4084>*	Displays the FIB for the specified SPBM VLAN.

Use the data in the following table to use the **show isis spbm unicast-tree** command.

Variable	Value
<2–4084>	Specifies the SPBM B-VLAN ID.
destination *<xxxx.xxxx.xxxx>*	Displays the unicast tree for the specified destination.

Job aid

The following sections describe the fields in the outputs for SPBM multicast FIB, unicast FIB, and unicast tree show commands.

show isis spbm ip-unicast-fib

The following table describes the fields in the output for the **`show isis spbm ip-unicast-fib`** command.

Parameter	Dsecription
VRF	Specifies the VRF ID of the IP unicast FIB entry, 0 indicates NRE.
VRF ISID	Specifies the I-SID of the IP unicast FIB entry.
DEST ISID	Specifies the destination I-SID.
Destination	Specifies the destination IP address of the IP unicast FIB entry.
NH BEB	Specifies the next hop B-MAC of the IP unicast FIB entry.
VLAN	Specifies the VLAN of the IP unicast FIB entry.
OUTGOING INTERFACE	Specifies the outgoing port of the IP unicast FIB.
SPBM COST	Specifies the B-MAC cost of the IP unicast FIB entry.
PREFIX COST	Specifies the prefix cost of the IP unicast FIB entry.
IP ROUTE PREFERENCE	Specifies the IP route preference.

show isis spbm multicast-fib

The following table describes the fields in the output for the **`show isis spbm multicast-fib`** command.

Parameter	Description
MCAST DA-INTERFACES	Indicates the multicast destination MAC address of the multicast FIB entry.
ISID	Indicates the I-SID of the multicast FIB entry.
BVLAN	Indicates the B-VLAN of the multicast FIB entry.
SYSID	Indicates the system identifier of the multicast FIB entry.
HOST-NAME	Indicates the host name of the multicast FIB entry.
OUTGOING-INTERFACES	Indicates the outgoing interface of the multicast FIB entry.
INCOMING INTERFACE	Indicates the incoming interface of the multicast FIB entry.

show isis spbm unicast-fib

The following table describes the fields in the output for the **show isis spbm unicast-fib** command.

Parameter	Description
DESTINATION ADDRESS	Indicates the destination MAC Address of the unicast FIB entry.
BVLAN	Indicates the B-VLAN of the unicast FIB entry.
SYSID	Indicates the destination system identifier of the unicast FIB entry.
HOST-NAME	Indicates the destination host name of the unicast FIB entry.
OUTGOING INTERFACE	Indicates the outgoing interface of the unicast FIB entry.
COST	Indicates the cost of the unicast FIB entry.

Displaying IS-IS LSDB and adjacencies

Use the following procedure to display the IS-IS LSDB and adjacencies.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the IS-IS LSDB:

   ```
   show isis lsdb [detail][ip-unicast][level {l1|l2|l12}] [local][lspid
   <xxxx.xxxx.xxxx.xx-xx>] [sysid <xxxx.xxxx.xxxx>] [tlv <1-186>]
   ```

3. Display IS-IS LSDB IP unicast information:

   ```
   show isis lsdb ip-unicast [i-sid <0-16777215>][lspid
   <xxxx.xxxx.xxxx.xx-xx>][sysid <xxxx.xxxx.xxxx>]
   ```

4. Display IS-IS adjacencies:

   ```
   show isis adjacencies
   ```

5. Enter Privileged EXEC mode:

   ```
   enable
   ```

6. Clear IS-IS LSDB:

   ```
   clear isis lsdb
   ```

Example

Display LSDB information:

```
VSP-9012:1#show isis lsdb
================================================================================
                                   ISIS
LSDB
================================================================================
LSP ID                          LEVEL     LIFETIME  SEQNUM     CHKSUM    HOST-NAME
--------------------------------------------------------------------------------
0014.c7e1.33df.00-00            1         545       0xb1       0xed28    NewYork
```

```
0016.ca23.73df.00-00           1         1119      0x9f       0x9c9d   VSP-Lab2
0018.b0bb.b3df.00-00           1         708       0xb9       0xcb1a   VSP-Lab1

--------------------------------------------------------------------------------
Level-1 : 3 out of 3 Total Num of LSP Entries
Level-2 : 0 out of 0 Total Num of LSP Entries
```

Display IS-IS adjacencies:

```
VSP-9012:1# show isis adjacencies
================================================================================
                               ISIS Adjacencies
================================================================================
INTERFACE   L STATE   UPTIME         PRI    HOLDTIME SYSID                HOST-NAME
--------------------------------------------------------------------------------
Mlt2          1 UP      1d 03:57:25 127    20             0018.b0bb.b3df       VSP-Lab1
Port3/21      1 UP      1d 03:57:16 127    27             0016.ca23.73df       VSP-Lab2
--------------------------------------------------------------------------------
 2 out of 2 Total Num of Adjacencies
--------------------------------------------------------------------------------
```

Display IS-IS LSDB detail:

```
VSP-9012:1>show isis lsdb detail

================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------

Level-1 LspID: 0001.bcb0.0003.00-001     SeqNum: 0x00000522      Lifetime:  1144
        Chksum: 0x32f7  PDU Length: 312
        Host_name: C0
        Attributes:     IS-Type 1
TLV:1   Area Addresses: 1
                c1.3000.0000.00

TLV:22  Extended IS reachability:
        Adjacencies: 7
        TE Neighbors: 7
                0000.beb1.0007.01 (VSP0)         Metric:10
                        SPBM Sub TLV:
                                port id: 640 num_port 1
                                Metric: 10
                0000.beb1.00b1.01 (VSP1)        Metric:10

                        SPBM Sub TLV:

                                port id: 643 num_port 1

                                Metric: 10

                0000.bcb1.0004.01 (C1)  Metric:10

                        SPBM Sub TLV:

                                port id: 6144 num_port 1

                                Metric: 10

                0000.beb1.00ca.01 (VSP2)        Metric:10

                        SPBM Sub TLV:
```

```
                                port id: 6156 num_port 1

                                Metric: 10

                0000.beb1.00a5.01 (VSS0)        Metric:10

                        SPBM Sub TLV:

                                port id: 651 num_port 1

                                Metric: 10

                0000.beb1.00b2.01 (VSS1)        Metric:10

                        SPBM Sub TLV:

                                port id: 645 num_port 1

                                Metric: 10

                0000.beb1.0008.01 (VSP1)         Metric:10

                        SPBM Sub TLV:

                                port id: 652 num_port 1

                                Metric: 10

TLV:129 Protocol Supported: SPBM

TLV:137 Host_name: C0#

TLV:144 SUB-TLV 1       SPBM INSTANCE:
                Instance: 0
                bridge_pri: 0
                OUI: 00-33-33
                num of trees: 2
                vid tuple : u-bit 1 m-bit 1 ect-alg 0x80c201 base vid 1000
                vid tuple : u-bit 1 m-bit 1 ect-alg 0x80c202 base vid 1001
TLV:144 SUB-TLV 3       ISID:
                Instance: 0
                Metric: 0
                B-MAC: 00-00-bc-b1-00-03
                BVID:1000
                Number of ISID's:8
                        3001(Both),3002(Rx),3003(Both),3004(Rx),4001(Both),4002(
Rx),4003(Both),4004(Rx)


                Instance: 0
                Metric: 0
                B-MAC: 00-00-bc-b1-00-03

--More-- (q = quit)
```

Display IS-IS LSDB IP unicast information:

```
VSP-9012:1#show isis lsdb ip-unicast

================================================================================
                         ISIS IP-UNICAST-ROUTE SUMMARY
================================================================================
                           PREFIX          TLV   LSP   HOST
I-SID     ADDRESS          LENGTH  METRIC  TYPE  FRAG  NAME
--------------------------------------------------------------------------------
```

```
-          192.0.2.5         32       1        135    0x2    VSP5
-          192.0.2.0         24       1        135    0x2    VSP5
111        192.0.2.111       32       1        184    0x3    VSP5
111        192.0.2.0         24       1        184    0x3    VSP5
7          192.0.2.0         24       1        184    0x3    VSP5
-          198.51.100.200    32       1        135    0x2    VSP200
7          198.51.100.7      32       1        184    0x3    VSP200
7          198.51.100.10     32       1        184    0x3    VSP200
11         198.51.100.11     32       1        184    0x3    VSP200
11         198.51.100.25     32       1        184    0x3    VSP200
11         198.51.100.0      24       1        184    0x3    VSP200
--------------------------------------------------------------------------------
11 out of 11 Total Num of Entries
```

Display IS-IS LSDB IP unicast information by I-SID:

```
VSP-9012:1#show isis lsdb ip-unicast i-sid 7

================================================================================
                          ISIS IP-UNICAST-ROUTE SUMMARY
================================================================================
                            PREFIX          TLV   LSP   HOST
I-SID     ADDRESS           LENGTH  METRIC  TYPE  FRAG  NAME
--------------------------------------------------------------------------------
7          192.0.2.0         24       1        184    0x3    VSP5
7          198.51.100.7      32       1        184    0x3    VSP200
7          198.51.100.10     32       1        184    0x3    VSP200
--------------------------------------------------------------------------------
3 out of 11 Total Num of Entries
```

Display IS-IS LSDB IP unicast information by LSP ID:

```
VSP-9012:1#show isis lsdb ip-unicast lspid 0200.0200.0200.00-03

================================================================================
                          ISIS IP-UNICAST-ROUTE SUMMARY
================================================================================
                            PREFIX          TLV   LSP   HOST
I-SID     ADDRESS           LENGTH  METRIC  TYPE  FRAG  NAME
--------------------------------------------------------------------------------
7          198.51.100.7      32       1        184    0x3    VSP200
7          198.51.100.10     32       1        184    0x3    VSP200
11         198.51.100.11     32       1        184    0x3    VSP200
11         198.51.100.25     32       1        184    0x3    VSP200
11         198.51.100.0      24       1        184    0x3    VSP200
--------------------------------------------------------------------------------
5 out of 11 Total Num of Entries
```

Display IS-IS LSDB IP unicast information by system ID:

```
VSP-9012:1#show isis lsdb ip-unicast sysid 0005.0005.0005

================================================================================
                          ISIS IP-UNICAST-ROUTE SUMMARY
================================================================================
                            PREFIX          TLV   LSP   HOST
I-SID     ADDRESS           LENGTH  METRIC  TYPE  FRAG  NAME
--------------------------------------------------------------------------------
-          192.0.2.5         32       1        135    0x2    VSP5
-          192.0.2.0         24       1        135    0x2    VSP5
111        192.0.2.111       32       1        184    0x3    VSP5
111        192.0.2.0         24       1        184    0x3    VSP5
7          192.0.2.0         24       1        184    0x3    VSP5
--------------------------------------------------------------------------------
5 out of 11 Total Num of Entries
```

Variable definitions

Use the data in the following table to use the **show isis lsdb** command.

Variable	Value
detail	Displays detailed information.
ip-unicast	Displays IS-IS LSDB IP unicast information.
level *{l1|l2|l12}]*	Displays the LSDB for the specified level: l1, l2, or l12.
local	Displays information on the local LSDB.
lspid <*xxxx.xxxx.xxxx.xx-xx*>	Displays the LSDB for the specified LSP ID.
sysid <*xxxx.xxxx.xxxx*>	Displays the LSDB for the specified system ID.
tlv <*1–186*>	Displays the LSDB by TLV type.

Use the data in the following table to use the **show isis lsdb ip-unicast** command.

Variable	Value
i-sid <0–16777215>	Displays information on the local LSDB.
lspid <*xxxx.xxxx.xxxx.xx-xx*>	Displays the LSDB for the specified LSP ID.
sysid <*xxxx.xxxx.xxxx*>	Displays the LSDB for the specified system ID.

Use the data in the following table to use the **clear isis** command.

Variable	Value
lsdb	Clears the IS-IS Link State Database (LSDB). The command clears learned LSPs only. The command does not clear local generated LSPs. As soon as the platform clears the LSDB the LSP synchronization process starts immediately and the LSDB synchronizes with its neighbors.

Job aid

The following sections describe the fields in the outputs for the IS-IS LSDB and adjacencies show commands.

show isis lsdb

The following table describes the fields in the output for the **show isis lsdb** command.

Parameter	Description
LSP ID	Indicates the LSP ID assigned to external IS-IS routing devices.
LEVEL	Indicates the level of the external router: l1, l2, or l12.
LIFETIME	Indicates the maximum age of the LSP. If the max-lsp-gen-interval is set to 900 (default) then the lifetime value begins to count down from 1200 seconds and updates after 300 seconds if connectivity remains. If the timer counts down to zero, the counter adds on an additional 60 seconds, then

Parameter	Description
	the LSP for that router is lost. This happens because of the zero age lifetime, which is detailed in the RFC standards.
SEQNUM	Indicates the LSP sequence number. This number changes each time the LSP is updated.
CHKSUM	Indicates the LSP checksum. This is an error checking mechanism used to verify the validity of the IP packet.
HOST-NAME	Indicates the hostname listed in the LSP. If the host name is not configured, then the system name is displayed.

show isis adjacencies

The following table describes the fields in the output for the **`show isis adjacencies`** command.

Parameter	Description
INTERFACE	Indicates the interface port or MLT on which IS-IS exists.
L	Indicates the level of the adjacent router.
STATE	Indicates the state of IS-IS on the interface (enabled [UP] or disabled [DOWN]). The state is non-configurable.
UPTIME	Indicates the length of time the adjacency has been up in ddd hh:mm:ss format.
PRI	Indicates the priority of the neighboring Intermediate System for becoming the Designated Intermediate System (DIS).
HOLDTIME	Indicates the calculated hold time for the Hello (hello multiplier x hello interval); if the route is determined to be a designated router, then the product is divided by 3.
SYSID	Indicates the adjacent system ID of the router.
HOST-NAME	Indicates the hostname listed in the LSP. If the host name is not configured, then the system name is displayed.

Displaying IS-IS statistics and counters

Use the following procedure to display the IS-IS statistics and counters.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display IS-IS system statistics:

   ```
   show isis statistics
   ```

3. Display IS-IS interface counters:

   ```
   show isis int-counters
   ```

4. Display IS-IS level 1 control packet counters:

```
show isis int-l1-cntl-pkts
```

Note:

The current release uses level 1 IS-IS. The current release does not support level 2 IS-IS. The ACLI command **`show isis int-l2-contl-pkts`** is not supported in the current release because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.

5. Enter Privileged EXEC mode:

   ```
   enable
   ```

6. Clear IS-IS statistics:

   ```
   clear isis stats [error-counters] [packet-counters]
   ```

Example

```
VSP-9012:1# show isis statistics
================================================================================
                              ISIS System Stats
================================================================================

LEVEL    CORR   AUTH   AREA   MAX SEQ   SEQ NUM  OWN LSP  BAD ID  PART   LSP    DB
            LSPs  FAILS  DROP   EXCEEDED   SKIPS    PURGE    LEN    CHANGES OLOAD
--------------------------------------------------------------------------------
Level-1  0      0      0      0        1        0        0       0       0
```

```
VSP-9012:1#show isis int-counters
================================================================================
                            ISIS Interface Counters
================================================================================

IFIDX    LEVEL   AUTH     ADJ          INIT      REJ     ID LEN  MAX AREA LAN    DIS
                 FAILS    CHANGES           FAILS       ADJ                      CHANGES
--------------------------------------------------------------------------------
Mlt2     Level 1-2  0        1               0       0           0       0        0
Port3/21 Level 1-2  0        1               0       0           0       0        0
```

```
VSP-9012:1#show isis int-l1-cntl-pkts
================================================================================
                        ISIS L1 Control Packet counters
================================================================================
IFIDX         DIRECTION     HELLO        LSP          CSNP         PSNP

--------------------------------------------------------------------------------
Mlt2           Transmitted       13346        231          2            229
Mlt2           Received          13329        230          1            230
Port3/21       Transmitted       13340        227          2            226
Port3/21       Received          13335        226          1            227
```

Variable definitions

Use the data in the following table to use the **`clear isis stats`** command.

Variable	Value
error-counters	Clears IS-IS stats error-counters.

Variable	Value
packet-counters	Clears IS-IS stats packet-counters.

Job aid

show isis statistics

The following table describes the fields in the output for the **`show isis statistics`** command.

Parameter	Description
LEVEL	Shows the level of the IS-IS interface (Level 1 in the current release).
CORR LSPs	Shows the number of corrupted LSPs detected.
AUTH FAILS	Shows the number of times authentication has failed on the global level.
AREA DROP	Shows the number of manual addresses dropped from the area.
MAX SEQ EXCEEDED	Shows the number of attempts to exceed the maximum sequence number.
SEQ NUM SKIPS	Shows the number of times the sequence number was skipped.
OWN LSP PURGE	Shows how many times the local LSP was purged.
BAD ID LEN	Shows the number of ID field length mismatches.
PART CHANGES	Shows the number of partition link changes.
LSP DB OLOAD	Show the number of times the Virtual Services Platform 9000 was in the overload state.

show isis int-counters

The following table describes the fields in the output for the **`show isis int-counters`** command.

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.
LEVEL	Shows the level of the IS-IS interface (Level 1 in the current release).
AUTH FAILS	Shows the number of times authentication has failed per interface.
ADJ CHANGES	Shows the number of times the adjacencies have changed.
INIT FAILS	Shows the number of times the adjacency has failed to establish.
REJ ADJ	Shows the number of times the adjacency was rejected by another router.
ID LEN	Shows the ID field length mismatches.
MAX AREA	Shows the maximum area address mismatches.
LAN DIS CHANGES	Shows the number of times the DIS has changed.

show isis int-l1-cntl-pkts

The following table describes the fields in the output for the **`show isis int-l1-cntl-pkts`** command.

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.

Parameter	Description
DIRECTION	Shows the packet flow (Transmitted or Received).
HELLO	Shows the amount of interface-level Hello packets.
LSP	Shows the amount of LSP packets.
CSNP	Shows the amount of CSNPs.
PSNP	Shows the amount of PSNPs.

Displaying SPBM packet drop statistics by port

Use this procedure to display the SPBM drop statistics.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the last dropped statistics:

   ```
   show isis spbm drop-stats port last-drop [{slot/port[-slot/port]
   [,...]}]
   ```

3. Display the RPFC multicast source MAC drop statistics:

   ```
   show isis spbm drop-stats port rpfc-multicast-sa [{slot/port[-slot/
   port][,...]}]
   ```

4. Display the RPFC unicast source MAC drop statistics:

   ```
   show isis spbm drop-stats port rpfc-unicast-sa [{slot/port[-slot/
   port][,...]}]
   ```

5. Display the unknown multicast destination MAC drop statistics:

   ```
   show isis spbm drop-stats port unknown-multicast-da [{slot/port[-
   slot/port][,...]}]
   ```

6. Display the unknown unicast destination MAC drop statistics:

   ```
   show isis spbm drop-stats port unknown-unicast-da [{slot/port[-slot/
   port][,...]}]
   ```

7. Display the unknown unicast source MAC drop statistics:

   ```
   show isis spbm drop-stats port unknown-unicast-sa [{slot/port[-slot/
   port][,...]}]
   ```

   Note:

   Virtual Services Platform 9000 does not support the **`show isis spbm drop-stats port unknown-unicast-sa`** drop count parameter for second generation modules in

this release. The device always displays the ACLI output for second generation modules as 0 for this counter.

Example

The following output shows examples for the **show isis spbm drop-stats port** command.

```
VSP-9012:1#show isis spbm drop-stats port last-drop

================================================================================
                          SPBM Drop Stats By Port
                                Last Drop
================================================================================
PORT  PRIMARY B-VID                          SECONDARY B-VID
NUM   HOST NAME        DA  B-MAC            HOST NAME        DA  B-MAC
--------------------------------------------------------------------------------
3/11 evp              N  00:13:0a:e6:73:df  evp              N  00:13:0a:e6:73:df

3/13 evs              N  00:13:0a:e6:43:df  evs              N  00:13:0a:e6:43:df
```

```
VSP-9012:1#show isis spbm drop-stats port unknown-unicast-sa 3/11
================================================================================
                          SPBM Drop Stats By Port
                       Unknown Unicast Source Address
================================================================================
PORT     PRIMARY B-VID    SECONDARY B-VID
NUM      PKT DROP         PKT DROP
--------------------------------------------------------------------------------
3/11     4                4
```

```
VSP-9012:1#show isis spbm drop-stats port rpfc-unicast-sa

================================================================================
                          SPBM Drop Stats By Port
            Reverse Path Forwarding Check Unicast Source Address
================================================================================
PORT     PRIMARY B-VID    SECONDARY B-VID
NUM      PKT DROP         PKT DROP
--------------------------------------------------------------------------------
3/11     0                0

3/13     0                0
```

Variable definitions

Use the data in the following table to use the **show isis spbm drop-stats port** commands.

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).

Job aid

The following table describes the fields in the output for the **show isis spbm drop-stats port last-drop** command.

Table 1: show isis spbm drop-stats port last-drop field descriptions

Field	Description
PORT NUM	Shows the slot/port number that identifies the ingress port of the dropped packet.
PRIMARY B-VID HOST NAME	Shows the name of the primary SPBM B-VID.
DA	Shows whether there were dropped packets. • `Y` indicates that the last drop B-MAC is a destination MAC. • `N` indicates that the last drop B-MAC is a source MAC.
B-MAC	Shows the backbone MAC address of the primary SPBM B-VLAN.
SECONDARY B-VID HOST NAME	Shows the name of the secondary SPBM B-VID, if known.
DA	Shows whether there were dropped packets: Yes (Y) or No (N).
B-MAC	Shows the backbone MAC address of the secondary SPBM B-VLAN.

The following table describes the fields in the output for all of the other **`show isis spbm drop-stats port`** commands.

Table 2: show isis spbm drop-stats port field descriptions

Field	Description
PORT NO	Shows the slot/port number that identifies the ingress port of the dropped packet.
PRIMARY B-VID PKT DROP	Shows the total number of RPFC multicast drops for each primary SPBM B-VLAN.
SECONDARY B-VID PKT DROP	Shows the total number of RPFC multicast drops for each secondary SPBM B-VLAN.

Clearing SPBM packet drop statistics

Clear drop statistics to reset the counters.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Clear SPBM port-based drop statistics:

   ```
   clear isis spbm drop-stats [{slot/port[-slot/port][,...]} ]
   ```

Variable definitions

Use the data in the following table to use the **`clear isis spbm drop-stats`** 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).

SPBM and IS-IS infrastructure configuration using EDM

This section provides procedures to configure basic SPBM and IS-IS infrastructure using Enterprise Device Manager (EDM).

Configuring required SPBM and IS-IS parameters

Use the following procedure to configure the minimum required SPBM and IS-IS parameters to allow SPBM to operate on the switch. 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.

Note:

Virtual Services Platform 9000 does not support IPFIX on IS-IS interfaces.

Note:

SPB internally uses spanning tree group (STG) 63 or Multiple Spanning Tree Instance (MSTI) 62. STG 63 or MSTI 62 cannot be used by another VLAN or MSTI. For non-SPB customer networks, if you use STG 63 or MSTI 62 in the configuration, you must delete STG 63 or MSTI 62 before you can configure SPBM.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.
3. From the **Globals** tab, select **enable** to enable SPBM globally, and click **Apply**.
4. Click the **SPBM** tab.
5. Click **Insert** to create an SPBM instance (in this release, only one SPBM instance is supported).
6. In the **Id** field, specify the SPBM instance ID.
7. In the **NodeNickName** field, specify the node nickname (valid value is 2.5 bytes in the format <x.xx.xx>)
8. Click **Insert**.

9. From the navigation tree, select **Configuration** > **IS-IS** > **SPBM**.
10. Click the **SPBM** tab.
11. In the **Vlans** field, specify the IDs of the SPBM B-VLANs to add to the SPBM instance.
12. In the **PrimaryVlan** field, specify which of the SPBM B-VLANs specified in the previous step is the primary B-VLAN.
13. Click **Apply**.
14. From the navigation tree, select **Configuration** > **VLAN** > **VLANs**.
15. Click the **Basic** tab.
16. Click **Insert**.
17. In the **Type** field, click **spbm-bvlan**.
18. Click **Insert**.
19. In the navigation tree, expand the following folders:**Configuration** > **IS-IS** > **IS-IS**.
20. Click the **Manual Area** tab.
21. In the Manual Area tab, click **Insert** to add a manual area (in this release, only one manual area is supported).
22. Specify the Manual Area Address (valid value is 1–13 bytes in the format <xx.xxxx.xxxx...xxxx>).
23. Click **Insert**.
24. Under the IS-IS tab, click the **Globals** tab.

    Note:

    Although it is not strictly required for SPBM operation, Avaya recommends that you change the IS-IS system ID from the default B-MAC value to a recognizable address to easily identify a switch (using the **SystemID** field under the IS-IS Globals tab) . This helps to recognize source and destination addresses for troubleshooting purposes.
25. In the AdminState field, click **on**, and click **Apply**.
26. Under the IS-IS tab, click the **Interfaces** tab.
27. Click **Insert** to create an IS-IS circuit.
28. In the **IfIndex** field, specify the port or MLT on which to create the IS-IS circuit.
29. Click **Insert**.
30. Select the newly created IS-IS circuit entry, and click **SPBM**.
31. In the **Interfaces SPBM** tab, click **Insert**.
32. In the **SpbmId** field, type the correct SPBM ID.
33. In the **State** field, select **enable**.
34. Click **Insert** to enable the SPBM instance on the IS-IS circuit.

35. Under the IS-IS tab, click the **Interfaces** tab.
36. In the **AdminState** field for the IS-IS circuity entry, select **on** to enable the IS-IS circuit.
37. Click **Apply**.

SPBM field descriptions

Note:

The following tables list the minimum required SPBM and IS-IS parameters to allow SPBM to operate on the switch. For more detailed information on all of the parameters see the procedures that follow. For more information on how to configure VLANs, see *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500.

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

Name	Description
GlobalEnable	Enables or disables SPBM globally.
GlobalEtherType	Specifies the global Ethertype value as 0x8100 or 0x88a8. The default value is 0x8100.

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

Name	Description
Id	Specifies the SPBM instance ID. In this release, only one SPBM instance is supported.
NodeNickName	Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Specifies the primary SPBM B-VLANs to add to the SPBM instance.
Vlans	Specifies the SPBM B-VLANs to add to the SPBM instance.

Use the data in the following table to use the **VLANs Basic** tab.

Name	Description
Type	Specifies the type of VLAN: • byPort • byIpSubnet • byProtocolId • bySrcMac • spbm-bvlan

Use the data in the following table to use the **IS-IS Manual Area** tab.

Name	Description
AreaAddr	Specifies the IS-IS manual area. Valid value is 1–13 bytes in the format <xx.xxx.xxx...xxx>. In this release, only one manual area is supported. For IS-IS to operate, you must configure at least one manual area.

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

Name	Description
AdminState	Specifies the global status of IS-IS on the switch: on or off. The default is off.
SystemId	Specifies the system ID. **Note:** Although it is not strictly required for SPBM operation, Avaya recommends that you change the IS-IS system ID from the default B-MAC value to a recognizable address to easily identify a switch (using the **SystemID** field under the IS-IS Globals tab) . This helps to recognize source and destination addresses for troubleshooting purposes.

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

Name	Description
IfIndex	The identifier of this circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value. This object cannot be modified after creation.
AdminState	Specifies the administrative state of the circuit: on or off. The default is off.

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

Name	Description
SpbmId	Specifies the SPBM instance ID.
State	Specifies whether the SPBM interface is enabled or disabled.

Job aid

Important:

After you have configured the SPBM nickname and enabled IS-IS, if you require a change of the system ID, you must also change the nickname. However, for naming convention purposes or

configuration purposes, you may not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps:

1. Disable IS-IS.
2. Change the system ID.
3. Change the nickname to a temporary one.
4. Enable IS-IS.
5. Disable IS-IS.
6. Change the nickname to the original nickname.
7. Enable IS-IS.

Configuring IP multicast over SPBM globally

Use this procedure to globally enable IP multicast over SPBM on the Backbone Edge Bridges (BEBs) that directly or indirectly (using Layer 2 switches) connect to IP multicast senders or receivers. By default, IP multicast over SPBM is disabled. There is no need to enable IP multicast over SPBM on the Backbone Core Bridges (BCBs).

You must configure IP multicast over SPBM at the global level, and then enable it on the service option or options you choose.

Important:

SPBM multicast uses I-SIDs that start at 16,000,000 and above. The device displays an error message if the Layer 2 and Layer 3 I-SIDs are within this range and the system does not enable IP multicast over SPBM.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology.

Procedure

1. Determine if any I-SIDs are within the default range reserved for multicast. From the navigation tree, expand the following folders: **Configuration** > **IS-IS** > **SPBM**.
2. Click the **I-SID** tab to determine if the I-SIDs are within the default range reserved for multicast.
3. From the navigation tree, expand the following folders: **Configuration** > **IS-IS** > **SPBM**.
4. Click the **SPBM** tab.
5. If you want to enable multicast on an SPBM instance that already exists, in the **Mcast** column in the table, select **enable**.

6. If you want to enable multicast on an SPBM instance that does not yet exit, click **Insert**.
7. In the **Mcast** box, select **enable** to enable IP multicast over SPBM globally.
8. Click **Insert**.
9. Click **Apply**.

SPBM SPBM field descriptions

Use the data in the following table to use the **SPB Multicast** tab.

Name	Description
Id	Specifies the SPBM instance ID. In this release, only one SPBM instance is supported.
NodeNickName	Specifies a nickname for the SPBM instance globally.
PrimaryVlan	Specifies the primary SPBM B-VLAN to add to the SPBM instance.
Vlans	Specifies the SPBM B-VLANs to add to the SPBM instance.
LsdbTrap	Specifies if the LSDB update trap is enabled on this SPBM instance. The default is disabled.
IpShortcut	Specifies if SPBM IP Shortcuts is enabled. The default is disabled.
SmltSplitBEB	Specifies the SMLT split BEB for this SPBM instance.
SmltVirtualBmac	Specifies the SMLT virtual MAC for this SPBM instance.
SmltPeerSysId	Specifies the SMLT peer system ID for this SPBM instance.
Mcast	Specifies if IP multicast over SPBM is enabled. The default is disabled.
McastFwdCacheTimeout	Specifies the global forward cache timeout in seconds. The default is 210 seconds.

Modifying IP multicast over SPBM globally

Use this procedure to modify IP multicast over SPBM globally on the Backbone Edge Bridges (BEBs) that directly or indirectly (using Layer 2 switches) connect to IP multicast senders or receivers. By default, IP multicast over SPBM is disabled. There is no need to enable IP multicast over SPBM on the Backbone Core Bridges (BCBs).

You must configure IP multicast over SPBM at the global level, and then enable it on the service option or options you choose.

Important:

SPBM multicast uses I-SIDs that start at 16,000,000 and above. The device displays an error message if the Layer 2 and Layer 3 I-SIDs are within this range and the system does not enable IP multicast over SPBM.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.

- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **IS-IS** > **SPBM**.
2. Click the **SPBM** tab.
3. Double-click in the **Mcast** cell, select **enable** or **disable**.
4. Click **Apply**.

Displaying IP multicast over SPBM routes

Use this procedure to display IP multicast over SPBM routes.

Procedure

1. From the navigation tree, expand the following: **Configuration** > **IS-IS** > **SPBM**.
2. Click the **IpMcastRoutes** tab.

IpMcastRoutes field descriptions

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

Name	Description
Group	Specifies the group IP address for the IP multicast over SPBM route.
Source	Specifies the IP address where the IP multicast over SPBM route originated.
VsnIsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.
SourceBeb	Specifies the source BEB for the IP multicast route.
VlanId	Specifies the ID for the C-VLAN.
VrfName	Specifies the VRF name.
DataIsid	Specifies the data I-SID for the IP multicast over SPBM route. A a BEB receives IP multicast data from a sender, a BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
NniPorts	Specifies the NNI ports for the IP multicast route. SPBM runs in the core on the ports that connect to the core. These ports are NNI ports. Ports that face a

Name	Description
	customer VLAN are user-to-network interface (UNI) ports.
Type	Specifies the type for the IP multicast over SPBM route.
Bvlan	Specifies the B-VLAN for the IP multicast over SPBM route.

Displaying the UNI ports for IP multicast routes

Use this procedure to display UNI ports associated with particular IP multicast routes.

Procedure

1. From the navigation tree, expand the following: **Configuration** > **IS-IS** > **SPBM**.
2. Click the **IpMcastRoutes** tab.
3. Select the desired row and click the **UNI Ports** tab to display the UNI ports associated with a particular stream.

IpMcastRoutes Uni Ports field descriptions

Use the data in the following table to use the **IpMcastRoutes Uni Ports** tab.

Name	Description
Group	Specifies the group IP address for the IP multicast over SPBM route.
Source	Specifies the IP address where the IP multicast over SPBM route originated.
VsnIsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.
DataIsid	Specifies the data I-SID for the IP multicast route. After a BEB receives the IP multicast data from a sender, a BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
SourceBeb	Specifies the source BEB for the IP multicast route.
VlanId	Specifies the ID for the C-VLAN.
VrfName	Specifies the VRF name.
NniPorts	Specifies the NNI ports for the IP multicast route. SPBM runs in the core on the ports that connect to the core. These ports are NNI ports. Ports facing a

Name	Description
	customer VLAN are user-to-network interface (UNI) ports.
Type	Specifies the type for the IP multicast route.
Bvlan	Specifies the B-VLANs for the IP multicast route.

Displaying SPBM and IS-IS summary information

Use the following procedure to view a summary of SPBM and IS-IS protocol information.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Protocol Summary** tab.

Protocol Summary field descriptions

Use the data in the following table to use the **Protocol Summary** tab.

Name	Description
Globals ISIS	
AdminState	Indicates the global status of IS-IS on the switch.
SystemId	Indicates the IS-IS system ID for the switch. Valid value is a 6–byte value in the format <xxxx.xxxx.xxxx>
HostName	Indicates a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system name comes from the host name configured at the system level.
Globals SPBM	
GlobalEnable	Indicates whether SPBM is enabled or disabled at the global level.
NodeNickName	Indicates the nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Indicates the primary VLAN ID for this SPBM instance.
SmltSplitBEB	Indicates whether the switch is the primary or secondary IST peer.
ISIS Interfaces	
Circuit Index	Displays the identifier of this IS-IS circuit, unique within the Intermediate System. This is for SNMP Indexing purposes only and need not have any relation to any protocol value.
IfIndex	Indicates the interface to which this circuit corresponds.

Name	Description
AdminState	Indicates the administrative state of the circuit: on or off.
OperState	Indicates the operational state of the circuit: up or down.
ISIS Adjacency View	
Circuit Index	Displays the identifier of this IS-IS circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
AdjIndex	Displays a unique value identifying the IS adjacency from all other such adjacencies on this circuit. This value is automatically assigned by the system when the adjacency is created
AdjIfIndex	Indicates the interface to which this circuit corresponds.
AdjState	Indicates the state of the adjacency: • down • initializing • up • failed
AdjNeighSysID	Indicates the system ID of the neighboring Intermediate System.
AdjHostName	Indicates the host name listed in the LSP, or the system name if the host name is not configured.

Displaying the SPBM I-SID information

Use the following procedure to display the SPBM Service Instance Identifier (I-SID) information. The SPBM B-MAC header includes an I-SID with a length of 24 bits. This I-SID can be used to identify and transmit any virtualized traffic in an encapsulated SPBM frame.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.
3. Click the **I-SID** tab.

I-SID field descriptions

Use the data in the following table to use the **I-SID** tab.

Name	Description
SysId	Indicates the system identifier.
Vlan	Indicates the B-VLAN where this I-SID was configured or discovered.
Isid	Indicates the IS-IS SPBM I-SID identifier.

Name	Description
NickName	Indicates the nickname of the node where this I-SID was configured or discovered.
HostName	Indicates the host name listed in the LSP, or the system name if the host name is not configured.
Type	Indicates the SPBM I-SID type; either configured or discovered.

Displaying Level 1 Area information

Use the following procedure to display Level 1 area information. IS-IS provides support for hierarchical routing, which enables you to partition large routing domains into smaller areas. IS-IS uses a two-level hierarchy, dividing the domain into multiple Level 1 areas and one Level 2 area. The Level 2 area serves as backbone of the domain, connecting to all the Level 1 areas.

Important:

The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled in the current release.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **L1 Area** tab.

L1 Area field descriptions

Use the data in the following table to use the **L1 Area** tab.

Name	Description
AreaAddr	Specifies an area address reported in a Level 1 link-state packets (LSP) generated or received by this Intermediate System.

Configuring SMLT parameters for SPBM

Use the following procedure to configure the required Split MultiLink Trunking (SMLT) parameters to allow SPBM to interoperate with SMLT on the switch.

Note:

If you want to modify SMLTs that contain NNI ports, Avaya recommends you do the modification during maintenance windows. Otherwise, if you create or delete SMLTs that contain NNI ports on Virtual Services Platform 9000 running MSTP, IS-IS adjacencies that connect to those ports can bounce even if the SMLT is not used.

Note:

- The assignment of primary and secondary roles to the IST peers is automatic. The switch with the lower system ID (between the two IST peers) is primary, and the switch with the higher system ID is secondary.
- SMLT peer system ID is part of the required configuration. You must configure the SMLT peer system ID as the nodal MAC of the peer device. In the IS-IS network, the nodal MAC of devices should be eight apart from each other.
- SMLT virtual B-MAC is an optional configuration. If SMLT virtual B-MAC is not configured, the system derives SMLT virtual B-MAC from the configured SMLT peer system ID and the nodal MAC of the device (IS-IS system ID). The system compares the nodal MAC of the device with the SMLT peer system ID configured and takes the small one, plus 0x01, as the SMLT virtual B-MAC.
- Different from Ethernet Routing Switch 8800 SPBM implementation, on Virtual Services Platform 9000 when you map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID), you must include the IST MLT peer switches in the VLAN, which means all IST member ports must be members of this C-VLAN. In ERS 8800, you cannot configure IST MLT to be part of the C-VLAN. For more information about this and other differences between the two products, see *Platform Migration Reference for Avaya Virtual Services Platform 9000,* NN46250-107.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS** > **SPBM**.
2. Click the **SPBM** tab.
3. Use the **SmltSplitBEB** field to see whether the switch is the primary or secondary IST peer. This field cannot be modified.
4. Use the **SmltVirtualBmac** field to specify a virtual MAC address that can be used by both peers.
5. Use the **SmltPeerSysId** field to specify the IST peer B-MAC address.
6. Click **Apply**.

Enabling or disabling SPBM at the global level

Use the following procedure to enable or disable SPBM at the global level. 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.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.

3. Click the **Globals** tab.
4. To enable or disable SPBM, click **enable** or **disable** in the **GlobalEnable** field.
5. To configure the global ethertype value, click the desired option in the **GlobalEtherType** field.
6. Click **Apply**.

Globals field descriptions

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

Name	Description
GlobalEnable	Enables or disables SPBM globally. The default is disabled.
GlobalEtherType	Specifies the global ethertype value as 0x8100 or 0x88a8. The default value is 0x8100.

Configuring SPBM parameters

Use the following procedure to configure SPBM global parameters. 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.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.
3. Click the **SPBM** tab.
4. To create an SPBM instance, click **Insert**.
5. Configure the SPBM parameters.
6. Click **Apply**.

SPBM field descriptions

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

Name	Description
Id	Specifies the SPBM instance ID. In this release, only one SPBM instance is supported.
NodeNickName	Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Specifies the primary SPBM B-VLANs to add to the SPBM instance.
Vlans	Specifies the SPBM B-VLANs to add to the SPBM instance.

Name	Description
LsdbTrap	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable.
IpShortcut	Enables or disables SPBM IP shortcut state. The default is disable.
SmltSplitBEB	Specifies whether the switch is the primary or secondary IST peer. The default is primary.
SmltVirtualBmac	Specifies a virtual MAC address that can be used by both peers.
SmltPeerSysId	Specifies the system ID of the SPBM SMLT for this SPBM instance.
Mcast	Specifies if multicast over SPBM is enabled.
McastFwdCacheTimeout	Specifies the multicast forward-cache timeout value, in seconds. The default is 210.

Displaying SPBM nicknames

Use the following procedure to display SPBM nicknames.

If you want to display link-state packet (LSP) summary information, see Displaying LSP summary information on page 101.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.
3. Click the **Nick Names** tab.

Nickname field descriptions

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

Name	Description
Level	Indicates the level at which this LSP appears.
ID	Indicates the 8 byte LSP ID, consisting of the SystemID, Circuit ID, and Fragment Number.
LifetimeRemain	Indicates the remaining lifetime in seconds for the LSP.
NickName	Indicates the nickname for the SPBM node.
HostName	Indicates the hostname listed in the LSP, or the system name if the host name is not configured.

Configuring interface SPBM parameters

Use the following procedure to configure SPBM interface parameters.

Procedure

1. In the navigation tree, expand the following folders:**Configuration** > **IS-IS**.
2. Click **SPBM**.
3. Click the **Interfaces SPBM** tab.
4. Configure the SPBM interface parameters.
5. Click **Apply**.

SPBM field descriptions

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

Name	Description
Index	Specifies an Index value for the SPBM interface.
State	Specifies whether the SPBM interface is enabled or disabled.
Type	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT: ptpt or bcast. In this release, only the point-to-point (ptpt) interface type is supported.
L1Metric	Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10.

Configuring SPBM on an interface

Use the following procedure to configure SPBM on an interface.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Click the **SPBM** button.
5. In the **Interfaces SPBM** tab, click **Insert**.
6. Click **Insert**.

Interfaces SPBM field descriptions

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

Name	Description
Index	Specifies an Index value for the SPBM interface.
SpbmId	Specifies the SPBM instance ID.
State	Specifies whether the SPBM interface is enabled or disabled. The default is disabled.
Type	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT. In this release, only the pt-pt interface type is supported. The default is pt-pt.
L1Metric	Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10.

Displaying the IP unicast FIB

Use the following procedure to display the IP unicast Forwarding Information Base (FIB). The tab shows IP routes from remote Backbone Edge Bridges (BEBs)

About this task

In SPBM, each node has a System ID, which also serves as Backbone MAC address (B-MAC) of the switch. These Backbone MAC addresses are populated into the SPBM VLAN Forwarding Information Base (FIB). When the network topology is discovered and stored in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

I-SIDs are only used for virtual services (Layer 2 VSNs and Layer 3 VSNs). If you only enable IP Shortcuts on the Backbone Edge Bridges, I-SIDs are never exchanged in the network as IP Shortcuts allows for Global Routing Table (GRT) IP networks to be transported across IS-IS.

The **IP Unicast FIB** tab displays all of the IS-IS routes in the IS-IS LSDB. The Preference column in the **IP Unicast FIB** tab displays the IP route preference.

Routes within the same VSN are added to the LSDB with a default preference of 7. Inter-VSN routes are added to the LSDB with a route preference of 200. IS-IS accept policies allow you to change the route preference for incoming routes. If the same route is learned from multiple sources with different route preferences, then the routes are not considered equal cost multipath (ECMP) routes. The route with the lowest route preference is the preferred route.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.
3. Click the **IP Unicast FIB** tab.

IP Unicast FIB field descriptions

Use the data in the following table to use the **IP Unicast FIB** tab.

Name	Description
VrfId	Specifies the VRF ID of the IP unicast FIB entry, 0 indicates NRE.
DestinationIpAddrType	Specifies the address type of the destination IP address.
DestinationIpAddr	Specifies the destination IP address of the IP unicast FIB entry.
DestinationMask	Specifies the destination IP mask of the IP unicast FIB entry
NextHopBmac	Specifies the nexthop B-MAC of the IP unicast FIB entry.
DestIsid	Specifies the destination I-SID of the IP unicast FIB entry.
Vlan	Specifies the VLAN of the IP unicast FIB entry.
Isid	Specifies the I-SID of the IP unicast FIB entry.
NextHopName	Specifies the nexthop hostname of the IP unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the IP unicast FIB entry.
PrefixCost	Specifies the prefix cost of the IP unicast FIB entry.
SpbmCost	Specifies the B-MAC cost of the IP unicast FIB entry.
Preference	Specifies the IP route preference.

Displaying the unicast FIB

Use the following procedure to display the unicast FIB.

In SPBM, B-MAC addresses are carried within the IS-IS link-state database. To do this, SPBM supports an IS-IS TLV that advertises the I-SID and B-MAC information across the network. Each node has a System ID, which also serves as Backbone MAC address (B-MAC) of the switch. These Backbone MAC addresses are populated into the SPBM VLAN Forwarding Information Base (FIB).

When the network topology is discovered and stored in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.
3. Click the **Unicast FIB** tab.

Unicast FIB field descriptions

Use the data in the following table to use the **Unicast FIB** tab.

Name	Description
SysId	Specifies the system ID of the node where the unicast FIB entry originated.

Name	Description
Vlan	Specifies the VLAN of the unicast FIB entry.
DestinationMacAddr	Specifies the destination MAC Address of the unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the unicast FIB entry.
HostName	Specifies the host name of the node where unicast FIB entry originated.
Cost	Specifies the cost of the unicast FIB entry.

Displaying the multicast FIB

Use the following procedure to display the multicast FIB.

In SPBM, B-MAC addresses are carried within the IS-IS link-state database. To do this, SPBM supports an IS-IS TLV that advertises the I-SID and B-MAC information across the network. Each node has a System ID, which also serves as Backbone MAC address (B-MAC) of the switch. These Backbone MAC addresses are populated into the SPBM VLAN Forwarding Information Base (FIB).

When the network topology is discovered and stored in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

The multicast FIB is not produced until virtual services are configured and learned.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.
3. Click the **Multicast FIB** tab.

Multicast FIB field descriptions

Use the data in the following table to use the **Multicast FIB** tab.

Name	Description
SysId	System ID of the node where the multicast FIB entry originated.
Vlan	VLAN of the multicast FIB entry.
McastDestMacAddr	Multicast destination MAC Address of the multicast FIB entry
Isid	I-SID of the multicast FIB entry.
OutgoingPorts	NNI port of the multicast FIB entry.
HostName	Host name of the node where the multicast FIB entry originated.

Displaying LSP summary information

Use the following procedure to display link-state packet (LSP) summary information. Link State Packets (LSP) contain information about the state of adjacencies or defined and distributed static routes. Intermediate System to Intermediate System (IS-IS) exchanges this information with neighboring IS-IS routers at periodic intervals.

Procedure

1. From the navigation tree, choose **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **LSP Summary** tab.

LSP Summary field descriptions

Use the data in the following table to use the **LSP Summary** tab.

Name	Description
Level	Specifies the level at which this LSP appears.
ID	Specifies the 8 byte LSP ID, consisting of the SystemID, Circuit ID, and Fragment Number.
Seq	Specifies the sequence number for this LSP.
Checksum	Specifies the 16 bit Fletcher Checksum for this LSP.
LifetimeRemain	The remaining lifetime in seconds for this LSP.
HostName	The hostname listed in LSP, or the system name if host name is not configured.

Displaying IS-IS adjacencies

Use the following procedure to display IS-IS adjacency information. The platform sends IS-IS Hello (IIH) packets to discover IS-IS neighbors and establish and maintain IS-IS adjacencies. The platform continues to send IIH packets to maintain the established adjacencies. For two nodes to form an adjacency the B-VLAN pairs for the primary B-VLAN and secondary B-VLAN must match.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Adjacency** tab.

Adjacency field descriptions

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

Name	Description
Index	A unique value identifying the IS adjacency from all other such adjacencies on this circuit. This value is automatically assigned by the system when the adjacency is created.
AdjIfIndex	Specifies the IS-IS interface on which the adjacency is found.
Usage	Specifies how the adjacency is used. On a point-to-point link, this can be level 1 and 2. But on a LAN, the usage is level 1 on the adjacency between peers at level 1, and level 2 for the adjacency between peers at level 2.
State	Specifies the state of the adjacency: • down • initializing • up • failed
LastUpTime	Indicates when the adjacency most recently entered the state **up**, measured in hundredths of a second since the last re-initialization of the network management subsystem. Displays 0 if the adjacency has never been in state **up**.
NeighPriority	Specifies the priority of the neighboring Intermediate System for becoming the Designated Intermediate System.
HoldTimer	Specifies the holding time in seconds for this adjacency. This value is based on received IS-IS Hello (IIH) PDUs and the elapsed time since receipt.
NeighSysID	Specifies the system ID of the neighboring Intermediate System.
AdjHostName	Specifies the host name listed in the LSP, or the system name if host name is not configured.

Configuring IS-IS global parameters

Use the following procedure to configure IS-IS global parameters. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. In the navigation tree, expand the following folders:**Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. From the **Globals** tab, configure the global IS-IS parameters.
4. Click **Apply**.

Globals field descriptions

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

Name	Description
AdminState	Specifies the global status of IS-IS on the switch: on or off. The default is off.
LevelType	Sets the router type globally: • level1: Level-1 router type • level1and2: Level–1/2 router type is not supported in this release. The default value is level1.
SystemId	Specifies the IS-IS system ID for the switch. Valid value is a 6–byte value in the format <xxxx.xxxx.xxxx>. **Important:** After you have configured the SPBM nickname and enabled IS-IS, if you require a change of the system ID, you must also change the nickname. However, for naming convention purposes or configuration purposes, you may not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps: 1. Disable IS-IS. 2. Change the system ID. 3. Change the nickname to a temporary one. 4. Enable IS-IS. 5. Disable IS-IS. 6. Change the nickname to the original nickname. 7. Enable IS-IS.
MaxLspGenInt	Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate system. The value must be greater than any value configured for RxmtLspInt. The default value is 900 seconds.
CsnpInt	Specifies the Complete Sequence Number Packet (CSNP) interval in seconds. This is a system level parameter that applies for L1 CSNP generation on all interfaces. The default value is 10.
RxmtLspInt	Specifies the minimum time between retransmission of an LSP. This defines how fast the switch resends the same LSP. This is a system level parameter that applies for L1 retransmission of LSPs.

Name	Description
	The default value is 5 seconds.
PSNPInterval	Specifies the Partial Sequence Number Packet (PSNP) interval in seconds. This is a system level parameter that applies for L1 PSNP generation on all interfaces. The default value is 2.
SpfDelay	Specifies the SPF delay in milliseconds. This value is used to pace successive SPF runs. The timer prevents two SPF runs from being scheduled very closely. The default value is 100 milliseconds.
HostName	Specifies a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system name comes from the host name configured at the system level.
IpSourceAddress	Specifies IP source address for SPBM IP shortcuts.

Configuring system-level IS-IS parameters

Use the following procedure to configure system-level IS-IS parameters.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **System Level** tab.
3. Configure the IS-IS system level parameters.
4. Click **Apply**.

System Level field descriptions

Use the data in the following table to use the **System Level** tab.

Name	Description
Index	Specifies the level: l1 or l2. In this release, only l1 is supported.
State	Specifies the state of the database at this level. The value 'off' indicates that IS-IS is not active at this level. The value 'on' indicates that IS-IS is active at this level, and not overloaded. The value 'waiting' indicates a database that is low on an essential resources, such as memory. The administrator may force the state to 'overloaded' by setting the object **SetOverload**. If the state is 'waiting' or 'overloaded', you originate LSPs with the Overload bit set.

Name	Description
SetOverload	Sets or clears the overload condition. The possible values are true or false. The default value is false.
SetOverloadUntil	Sets the IS-IS overload-on-startup value in seconds. The overload-on-startup value is used as a timer to control when to send out LSPs with the overload bit cleared after IS-IS startup. **Note:** If you configure **SetOverloadUntil** to a number other than zero, then the overload bit is set at this level when the **AdminState** variable goes to the state 'on' for this Intermediate System. After the **SetOverloadUntil** seconds elapse, the overload flag remains set if the implementation runs out of memory or if you configured it manually using **SetOverload** to true. If **SetOverload** is false, the system clears the overload bit after **SetOverloadUntil** seconds elapse, if the system has not run out of memory. The default value is 20.
MetricStyle	Specifies the IS-IS metric type. Available values are narrow, wide or both. Only wide is supported in this release.

Displaying IS-IS system statistics

Use the following procedure to display Intermediate-System-to-Intermediate-System (IS-IS) system statistics.

Procedure

1. In the navigation pane, choose **Configuration** > **IS-IS**.
2. Click **Stats**.
3. Click the **System Stats** tab.

System Stats field descriptions

Use the data in the following table to use the **System Stats** tab.

Name	Description
CorrLSPs	Indicates the number of corrupted in-memory link-state packets (LSPs) detected. LSPs received from the wire with a bad checksum are silently dropped and not counted.
AuthFails	Indicates the number of authentication key failures recognized by this Intermediate System.

Name	Description
LSPDbaseOloads	Indicates the number of times the LSP database has become overloaded.
ManAddrDropFromAreas	Indicates the number of times a manual address has been dropped from the area.
AttmptToExMaxSeqNums	Indicates the number of times the IS has attempted to exceed the maximum sequence number.
SeqNumSkips	Indicates the number of times a sequence number skip has occurred.
OwnLSPPurges	Indicates the number of times a zero-aged copy of the system's own LSP is received from some other node.
IDFieldLenMismatches	Indicates the number of times a PDU is received with a different value for ID field length to that of the receiving system.
PartChanges	Indicates partition changes.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you expanded the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.

Configuring IS-IS interfaces

Use the following procedure to configure IS-IS interfaces. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Configure the IS-IS interface parameters.
5. Click **Apply**.

Interfaces field descriptions

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

Name	Description
Index	The identifier of this circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
IfIndex	Specifies the interface on which the circuit is configured (port or MLT).
Type	Specifies the IS-IS circuit type. In this release, only the point-to-point (PtToPt) interface type is supported.
AdminState	Specifies the administrative state of the circuit: on or off.
OperState	Specifies the operational state of the circuit.
AuthType	Specifies the authentication type: • none • simple: If selected, you must also specify a key value but the key id is optional. Simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5: hmac-md5: If selected, you must also specify a key value but the key-id is optional. MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. There is an optional key ID. The default is none.
AuthKey	Specifies the authentication key.
KeyId	Specifies the authentication key ID.
LevelType	Specifies the router type globally: • level1: Level-1 router type • level 1and2: Level–1/2 router type. Not supported in this release. The default value is level1.
NumAdj	Specifies the number of adjacencies on this circuit.
NumUpAdj	Specifies the number of adjacencies that are up.

Configuring IS-IS interface level parameters

Use the following procedure to configure IS-IS interface level parameters. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. From the navigation tree, choose **Configuration** > **IS-IS**.

2. Click **IS-IS**.
3. Click the **Interfaces Level** tab.
4. Configure the IS-IS interface level parameters.
5. Click **Apply**.

Interfaces field descriptions

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

Name	Description
Index	The identifier of this circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
Level	Specifies the router type globally: • l1: Level1 router type • l12: Level1/Level2 router type. Not supported in this release. The default value is l1.
ISPriority	Specifies an integer sub-range for IS-IS priority. The default is 64.
HelloTimer	Configures the level 1 hello interval. Specifies the maximum period, in milliseconds, between IS-IS Hello Packets (IIH) PDUs on multiaccess networks at this level for LANs. The value at Level1 is used as the period between Hellos on Level1/Level2 point to point circuits. Setting this value at Level 2 on an Level1/Level2 point-to-point circuit results in an error of InconsistentValue. The default value is 9000 milliseconds or 9 seconds.
HelloMultiplier	Configures the level 1 hello multiplier. The default value is 3 seconds.
DRHelloTimer	Indicates the period, in milliseconds, between Hello PDUs on multiaccess networks when this Intermediate System is the Designated Intermediate System. The default is 3.

Displaying IS-IS interface counters

Use the following procedure to display IS-IS interface counters.

Procedure

1. From the navigation pane, choose **Configuration** > **IS-IS**.
2. Click **Stats**.
3. Click the **Interface Counters** tab.

Interface Counters field descriptions

Use the data in the following table to use the **Interface Counters** tab.

Name	Description
Index	Shows a unique value identifying the IS-IS interface.
Type	Shows the type of interface.
AdjChanges	Shows the number of times an adjacency state change has occurred on this circuit.
InitFails	Shows the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures. Failures to form an adjacency are counted by isisCircRejAdjs.
RejAdjs	Shows the number of times an adjacency has been rejected on this circuit.
IDFieldLenMismatches	Shows the number of times an IS-IS control PDU with an ID field length different to that for this system has been received.
MaxAreaAddrMismatches	Shows the number of times an IS-IS control PDU with a max area address field different to that for this system has been received.
AuthFails	Shows the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	Shows the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.

Displaying IS-IS interface control packets

Use the following procedure to display IS-IS interface control packets.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IS-IS**.
2. Click **Stats**.
3. Click the **Interface Control Packets** tab.

Interface Control Packets field descriptions

Use the data in the following table to use the **Interface Control Packets** tab.

Name	Description
Index	Shows a unique value identifying the Intermediate-System-to-Intermediate-System (IS-IS) interface.
Level	Shows the level at which the system collects the counts.
Direction	Indicates whether the switch is sending or receiving the PDUs.

Name	Description
Hello	Indicates the number of IS-IS Hello frames seen in this direction at this level.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packets (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNP) frames seen in this direction at this level.

Graphing IS-IS interface counters

Use the following procedure to graph IS-IS interface counters.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.

Interface Counters field descriptions

The following table describes the fields in the **Interface Counters** tab.

Name	Description
InitFails	Indicates the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures.
RejAdjs	Indicates the number of times an adjacency has been rejected on this circuit.
IDFieldLenMismatches	Indicates the number of times an Intermediate-System-to-Intermediate-System (IS-IS) control PDU with an ID field length different from that for this system has been received.
MaxAreaAddrMismatches	Indicates the number of times an IS-IS control PDU with a max area address field different from that for this system has been received.
AuthFails	Indicates the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	Indicates the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.
AbsoluteValue	Displays the counter value.

Name	Description
Cumulative	Displays the total value since you expanded the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Graphing IS-IS interface sending control packet statistics

Use the following procedure to graph IS-IS interface receiving control packet statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.
6. Click the **Interface Sending Control Packets** tab.

Interface Sending Control Packets field descriptions

The following table describes the fields in the **Interface Sending Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello (IIH) PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you expanded the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Graphing IS-IS interface receiving control packet statistics

Use the following procedure to graph IS-IS interface sending control packet statistics.

Procedure

1. From the navigation pane, choose **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces** tab.
4. Select an existing interface.
5. Click the **Graph** button.
6. Click the **Interface Receiving Control Packets** tab.

Interface Receiving Control Packets field descriptions

The following table describes the fields in the **Interface Receiving Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
LSP	Indicates the number of IS-IS link-state packet (LSP) frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you expanded the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Displaying SPBM packet drop statistics by port

Use this procedure to display the SPBM drop statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IS-IS**.
2. Click **SPBM**.

3. Click the **Drop Stats** tab.

Drop Stats field descriptions

Use the data in the following table to use the **Drop Stats** tab.

Name	Description
PortIndex	Shows the slot/port number that identifies the ingress port of the dropped packet.
VlanId	Shows the VLAN ID.
VlanType	Shows the VLAN type as either primary or secondary.
UknownUcastSrcAddr	Shows the total number unknown source unicast packets.
RpfcUcastSrcAddr	Shows the total number of RPFC source unicast packets.
UnknownUcastDestAddr	Shows the total number of unknown destination unicast packets.
Unknown McastDesAddr	Shows the total number of unknown multicast destination packets.
RpfcMcastSrcAddr	Shows the total number of RPFC multicast source packets.
LastDropMac	Shows the time of the last drop backbone MAC.
IsMacDestAddr	Shows the IS MAC destination address.
LastDropMacHostName	Shows the last drop MAC host name.

Configuring an IS-IS Manual Area

Use the following procedure to configure an IS-IS manual area.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IS-IS**.
2. Click **IS-IS**.
3. Click the **Manual Area** tab.
4. Click **Insert**.
5. Specify an Area Address in the **AreaAddr** field, and click **Insert**.

Manual Area field descriptions

Use the data in the following table to use the **Manual Area** tab.

Name	Description
AreaAddr	Specifies the IS-IS manual area. Valid value is 1-13 bytes in the format <xx.xxxx.xxxx...xxxx>. In this release, only one manual area is supported. Use the same manual area across the entire SPBM cloud. For IS-IS to operate, you must configure at least one manual area.

SPBM configuration examples

This section provides configuration examples to configure basic SPBM and IS-IS infrastructure.

Basic SPBM configuration example

The following figure shows a sample greenfield deployment for SPBM.

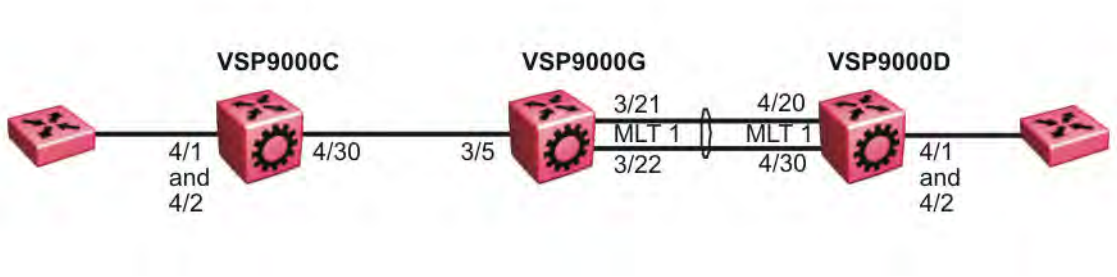

Figure 5: Greenfield SPBM deployment

Note the following:

- For migration purposes, SPBM can coexist with existing SMLT config

Ethernet and MLT configuration

The following sections show the steps required to configure the Ethernet and MLT interfaces in this example.

VSP9000C

```
PORT CONFIGURATION - PHASE 1

interface GigabitEthernet 4/30
encapsulation dot1q
exit
```

VSP9000G

```
PORT CONFIGURATION - PHASE 1

interface GigabitEthernet 3/5
encapsulation dot1q
exit

MLT CONFIGURATION

mlt 1 enable
mlt 1 member 3/21-3/22
mlt 1 encapsulation dot1q
```

VSP9000D

```
MLT CONFIGURATION

mlt 1 enable
mlt 1 member 4/20,4/30
mlt 1 encapsulation dot1q
```

IS-IS SPBM global configuration

The following figure shows the IS-IS area information added to the network.

Figure 6: IS-IS SPBM global

The following sections show the steps required to configure the global IS-IS SPBM parameters in this example.

VSP9000C

```
enable
configure terminal
prompt VSP9000C

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.13
spbm 1 b-vid 4000
```

```
VLAN CONFIGURATION

vlan create 4000 name "B-VLAN" type spbm-bvlan

ISIS CONFIGURATION

is-type l1
manual-area 30.0000
sys-name VSP9000C
exit
router isis enable
```

VSP9000G

```
enable
configure terminal
prompt VSP9000G

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.10
spbm 1 b-vid 4000

VLAN CONFIGURATION

vlan create 4000 name "B-VLAN" type spbm-bvlan

ISIS CONFIGURATION

is-type l1
manual-area 30.0000
sys-name VSP9000G
exit
router isis enable
```

VSP9000D

```
enable
configure terminal
prompt VSP9000D

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.14
spbm 1 b-vid 4000

VLAN CONFIGURATION

vlan create 4000 name "B-VLAN" type spbm-bvlan
```

```
ISIS CONFIGURATION

is-type l1
manual-area 30.0000
sys-name VSP9000D
exit
router isis enable
```

IS-IS SPBM Interface Configuration

The following figure shows the IS-IS area information and interfaces in the network.

Figure 7: IS-IS SPBM interface

The following sections show the steps required to configure the IS-IS SPBM interfaces in this example.

VSP9000C

```
PORT CONFIGURATION - PHASE II

interface GigabitEthernet 4/30
isis
isis spbm 1
isis enable
exit
```

VSP9000G

```
PORT CONFIGURATION - PHASE II

interface GigabitEthernet 3/5
isis
isis spbm 1
isis enable
exit

MLT INTERFACE CONFIGURATION

interface mlt 1
isis
isis spbm 1
isis enable
exit
```

VSP9000D

```
MLT INTERFACE CONFIGURATION
```

```
interface mlt 1
isis
isis spbm 1
isis enable
exit
```

IP multicast over SPBM global configuration

The following sections show the steps required to configure IP multicast over SPBM at a global level

VSP9000C

```
enable
configure terminal
prompt VSP9000C

ISIS SPBM CONFIGURATION
router isis
spbm 1 multicast enable
exit
```

VSP9000G

```
enable
configure terminal
prompt VSP9000G

ISIS SPBM CONFIGURATION
router isis
spbm 1 multicast enable
exit
```

VSP9000D

```
enable
configure terminal
prompt VSP9000D

ISIS SPBM CONFIGURATION
router isis
spbm 1 multicast enable
exit
```

Verifying SPBM operations

The following sections show the output from verifying the sample IS-IS SPBM configuration.

Checking operation — VSP9000C

```
VSP9000C:1# show isis interface
================================================================================
                                ISIS Interfaces
================================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
Port4/30   pt-pt   Level 1   UP        UP         1      1       10

VSP9000C:1# show isis adjacencies
================================================================================
                                ISIS Adjacencies
```

```
================================================================================
INTERFACE     L STATE   UPTIME      PRI   HOLDTIME SYSID            HOST-NAME
--------------------------------------------------------------------------------
Port4/30      1 UP      1d 19:11:30 127   26       000e.6225.a3df   VSP9000G
--------------------------------------------------------------------------------
 1 out of 1 interfaces have formed an adjacency
--------------------------------------------------------------------------------


VSP9000C:1# show isis spbm unicast-fib
================================================================================
                          SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION         BVLAN  SYSID          HOST-NAME   OUTGOING-INTERFACE  COST
ADDRESS
--------------------------------------------------------------------------------
00:0e:62:25:a3:df  4000   000e.6225.a3df  VSP9000G     4/30
00:14:0d:a0:13:df  4000   0014.0da0.13df  VSP9000D     4/30

--------------------------------------------------------------------------------
 Total number of SPBM UNICAST FIB entries 2
--------------------------------------------------------------------------------

VSP9000C:1# show isis spbm unicast-tree 4000
Node:000e.6225.a3df.00 (VSP9000G) -> ROOT
Node:0014.0da0.13df.00 (VSP9000D) -> Node:000e.6225.a3df.00 (VSP9000G) ->
ROOT
```

Checking operation — VSP9000G

```
VSP9000G:1# show isis interface
================================================================================
                                 ISIS Interfaces
================================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
Port3/5    pt-pt  Level 1   UP        UP         1      1        10
Mlt1       pt-pt  Level 1   UP        UP         1      1        10

VSP9000G:1# show isis adjacencies
================================================================================
                                 ISIS Adjacencies
================================================================================
INTERFACE L STATE   UPTIME      PRI HOLDTIME SYSID             HOST-NAME
--------------------------------------------------------------------------------
Port3/5   1 UP      1d 19:19:52 127 26       0015.e89f.e3df    VSP9000C
Mlt1      1 UP      04:57:34    127 20       0014.0da0.13df    VSP9000D

VSP9000G:1# show isis spbm unicast-fib
================================================================================
                          SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION ADDRESS  BVLAN  SYSID          HOST-NAME   OUTGOING-INTERFACE        COST
--------------------------------------------------------------------------------
00:14:0d:a0:13:df    4000   0014.0da0.13df VSP9000D           MLT-1
00:15:e8:9f:e3:df    4000   0015.e89f.e3df VSP9000C           3/5

VSP9000G:1# show isis spbm unicast-tree 4000
Node:0015.e89f.e3df.00 (VSP9000C) -> ROOT
Node:0014.0da0.13df.00 (VSP9000D) -> ROOT
```

Checking operation — VSP9000D

```
VSP9000D:1# show isis interface
================================================================================
```

```
                                ISIS Interfaces
================================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
Mlt1       pt-pt   Level 1   UP        UP         1      1       10

VSP9000D:1# show isis adjacencies
================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE L STATE      UPTIME   PRI HOLDTIME SYSID             HOST-NAME
--------------------------------------------------------------------------------
Mlt1      1 UP         05:03:59 127 21       000e.6225.a3df    VSP9000G

VSP9000D:1# show isis spbm unicast-fib
================================================================================
                          SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION ADDRESS  BVLAN  SYSID          HOST-NAME   OUTGOING-INTERFACE COST
--------------------------------------------------------------------------------
00:0e:62:25:a3:df    4000   000e.6225.a3df VSP9000G   MLT-1
00:15:e8:9f:e3:df    4000   0015.e89f.e3df VSP9000C   MLT-1

VSP9000D:1# show isis spbm unicast-tree 4000
Node:000e.6225.a3df.00 (VSP9000G) -> ROOT
Node:0015.e89f.e3df.00 (VSP9000C) -> Node:000e.6225.a3df.00 (VSP9000G) ->
ROOT
```

Chapter 4: SPBM and IS-IS services configuration

This chapter provides concepts and procedures to configure Layer 2 Virtual Services Networks (VSNs), IP Shortcuts, Layer 3 Virtual Services Networks (VSNs), and Inter-Virtual Services Networks (VSNs) routing.

Layer 2 VSN configuration

This section provides concepts and procedures to configure Layer 2 Virtual Services Networks (VSNs).

Layer 2 VSN configuration fundamentals

This section provides fundamentals concepts for Layer 2 VSN.

SPBM L2 VSN

SPBM supports Layer 2 VSN functionality where customer VLANs (C-VLANs) are bridged over the SPBM core infrastructure.

At the Backbone Edge Bridges (BEBs), customer VLANs (C-VLAN) are mapped to I-SIDs based on the local service provisioning. Outgoing frames are encapsulated in a MAC-in-MAC header, and then forwarded across the core to the far-end BEB, which strips off the encapsulation and forwards the frame to the destination network based on the I-SID-to-C-VLAN provisioning.

In the backbone VLAN (B-VLAN), Backbone Core Bridges (BCBs) forward the encapsulated traffic based on the BMAC-DA, using the shortest path topology learned using IS-IS.

The following figure shows a sample campus SPBM Layer 2 VSN network.

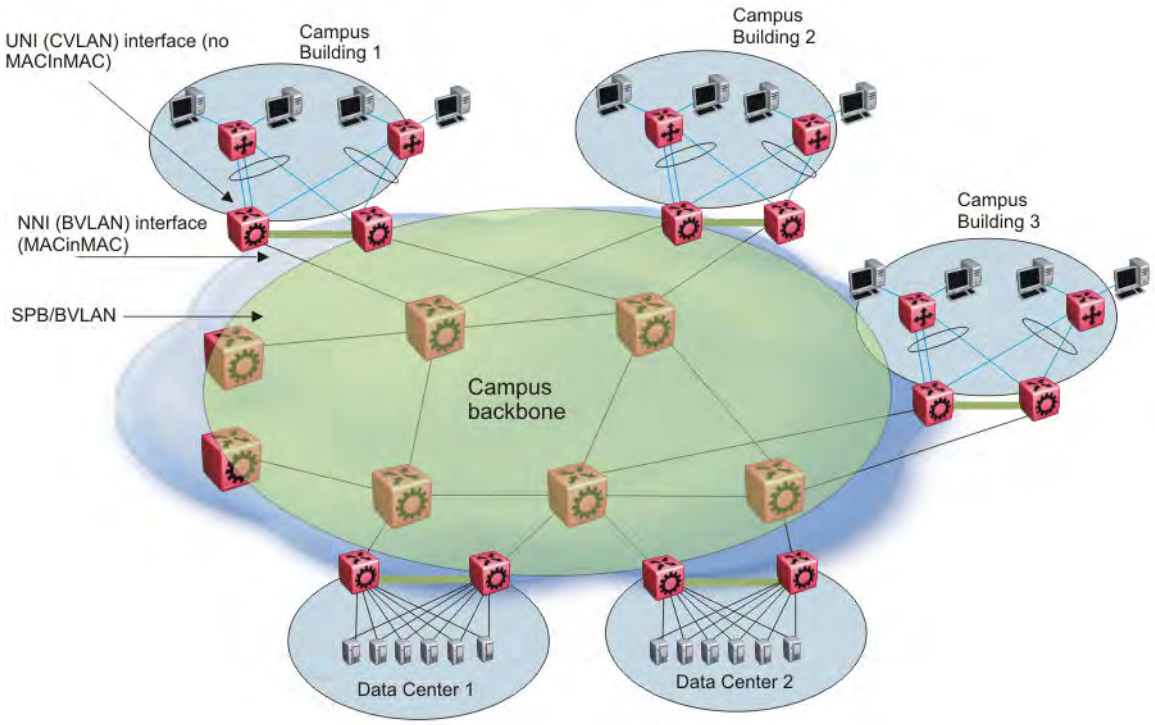

Figure 8: SPBM L2 VSN in a campus

One of the key advantages of the SPBM Layer 2 VSN is that network virtualization provisioning is achieved by configuring only the edge of the network (BEBs). As a result, the intrusive core provisioning that other Layer 2 virtualization technologies require is not needed when new connectivity services are added to the SPBM network. For example, when new virtual server instances are created and need their own VLAN instances, they are provisioned at the network edge only and do not need to be configured throughout the rest of the network infrastructure.

Based on its I-SID scalability, this solution can scale much higher than any 802.1Q tagging based solution. Also, due to the fact that there is no need for Spanning Tree in the core, this solution does not need any core link provisioning for normal operation.

Redundant connectivity between the C-VLAN domain and the SPBM infrastructure can be achieved by operating two SPBM switches in switch clustering (SMLT) mode. This allows the dual homing of any traditional link aggregation capable device into an SPBM network.

VSP 9000 configuration difference from ERS 8800

Important:

Different from Ethernet Routing Switch 8800 SPBM implementation, on Virtual Services Platform 9000 when you map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID), you must include the IST MLT peer switches in the VLAN, which means all IST member ports must be members of this C-VLAN. In ERS 8800, you cannot configure IST MLT to be part of the C-VLAN. For more information about this and other differences between the two products, see *Platform Migration Reference for Avaya Virtual Services Platform 9000,* NN46250-107.

If you add the I-SID to a VLAN with IST already configured, the following caution appears:

```
CAUTION: Adding I-SID to a VLAN
on an IST switch requires configuring this ISID-VLAN pair on both
IST peers and the IST MLT must be a member of the VLAN.
```

If you add IST to a switch, which has existing VLANs with I-SIDs, the following caution appears:

```
CAUTION: All VLANs with I-SIDs
MUST be configured on both IST peers and IST MLT MUST be a member
of all these VLANs.
```

SPBM sample operation—L2 VSN

The following section shows how a SPBM network is established, in this case, a Layer 2 VSN.

1. *Discover network topology*

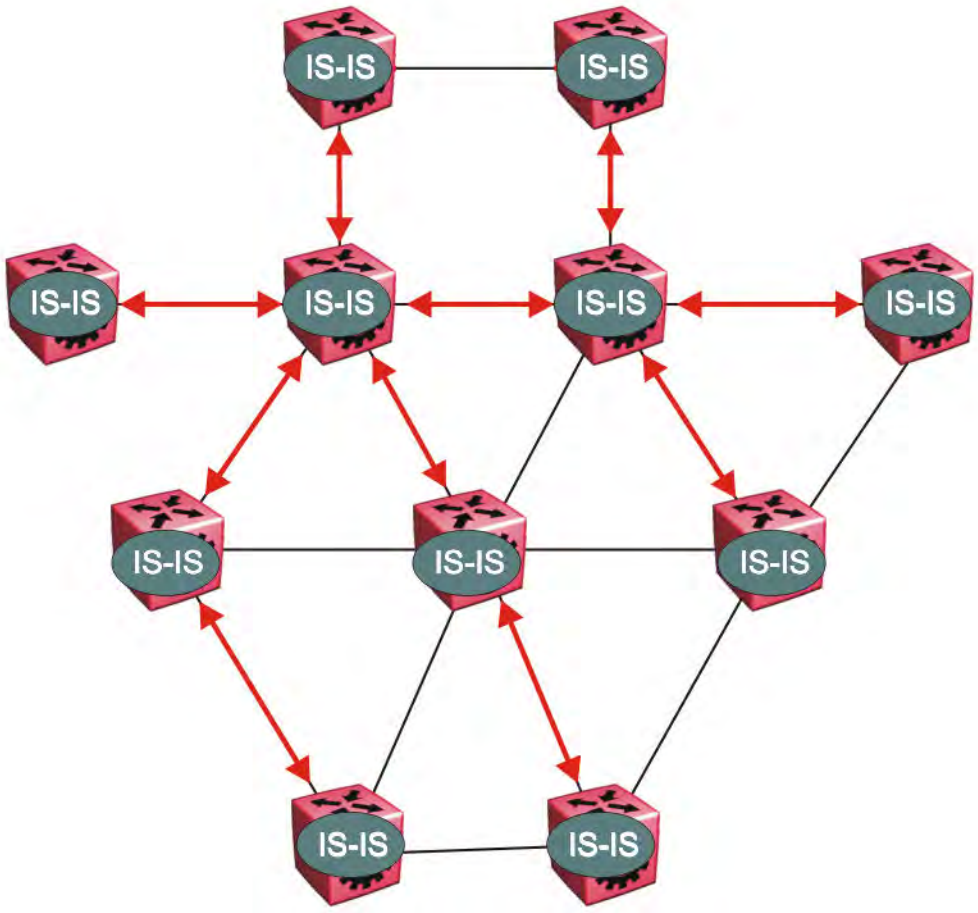

Figure 9: SPBM topology discover

IS-IS runs on all nodes of the SPBM domain. Since IS-IS is the basis of SPBM, the IS-IS adjacency must be formed first. After the neighboring nodes see hellos from each other they

look for the same Level (Level 1) and the same area (for example, Area 2f.8700.0000.00). After the hellos are confirmed both nodes send Link State Protocol Data Units, which contain connectivity information for the SPBM node. These nodes also send copies of all other LSPs they have in their databases. This establishes a network of connectivity providing the necessary information for each node to find the best and proper path to all destinations in the network.

Each node has a system ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network.

2. *Each IS-IS node automatically builds trees from itself to all other nodes*

   When the network topology is discovered and stored in the IS-IS link state database (LSDB), each node calculates shortest path trees for each source node. A unicast path now exists from every node to every other node

   With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes. Multicast FIB is not produced until Layer 2 VSN services are configured and learned.

3. *IS-IS advertises new service communities of interest*

   When a new service is provisioned, its membership is flooded throughout the topology with an IS-IS advertisement.

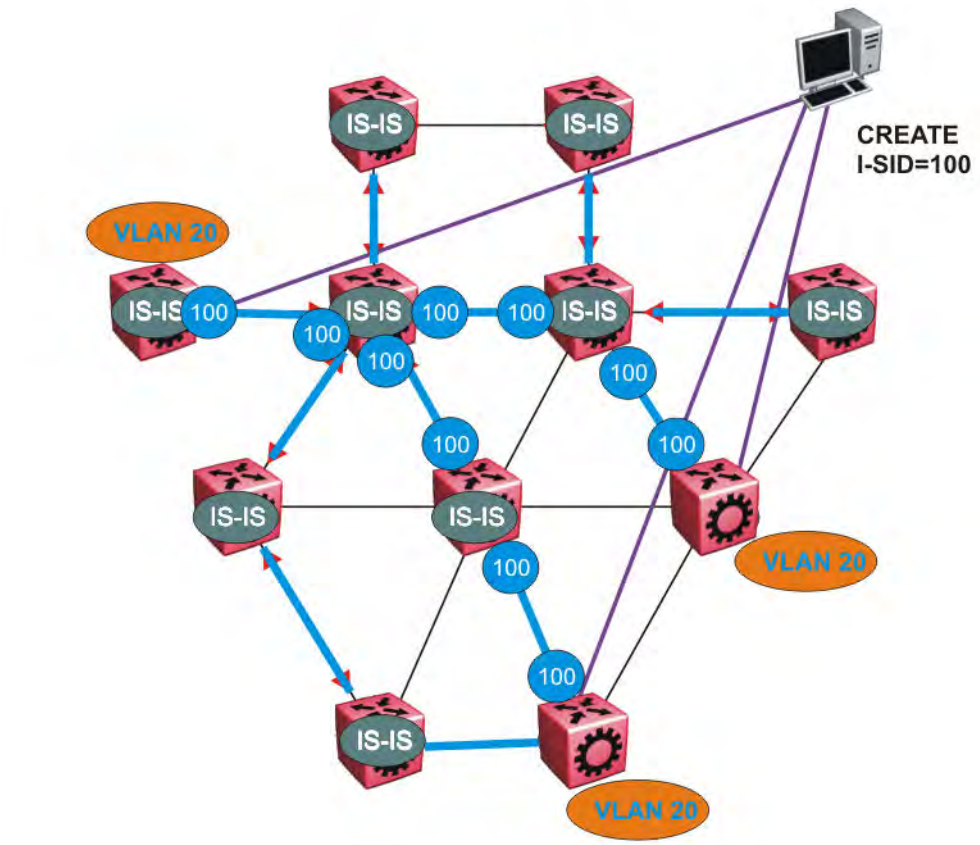

Figure 10: SPBM BMAC and I-SID population

BMAC and I-SID information is flooded throughout the network to announce new I-SID memberships. In this case, VLAN 20 is mapped to I-SID 100.

Note:

I-SIDs are only used for virtual services (Layer 2 VSNs and Layer 3 VSNs). If IP Shortcuts only is enabled on the BEBs, I-SIDs are never exchanged in the network as IP Shortcuts allow for IP networks to be transported across IS-IS.

Each node populates its FDB with the BMAC information derived from the IS-IS shortest path tree calculations. Thus there is no traditional flooding and learning mechanism in place for the B-VLAN, but FDBs are programmed by the IS-IS protocol.

4. *When a node receives notice of a new service AND is on the shortest path, it updates the FDB*

   In this scenario, where there are three source nodes having a membership on I-SID 100, there are three shortest path trees calculated (not counting the Equal Cost Trees (ECTs).

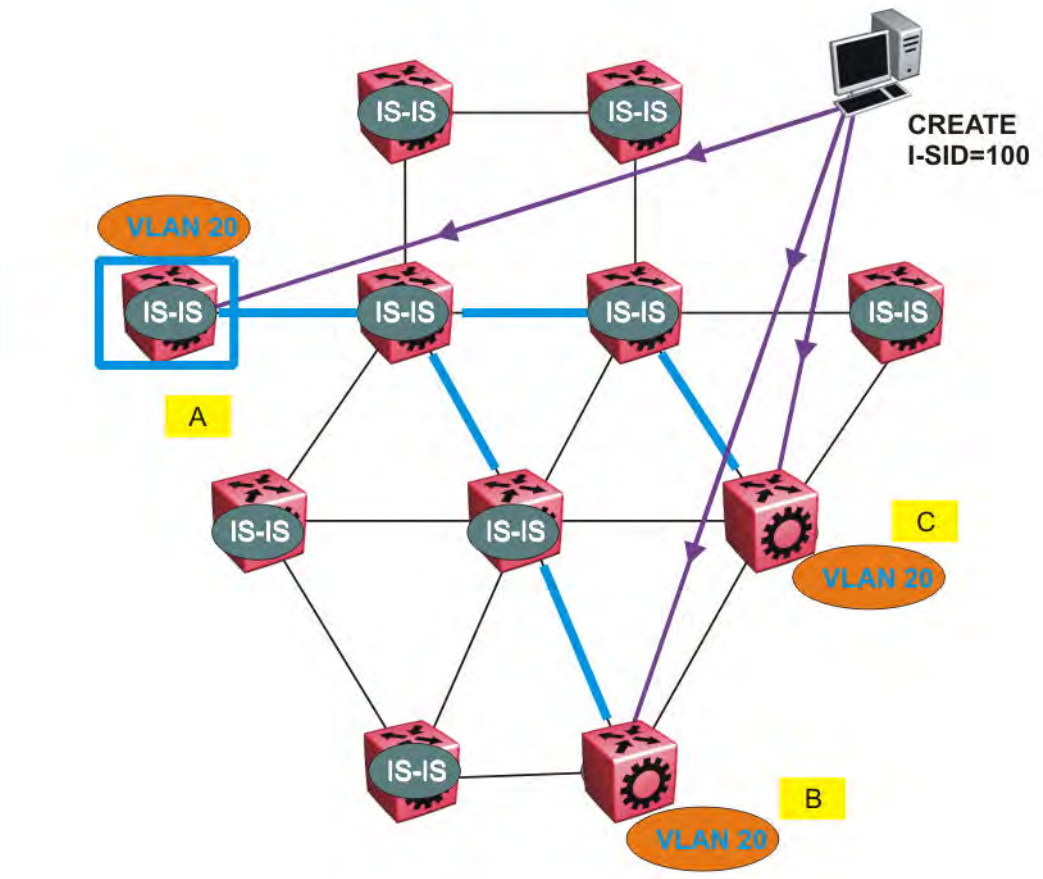

Figure 11: Shortest path tree for source node A

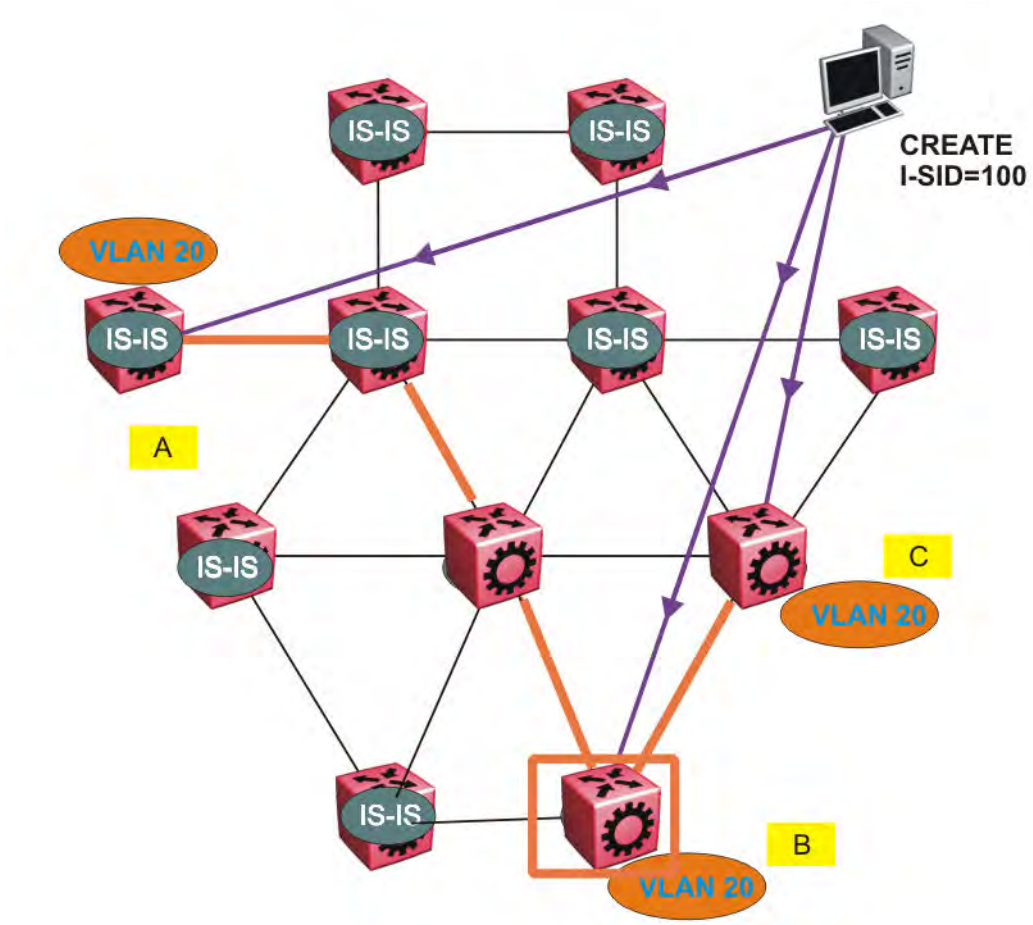

Figure 12: Shortest path tree for source node B

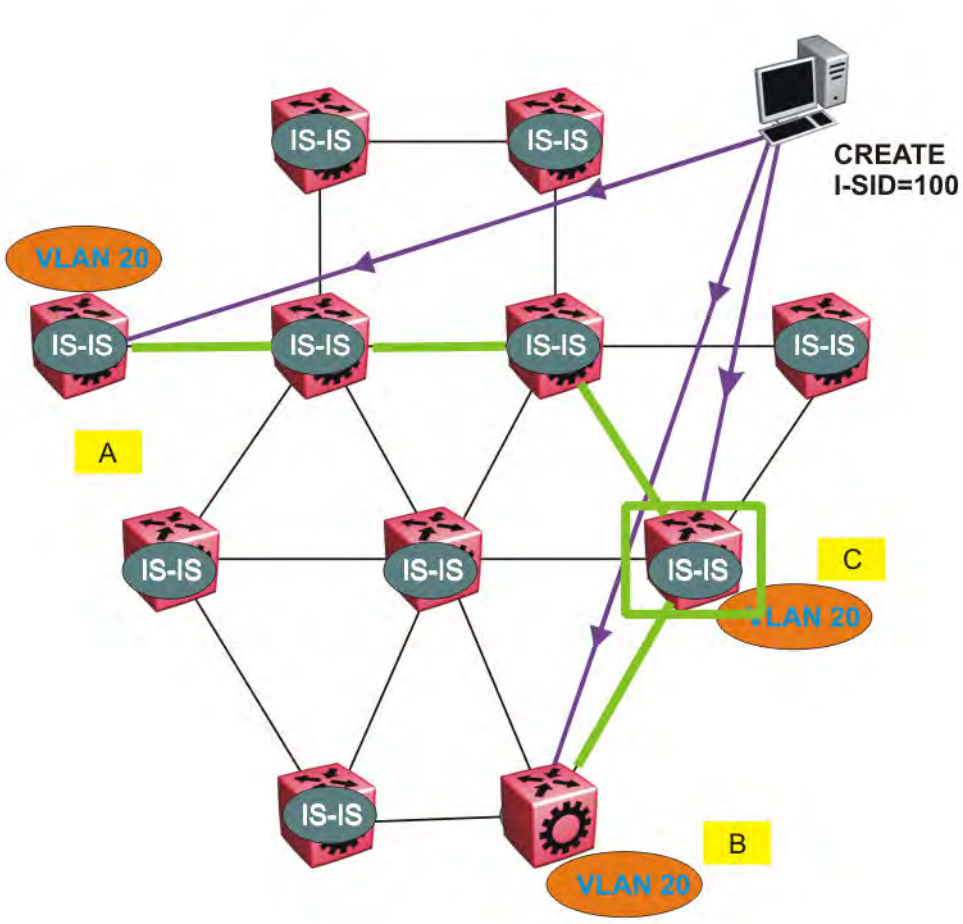

Figure 13: Shortest path tree for source node C

The paths between any two nodes are always the shortest paths. Also, the paths in either direction are congruent, thus a bidirectional communication stream can be monitored easily by mirroring ingress and egress on a link to a network analyzer.

VLAN traffic arriving on switch A and VLAN 20 is forwarded following the blue path, traffic arriving on switch B and VLAN 20 the orange path and on switch C VLAN 20 traffic is following the green path.

If the destination CMAC is unknown at the SPBM ingress node or the traffic is of type broadcast or multicast, then it is flooded to all members of the topology which spans VLAN 20. If the destination CMAC is already known, then the traffic is only forwarded as a unicast to the appropriate destination. In the SPBM domain, the traffic is switched on the BMAC header only. The bridge filtering database (FDB) at the VLAN to I-SID boundary (backbone edge bridge BEB), maintains a mapping between CMACs and corresponding BMACs.

For example, Switch B learns all CMACs which are on VLAN 20 connected to switch A with the BMAC of A in its FDB and the CMACs which are behind C are learned with the BMAC of C.

Layer 2 VSN IP multicast over SPBM

IP multicast over SPBM supports Layer 2 VSN functionality where multicast traffic is bridged over the SPBM core infrastructure. An application for Layer 2 VSNs using IP multicast over SPBM is multicast traffic in data centers.

After you configure **`ip igmp snooping`** on a VLAN that has an I-SID configured (a C-VLAN), that VLAN is automatically enabled for IP multicast over SPBM services. No explicit configuration exists separate from that to enable Layer 2 VSN IP multicast over SPBM.

Multicast traffic remains in the same Layer 2 VSN across the SPBM cloud for Layer 2 VSN IP multicast over SPBM. IP multicast over SPBM constrains all multicast streams within the scope level in which they originate. If a sender transmits a multicast stream to a BEB on a Layer 2 VSN with IP multicast over SPBM enabled, only receivers that are part of the same Layer 2 VSN can receive that stream.

I-SIDs

After a BEB receives IP multicast data from a sender, the BEB allocates a data service instance identifier (I-SID) in the range of 16,000,000 to 16,512,000 for the multicast stream. The stream is identified by the S, G, V tuple, which is the source IP address, the group IP address and the local VLAN the multicast stream is received on. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID.

In the context of Layer 2 VSNs with IP multicast over SPBM, the scope is the I-SID value of the Layer 2 VSN associated with the local VLAN on which the IP multicast data was received.

TLVs

This information is propagated through the SPBM cloud using IS-IS Link State Packets (LSPs), which carry TLV updates, that result in the multicast tree creation for that stream. For Layer 2 VSNs, the LSPs carry I-SID information and information about where IP multicast stream senders and receivers exist using TLV 144 and TLV 185.

IS-IS acts dynamically using the TLV information received from BEBs that connect to the sender and the receivers to create a multicast tree between them.

IGMP

After a BEB receives an IGMP join message from a receiver, a BEB queries the IS-IS database to check if a sender exists for the requested stream within the scope of the receiver. If the requested stream does not exist, the IGMP information is kept, but no further action is taken. If the request stream exists, the BEB sends an IS-IS TLV update to its neighbors to inform them of the presence of a receiver and this information is propagated through the SPBM cloud.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

Layer 2 VSN configuration using ACLI

This section provides procedures to configure Layer 2 VSNs using ACLI.

Configuring SPBM Layer 2 VSN

SPBM supports Layer 2 Virtual Service Network (VSN) functionality where customer VLANs (C-VLANs) are bridged over the SPBM core infrastructure.

At the BEBs, customer VLANs (C-VLAN) are mapped to I-SIDs based on the local service provisioning. Outgoing frames are encapsulated in a MAC-in-MAC header, and then forwarded across the core to the far-end BEB, which strips off the encapsulation and forwards the frame to the destination network based on the I-SID-to-C-VLAN provisioning.

Note:

Different from Ethernet Routing Switch 8800 SPBM implementation, on Virtual Services Platform 9000 when you map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID), you must include the IST MLT peer switches in the VLAN, which means all IST member ports must be members of this C-VLAN. In ERS 8800, you cannot configure IST MLT to be part of the C-VLAN. For more information about this and other differences between the two products, see *Platform Migration Reference for Avaya Virtual Services Platform 9000,* NN46250-107.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM BVLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology,

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID):

   ```
   vlan i-sid <1-4084> <0-16777215>
   ```

   Important:

   When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.

   Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

3. Display C-VLAN information:

   ```
   show vlan i-sid
   ```

Example

Map customer VLAN 5 to I-SID 5:.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#vlan i-sid 5 5
VSP-9012:1(config)#show vlan i-sid

================================================================================
                    Vlan I-SID
================================================================================
VLAN_ID    I-SID
--------------------------------------------------------------------------------
1
5                           5
50          200
51
52
53
54
55
56
57
58
59
60
61
62
63
65
17 out of 22 Total Num of Vlans displayed
```

Variable definitions

Use the data in the following table to use the **`vlan i-sid`** command.

Variable	Value
vlan i-sid *<1–4084>* *<0–16777215>*	Specifies the customer VLAN (CVLAN) to associate with the I-SID. Use the no or default options to remove the I-SID from the specified VLAN. **Note:** Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

Displaying C-VLAN I-SID information

Use the following procedure to display C-VLAN I-SID information.

Procedure

1. Log on to the switch to enter User EXEC mode.

2. Display the C-VLAN to I-SID associations:

   ```
   show vlan i-sid <1-4084>
   ```

3. Display the IS-IS SPBM multicast-FIB calculation results by I-SID:

   ```
   show isis spbm i-sid {all|config|discover} [vlan <0-4084>] [id <1-16777215>] [nick-name <x.xx.xx>]
   ```

4. Discover where entries are learned:

   ```
   show vlan mac-address-entry [spbm-tunnel-as-mac]
   ```

5. Display the VLAN remote MAC table for a C-VLAN:

   ```
   show vlan remote-mac-table <1-4084>
   ```

Example

```
VSP-9012:1#show vlan i-sid
================================================================================
                                  Vlan I-SID
================================================================================
VLAN_ID    I-SID
--------------------------------------------------------------------------------
1
2
5          5
10
20
```

```
VSP-9012:1#show isis spbm i-sid all
================================================================================
                                SPBM ISID INFO
================================================================================
ISID   SOURCE NAME    VLAN   SYSID                TYPE         HOST_NAME
--------------------------------------------------------------------------------
200    1.11.16        1000   0014.c7e1.33df       config   VSP-9000
300    1.11.16        1000   0014.c7e1.33df       config   VSP-9000
400    1.11.16        1000   0014.c7e1.33df       config   VSP-9000
200    1.11.16        2000   0014.c7e1.33df       config   VSP-9000
300    1.11.16        2000   0014.c7e1.33df       config   VSP-9000
400    1.11.16        2000   0014.c7e1.33df       config   VSP-9000
200    1.12.45        1000   0016.ca23.73df       discover VSP-9001
300    1.12.45        1000   0016.ca23.73df       discover VSP-9001

--------------------------------------------------------------------------------
 Total number of SPBM ISID entries configed: 6
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries discovered: 2
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries: 8
--------------------------------------------------------------------------------
```

```
VSP-9012:1#show vlan mac-address-entry
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                SMLT
ID   STATUS     ADDRESS            INTERFACE    REMOTE   TUNNEL
--------------------------------------------------------------------------------
1    learned    00:1d:42:6b:10:03  Port-3/9     false    VSP-9001
1    learned    00:80:2d:22:ac:46  Port-3/15    false    VSP-9001
2    self       00:01:81:08:92:00  -            false     —
```

```
2    learned    00:18:b0:55:52:00  smlt-ist     true     VSP-9001
5    learned    00:00:00:00:00:1a  access       false    VSP-9001
5    learned    00:00:00:00:00:62  VSP-9006     true     VSP-9001
10   self       00:00:00:00:49:50  Port-3/9     false     —
10   self       00:00:00:50:00:50  Port-3/9     false     —
```

```
VSP-9012:1#show vlan remote-mac-table 100
================================================================================================
====
                                    Vlan Remote Mac
Table
================================================================================================
====
VLAN STATUS  MAC-ADDRESS        DEST-MAC           BVLAN DEST-SYSNAME      PORTS
SMLTREMOTE
------------------------------------------------------------------------------------------------
----
100  learned 00:15:40:af:d2:00  00:74:00:00:00:00  20    VSP-9005          MLT-2
true

------------------------------------------------------------------------------------------------
----
 Total number of VLAN Remote MAC entries 1
```

Variable definitions

Use the data in the following table to use the **`show vlan`** commands.

Variable	Value
i-sid <*0–4084*>	Displays I-SID information for the specified C-VLAN.
mac-address-entry [spbm-tunnel-as-mac]	Displays the bridging forwarding database. Use the optional parameter, spbm-tunnel-as-mac to display the BMAC in the `TUNNEL` column. If you do not use this optional parameter, the `TUNNEL` column displays the host name. If an entry is not learned in the SPBM network, the TUNNEL column will be empty (–).
remote-mac-table <*1–4084*>	Displays C-VLAN remote-mac-table information.

Use the data in the following table to use the **`show isis`** commands.

Variable	Value
spbm i-sid {all|config|discover}	• all: displays all I-SID entries • config: displays configured I-SID entries • discover: displayes discovered I-SID entries
vlan <*0–4084*>	Displays I-SID information for the specified SPBM VLAN.
id <*1–16777215*>	Displays I-SID information for the specified I-SID.
nick-name <*x.xx.xx*>	Displays I-SID information for the specified nickname.

Job aid

The following sections describe the fields in the outputs for the C-VLAN I-SID show commands.

show vlan i-sid

The following table describes the fields in the output for the **show vlan i-sid** command.

Parameter	Description
VLAN_ID	Indicates the VLAN IDs.
I-SID	Indicates the I-SIDs associated with the specified C-VLANs.

show isis spbm i-sid

The following describes the fields in the output for the **show isis spbm i-sid** command.

Parameter	Description
ISID	Indicates the IS-IS SPBM I-SID identifier.
SOURCE NAME	Indicates the nickname of the node where this I-SID was configured or discovered. **Note:** SOURCE NAME is equivalent to nickname.
VLAN	Indicates the B-VLAN where this I-SID was configured or discovered.
SYSID	Indicates the system identifier.
TYPE	Indicates the SPBM I-SID type as either configured or discovered.
HOST_NAME	Indicates the host name of the multicast FIB entry.

show vlan mac-address-entry

The following table describes the fields in the output for the **show vlan mac-address-entry** command.

Parameter	Description
VLAN ID	Indicates the VLAN for this MAC address.
STATUS	Indicates the status of this entry: • other • invalid • learned • self • mgmt
MAC ADDRESS	Indicates the MAC address.
INTERFACE	Displays the network-to-network (NNI) interface.
SMLT REMOTE	Indicates the MAC address entry for the remote IST peer.

Parameter	Description
TUNNEL	Indicates the host name of the remote Backbone Edge Bridge (BEB).

show vlan remote-mac-table

The following table describes the fields in the output for the **`show vlan remote-mac-table`** command.

Parameter	Description
VLAN	Indicates the VLAN ID for this MAC address.
STATUS	Indicates the status of this entry: • other • invalid • learned • self • mgmt
MAC-ADDRESS	Indicates the customer MAC address for which the bridge has forwarding and/or filtering information.
DEST-MAC	Indicates the provide MAC address for which the bridge has forwarding and/or filtering information.
BVLAN	Indicates the B-VLAN ID for this MAC address.
DEST-SYSNAME	Indicates the system name of the node where the MAC address entry comes from.
PORTS	Either displays the value 0 or indicates the port in which a frame comes from.
SMLTREMOTE	Inidicates the MAC address entry for the remote IST peer.

Configuring Layer 2 VSN IP multicast over SPBM

Use this procedure to configure IP multicast over SPBM for Layer 2 VSN functionality. With Layer 2 VSN IP multicast over SPBM, multicast traffic remains in the same Layer 2 VSN across the SPBM cloud.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must assign the same I-SID to the C-VLANs on all the BEBs where you configure the C-VLAN.
- You must enable SPBM multicast globally.

About this task

Traffic is only delivered to UNIs on the Layer 2 VSN where the switch receives IGMP joins and reports. Traffic does not cross the Layer 2 VSN boundary.

Configuring **ip igmp snooping** on a VLAN that has an I-SID configured (a C-VLAN) automatically enables that VLAN for SPBM multicast services. No explicit configuration exists separate from that to enable Layer 2 VSN multicast over SPBM.

SPBM supports enabling IGMP Snooping on a C-VLAN, but it does not support enabling Protocol Independent Multicast (PIM) on a C-VLAN. If you enable IGMP snooping on a C-VLAN, then its operating mode is Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network.

In this release, the switch only supports IPv4 multicast traffic.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable proxy snoop:

   ```
   ip igmp proxy
   ```

3. Enable IGMP snooping:

   ```
   ip igmp snooping
   ```

4. **(Optional)** If you want to configure an address for the IGMP queries, enter the following command:

   ```
   ip igmp snoop-querier-addr <A.B.C.D>
   ```

   This step is not always required. The IGMP Querier on the BEB uses a source address 0.0.0.0 by default. When you do not configure this, a BEB sends IGMP queries on the UNI ports with 0.0.0.0 as the source IP address. Some Layer 2 edge switches do not support a 0.0.0.0 querier. You can use a fictitious IP address as the querier address, and use the same address on all BEBs in the network.

5. **(Optional)** Enable IGMPv3 at a VLAN level by enabling SSM-snooping and IGMPv3:

   ```
   ip igmp ssm-snoop
   ip igmp version 3
   ```

   You must enable SSM snoop before you configure IGMP version 3 and both ssm-snoop and snooping must be enabled for IGMPv3.

   For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

Example

Enable IGMPv2 at a VLAN level:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config-if)#interface vlan 501
Switch:1(config-if)#ip igmp proxy
Switch:1(config-if)#ip igmp snooping
Switch:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
```

Enable IGMPv3 at a VLAN level:

```
Switch:>enable
Switch:#configure terminal
Switch:1(config)#interface vlan 2256
Switch:1(config-if)#ip igmp proxy
Switch:1(config-if)#ip igmp snooping
Switch:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
Switch:1(config-if)#ip igmp version 3
Switch:1(config-if)#ip igmp ssm-snoop
```

Viewing Layer 2 VSN IP multicast over SPBM information

Use the following options to display Layer 2 VSN information to confirm proper configuration.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display all IP multicast over SPBM route information:

   `show isis spbm ip-multicast-route [all]`
3. Display detailed IP multicast over SPBM route information:

   `show isis spbm ip-multicast-route [detail]`
4. Display IP multicast route information by VLAN:

   `show isis spbm ip-multicast-route [vlan <1-4084>]`
5. Display IP multicast over SPBM route information by VSN I-SID:

   `show isis spbm ip-multicast-route [vsn-isid <1-16777215>]`
6. Display IP multicast over SPBM route information by group address:

   `show isis spbm ip-multicast-route [group {A.B.C.D}]`
7. Display IP multicast over SPBM route information by source address:

   `show isis spbm ip-multicast-route [source {A.B.C.D}]`
8. Display summary information for each S, G, V tuple with the corresponding scope, data I-SID, and the host name of the source:

   `show isis spb-mcast-summary [host-name WORD<0-255>][lspid <xxxx.xxxx.xxxx.xx-xx>]`

Example

Note:

The slot/port numbers in the following example are applicable to the VSP 9000. The slot/port numbers available for your platform may be different.

```
Switch:1#show isis spbm ip-multicast-route all
================================================================================
       SPBM IP-MULTICAST ROUTE INFO ALL
================================================================================
Type   VrfName   Vlan Source      Group       VSN-ISID  Data ISID  BVLAN Source-BEB
                 Id
--------------------------------------------------------------------------------
snoop   GRT      501  192.0.2.1   233.252.0.1   5010     16300001   10    el2
snoop   GRT      501  192.0.2.1   233.252.0.2   5010     16300002   20    el2
snoop   GRT      501  192.0.2.1   233.252.0.3   5010     16300003   10    el2
snoop   GRT      501  192.0.2.1   233.252.0.4   5010     16300004   20    el2
snoop   GRT      501  192.0.2.1   233.252.0.5   5010     16300005   10    el2
snoop   GRT      501  192.0.2.1   233.252.0.6   5010     16300006   20    el2
snoop   GRT      501  192.0.2.1   233.252.0.7   5010     16300007   10    el2
snoop   GRT      501  192.0.2.1   233.252.0.8   5010     16300008   20    el2
snoop   GRT      501  192.0.2.1   233.252.0.9   5010     16300009   10    el2
snoop   GRT      501  192.0.2.1   233.252.0.10  5010     16300010   20    el2

--------------------------------------------------------------------------------
 Total Number of SPBM IP MULTICAST ROUTE Entries: 10
--------------------------------------------------------------------------------


Switch:1#show isis spbm ip-multicast-route vlan 501
================================================================================
       SPBM IP-MULTICAST ROUTE INFO ALL
================================================================================
Type   VrfName  Vlan  Source      Group       VSN-ISID  Data ISID BVLAN Source-BEB
                 Id
--------------------------------------------------------------------------------
snoop   GRT     501  192.0.2.1   233.252.0.1   5010    16300001   10    el2
snoop   GRT     501  192.0.2.1   233.252.0.2   5010    16300002   20    el2
snoop   GRT     501  192.0.2.1   233.252.0.3   5010    16300003   10    el2
snoop   GRT     501  192.0.2.1   233.252.0.4   5010    16300004   20    el2
snoop   GRT     501  192.0.2.1   233.252.0.5   5010    16300005   10    el2
snoop   GRT     501  192.0.2.1   233.252.0.6   5010    16300006   20    el2
snoop   GRT     501  192.0.2.1   233.252.0.7   5010    16300007   10    el2
snoop   GRT     501  192.0.2.1   233.252.0.8   5010    16300008   20    el2
snoop   GRT     501  192.0.2.1   233.252.0.9   5010    16300009   10    el2
snoop   GRT     501  192.0.2.1   233.252.0.10  5010    16300010   20    el2

--------------------------------------------------------------------------------
 Total Number of SPBM IP MULTICAST ROUTE Entries: 10
--------------------------------------------------------------------------------


Switch:1# show isis spbm ip-multicast-route vsn-isid 5010
================================================================================
       SPBM IP-MULTICAST ROUTE INFO - VLAN ID : 501, VSN-ISID : 5010
================================================================================
Source          Group           Data ISID  BVLAN Source-BEB
--------------------------------------------------------------------------------
192.0.2.1      233.252.0.1    16300001   10    el2
192.0.2.1      233.252.0.2    16300002   20    el2
192.0.2.1      233.252.0.3    16300003   10    el2
192.0.2.1      233.252.0.4    16300004   20    el2
192.0.2.1      233.252.0.5    16300005   10    el2
192.0.2.1      233.252.0.6    16300006   20    el2
192.0.2.1      233.252.0.7    16300007   10    el2
```

```
192.0.2.1      233.252.0.8    16300008   20    el2
192.0.2.1      233.252.0.9    16300009   10    el2
192.0.2.1      233.252.0.10   16300010   20    el2

--------------------------------------------------------------------------------
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
--------------------------------------------------------------------------------

Switch:1# show isis spbm ip-multicast-route vsn-isid 5010 detail
================================================================================
SPBM IP-MULTICAST ROUTE INFO - TYPE : SNOOP , VLAN ID : 501, VSN-ISID : 5010
================================================================================
Source        Group     Data ISID BVLAN  NNI Rcvrs     UNI Rcvrs Source-BEB
--------------------------------------------------------------------------------
192.0.2.1 233.252.0.1  16300001 10     4/3         V501:9/38 el2
192.0.2.1 233.252.0.2  16300002 20     4/2,4/3     V501:9/38 el2
192.0.2.1 233.252.0.3  16300003 10     4/3         V501:9/38 el2
192.0.2.1 233.252.0.4  16300004 20     4/2,4/3     V501:9/38 el2
192.0.2.1 233.252.0.5  16300005 10     4/3         V501:9/38 el2
192.0.2.1 233.252.0.6  16300006 20     4/2,4/3     V501:9/38 el2
192.0.2.1 233.252.0.7  16300007 10     4/3         V501:9/38 el2
192.0.2.1 233.252.0.8  16300008 20     4/2,4/3     V501:9/38 el2
192.0.2.1 233.252.0.9  16300009 10     4/3         V501:9/38 el2
192.0.2.1 233.252.0.10 16300010 20     4/2,4/3     V501:9/38 el2
--------------------------------------------------------------------------------
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
--------------------------------------------------------------------------------

Switch:1# show isis spb-mcast-summary

================================================================================
                SPB Multicast - Summary
================================================================================
SCOPE    SOURCE           GROUP            DATA            LSP  HOST
I-SID    ADDRESS          ADDRESS          I-SID     BVID  FRAG NAME
--------------------------------------------------------------------------------
5010     192.0.2.1       233.252.0.1     16300001  10    0x0  el2
5010     192.0.2.1       233.252.0.3     16300003  10    0x0  el2
5010     192.0.2.1       233.252.0.5     16300005  10    0x0  el2
5010     192.0.2.1       233.252.0.7     16300007  10    0x0  el2
5010     192.0.2.1       233.252.0.9     16300009  10    0x0  el2
5010     192.0.2.1       233.252.0.2     16300002  20    0x0  el2
5010     192.0.2.1       233.252.0.4     16300004  20    0x0  el2
5010     192.0.2.1       233.252.0.6     16300006  20    0x0  el2
5010     192.0.2.1       233.252.0.8     16300008  20    0x0  el2
5010     192.0.2.1       233.252.0.10    16300010  20    0x0  el2


Switch:1# show isis spbm ip-multicast-route vsn-isid 5010 detail
================================================================================
SPBM IP-MULTICAST ROUTE INFO - TYPE : SNOOP , VLAN ID : 501, VSN-ISID : 5010
================================================================================
Source         Group     Data ISID  BVLAN NNI Rcvrs     UNI Rcvrs     Source-BEB
--------------------------------------------------------------------------------
192.0.2.1    233.252.0.1  16300001 10    4/4,MLT-35     V501:9/32-9/33     el2
192.0.2.1    233.252.0.3  16300002 20      -          V501:9/32-9/33     el2
192.0.2.1    233.252.0.5  16300003 10      4/4,MLT-35     V501:9/32-9/33     el2
192.0.2.1    233.252.0.7  16300004 20      -          V501:9/32-9/33     el2
192.0.2.1    233.252.0.9  16300005 10      4/4,MLT-35     V501:9/32-9/33     el2
192.0.2.1    233.252.0.2  16300006 20      -          V501:9/32-9/33     el2
192.0.2.1    233.252.0.4  16300007 10      4/4,MLT-35     V501:9/32-9/33     el2
192.0.2.1    233.252.0.6  16300008 20      -          V501:9/32-9/33     el2
192.0.2.1    233.252.0.8  16300009 10      4/4,MLT-35     V501:9/32-9/33     el2
192.0.2.1    233.252.0.10 16300010 20      -          V501:9/32=9/33     el2
--------------------------------------------------------------------------------
```

```
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
-------------------------------------------------------------------------------
```

Variable definitions

Use the data in the following table to use the **show isis spbm ip-multicast-route** command.

Variable	Value
all	Displays all IP multicast over SPBM route information.
detail	Displays detailed IP multicast over SPBM route information.
group {A.B.C.D} source {A.B.C.D}	Displays information on the group IP address for the IP multicast over SPBM route. If you select source it will also display the source IP address.
vlan *<0–4084>*	Displays IP multicast over SPBM route information by VLAN.
vrf *WORD<1–16>*	Displays IP multicast over SPBM route information by VRF.
vsn-isid *<1–16777215>*	Displays IP multicast over SPBM route information by I-SID.

Use the data in the following table to use the **show isis spb-mcast-summary** command.

Variable	Value
host-name *WORD<0–255>*	Displays the SPBM multicast summary for a given host-name.
lspid *<xxxx.xxxx.xxxx.xx-xx>*	Displays the SPBM multicast summary for a given LSP ID.

Job aid

The following table describes the fields for the **show isis spbm ip-multicast-route all** command.

Parameter	Description
Type	Specifies the type of interface. The options include: • routed— For IP Shortcuts and Layer 3 VSNs. • snoop— For Layer 2 VSNs.
VrfName	Specifies the VRF name of the interface.
Vlan Id	Specifies the VLAN ID of the interface.
Source	Specifies the group IP address for the IP multicast over SPBM route.
Group	Specifies the group IP address for the IP multicast over SPBM route.

Parameter	Description
VSN-ISID	Specifies the VSN I-SID for Layer 2 VSNs and Layer 3 VSNs. Specifies the GRT for IP Shortcuts with IP multicast over SPBM because IP Shortcuts multicast over SPBM does not use a VSN I-SID.
Data ISID	Specifies the data I-SID for the IP multicast over SPBM route. After a BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVLAN	Specifies the B-VLAN for the IP multicast over SPBM route.
Source-BEB	Specifies the source BEB for the IP multicast over SPBM route.

The following table describes the fields for the **`show isis spbm ip-multicast-route vsn-isid`** command.

Parameter	Description
Source	Specifies the group IP address for the IP multicast over SPBM route.
Group	Specifies the group IP address for the IP multicast over SPBM route.
Data ISID	Specifies the data I-SID for the IP multicast over SPBM route. After a BEB receives the IP multicast over SPBM data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVLAN	Specifies the B-VLAN for the IP multicast over SPBM route.
Source-BEB	Specifies the source BEB for the IP multicast over SPBM route.

The following table describes the fields for the **`show isis spbm ip-multicast-route vsn-isid <1-16777215> detail`** command.

Parameter	Description
Source	Specifies the group IP address for the IP multicast route.

Parameter	Description
Group	Specifies the group IP address for the IP multicast route.
Data ISID	Specifies the data I-SID for the IP multicast route. After a BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVLAN	Specifies the B-VLAN for the IP multicast route.
NNI Rcvrs	Specifies the NNI receivers.
UNI Rcvrs	Specifies the UNI receivers.
Source-BEB	Specifies the source BEB for the IP multicast route.

The following table describes the fields for the **`show isis spb-mcast-summary`** command.

Parameter	Description
SCOPE I-SID	Specifies the scope I-SID. Layer 2 VSN and Layer 3 VSN each require a scope I-SID.
SOURCE ADDRESS	Specifies the group IP address for the IP multicast over SPBM route.
GROUP ADDRESS	Specifies the group IP address for the IP multicast over SPBM route.
DATA I-SID	Specifies the data I-SID for the IP multicast over SPBM route. After a BEB receives the IP multicast over SPBM data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVID	Specifies the Backbone VLAN ID associated with the SPBM instance.
LSP FRAG	Specifies the LSP fragment number.
HOST NAME	Specifies the host name listed in the LSP, or the system name if the host is not configured.

Viewing IGMP information for Layer 2 VSN multicast

Use the following commands to display IGMP information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display information about the interfaces where IGMP is enabled:

   ```
   show ip igmp interface [gigabitethernet {slot/port [-slot/port] [,...]}][vlan <1-4084>[vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

   Ensure that the output displays `snoop-spb` under MODE.

3. Display information about the IGMP cache:

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

4. Display information about the IGMP group:

   ```
   show ip igmp group [count][group {A.B.C.D}][member-subnet {A.B.C.D/X}] [vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

5. Display information about the IGMP sender:

   ```
   show ip igmp sender [count][group {A.B.C.D}][member-subnet {A.B.C.D/X}][vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

6. Display information about IGMP snoop-trace information:

   ```
   show ip igmp snoop-trace [group {A.B.C.D}][source {A.B.C.D}][vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

Example

Note:

The slot/port numbers in the following example are applicable to the VSP 9000. The slot/port numbers available for your platform may be different.

```
Switch:#enable
Switch:1#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    snoop-spb

1 out of 1 entries displayed

Switch:1(config)#show ip igmp interface vlan 1

================================================================================
                                  Vlan Ip Igmp
================================================================================
VLAN QUERY QUERY ROBUST VERSION LAST  PROXY  SNOOP  SSM    FAST   FAST
ID   INTVL MAX                  MEMB  SNOOP  ENABLE SNOOP  LEAVE  LEAVE
           RESP                 QUERY ENABLE        ENABLE ENABLE PORTS
--------------------------------------------------------------------------------
1    125   100   2      2       10    false  false  false  false

VLAN SNOOP   SNOOP           DYNAMIC   COMPATIBILITY EXPLICIT
ID   QUERIER QUERIER         DOWNGRADE MODE          HOST
     ENABLE  ADDRESS         VERSION                 TRACKING
```

```
--------------------------------------------------------------------------------
1    false   0.0.0.0         enable    disable       disable

Switch:1# show ip igmp sender

================================================================================
                IGMP Sender - GlobalRouter
================================================================================
                                  PORT/
GRPADDR         IFINDEX    MEMBER          MLT        STATE
--------------------------------------------------------------------------------
233.252.0.1    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.2    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.3    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.4    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.5    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.6    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.7    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.8    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.9    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.10   Vlan 501   192.2.0.1      9/5        NOTFILTERED

10 out of 10 entries displayed
```

```
Switch:1# show ip igmp group

================================================================================
                           IGMP Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
233.252.0.1    V501-9/16       192.2.0.1      204        Dynamic
233.252.0.2    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.3    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.4    V501-9/16       192.2.0.1      207        Dynamic
233.252.0.5    V501-9/16       192.2.0.1      204        Dynamic
233.252.0.6    V501-9/16       192.2.0.1      209        Dynamic
233.252.0.7    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.8    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.9    V501-9/16       192.2.0.1      211        Dynamic
233.252.0.10   V501-9/16       192.2.0.1      207        Dynamic

10 out of 10 group Receivers displayed

Total number of unique groups 10


Switch:1# show ip igmp sender

================================================================================
                           IGMP Sender - GlobalRouter
================================================================================
                                             PORT/
GRPADDR         IFINDEX    MEMBER          MLT        STATE
--------------------------------------------------------------------------------
233.252.0.1    Vlan 501   192.2.0.1      spb        NOTFILTERED
233.252.0.2    Vlan 501   192.2.0.1      spb        NOTFILTERED
233.252.0.3    Vlan 501   192.2.0.1      spb        NOTFILTERED
233.252.0.4    Vlan 501   192.2.0.1      spb        NOTFILTERED
233.252.0.5    Vlan 501   192.2.0.1      spb        NOTFILTERED
233.252.0.6    Vlan 501   192.2.0.1      spb        NOTFILTERED
233.252.0.7    Vlan 501   192.2.0.1      spb        NOTFILTERED
233.252.0.8    Vlan 501   192.2.0.1      spb        NOTFILTERED
```

```
233.252.0.9     Vlan 501    192.2.0.1       spb         NOTFILTERED
233.252.0.10    Vlan 501    192.2.0.1       spb         NOTFILTERED

10 out of 10 entries displayed

Switch:1# show ip igmp snoop-trace

Switch: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       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.2       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.3       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.4       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.5       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.6       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.7       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.8       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.9       192.2.0.1          501    spb         501    4/28,4/30   NETWORK
233.252.0.10      192.2.0.1          501    spb         501    4/28,4/30   NETWORK
```

Variable definitions

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

Variable	Value
gigabitethernet *{slot/port [-slot/port][,...]}*	Specifies the GigabitEthernet interface. Use <slot/port [-slot/port][,...]> to specify the slot and port.
vlan *<1-4084>*	Specifies the VLAN.
vrf *WORD<1–16>*	Specifies the VRF by name.
vrfids *WORD<0–512>*	Specifies the VRF by VRF ID.

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

Variable	Value
vrf *WORD<1–16>*	Specifies the VRF by name.
vrfids *WORD<0–512>*	Specifies the VRF by VRF ID.

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

Variable	Value
count	Specifies the number of entries.
group *{A.B.C.D}*	Specifies the group address.
member-subnet *{A.B.C.D/X}*	Specifies the IP address and network mask.
vrf *WORD<1–16>*	Displays the multicast route configuration for a particular VRF by name.

Variable	Value
vrfids *WORD<0–512>*	Displays the multicast route configuration for a particular VRF by VRF ID.

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

Variable	Value
count	Specifies the number of entries.
group *{A.B.C.D}*	Specifies the group address.
member-subnet *{A.B.C.D/X}*	Specifies the IP address and network mask.
vrf *WORD<1–16>*	Displays the multicast route configuration for a particular VRF by name.
vrfids *WORD<0–512>*	Displays the multicast route configuration for a particular VRF by VRF ID.

Use the data in the following table to use the **`show ip igmp snoop-trace`** command.

Variable	Value
group *{A.B.C.D}*	Specifies the group address.
source *{A.B.C.D}*	Specifies the source address.
vrf *WORD<1–16>*	Displays the multicast route configuration for a particular VRF by name.
vrfids *WORD<0–512>*	Displays the multicast route configuration for a particular VRF by VRF ID.

Job aid

The following table describes the fields for the **`show ip igmp interface`** command.

Parameter	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.

Parameter	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. If routed-spb displays in the MODE column, then IP multicast over SPBM is enabled on the Layer 3 VSN or for IP shortcuts. If snoop-spb displays in the MODE column, then IGMP is enabled on a VLAN with an associated I-SID (Layer 2 VSN).

The following table describes the fields for the **`show ip igmp interface vlan`** command.

Parameter	Description
VLAN ID	Identifies the VLAN 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

Parameter	Description
	correctly, you must configure all routers on a LAN to run the same version of IGMP.
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	Indicates the set of ports that are enabled for fast leave.
VLAN ID	Identifies the VLAN where IGMP is configured.
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 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.

The following table describes the fields for the **`show ip igmp cache`** command.

Parameter	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.
EXPIRATION	Indicates the minimum amount of time that remains before this entry ages out.

Parameter	Description
V1HOSTIMER	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.

The following table describes the fields for the **`show ip igmp group`** command.

Parameter	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 minimum amount of time that remains before this entry ages out.
TYPE	Indicates whether the entry is learned dynamically or is added statically.

The following table describes the fields for the **`show ip igmp sender`** command.

Parameter	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.

Viewing TLV information for Layer 2 VSN IP multicast over SPBM

Use the following commands to check TLV information.

For Layer 2 VSN with IP multicast SPBM, TLV 185 on the BEB where the source is located, displays the multicast source and group addresses and has the Tx bit set. Each multicast group has its own unique data I-SID with a value between 16,000,000 to 16,512,000. TLV 144 on the BEB bridge, where the sender is located, has the Tx bit set. All BEB bridges, where a receiver exists, have the Rx bit set.

Procedure

1. Log on to the switch to enter User EXEC mode.

2. Display IS-IS Link State Database information by Type-Length-Value (TLV):

   ```
   show isis lsdb tlv <1-186> [sub-tlv <1-3>][detail]
   ```

3. Display IS-IS Link State Database information by Link State Protocol ID:

   ```
   show isis lsdb lspid <xxxx.xxxx.xxxx.xx-xx>tlv <1-186> [sub-tlv <1-3>] [detail]
   ```

Example

```
Switch:1# show isis lsdb tlv 185 detail
================================================================================
            ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------
Level-1LspID: 000c.f803.83df.00-00 SeqNum: 0x000001ae Lifetime: 898
Chksum: 0xcebe PDU Length: 522
Host_name: Switch
Attributes: IS-Type 1
TLV:185 SPBM IPVPN :
VSN ISID:5010
BVID :10
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.1
        Data ISID : 16300001
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.3
        Data ISID : 16300003
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.5
        Data ISID : 16300005
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.7
        Data ISID : 16300007
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.9
        Data ISID : 16300009
        TX : 1
        VSN ISID:5010
        BVID :20
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.2
        Data ISID : 16300002
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.4
        Data ISID : 16300004
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.6
        Data ISID : 16300006
```

```
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.8
        Data ISID : 16300008
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.1
        Group Address : 233.252.0.10
        Data ISID : 16300010
        TX : 1

Switch:1# show isis lsdb lspid 000c.f803.83df.00-05 tlv 144 detail
================================================================================
        ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------
Level-1 LspID: 000c.f803.83df.00-00 SeqNum: 0x00000477 Lifetime: 903
Chksum: 0x200b PDU Length: 522
Host_name: Switch
Attributes: IS-Type 1
    Instance: 0
    Metric: 0
    B-MAC: 03-00-00-00-00-00
    BVID:10
    Number of ISID's:5
        16000001(Tx),16000003(Tx),16000005(Tx),16000007    (Tx),16000009(Tx)
    Instance: 0
    Metric: 0
    B-MAC: 03-00-00-00-00-00
    BVID:20
    Number of ISID's:5
        16000002(Tx),16000004(Tx),16000006(Tx),16000008(Tx),16000010(Tx)
```

Variable definitions

Use the data in the following table to use the **`show isis lsdb`** command.

Variable	Value
detail	Displays detailed information about the IS-IS Link State database.
level *{l1, l2, l12}*	Displays information on the IS-IS level. The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 and combined Level 1 and 2 (l12) function is disabled in the current release.
local	Displays information on the local LSDB.
lspid <*xxxx.xxxx.xxxx.xx-xx*>	Specifies information about the IS-IS Link State database by LSP ID.
sub-tlv <*1–3*>	Specifies information about the IS-IS Link State database by sub-TLV.
sysid <*xxxx.xxxx.xxxx*>	Specifies information about the IS-IS Link State database by System ID.
tlv <*1–186*>	Specifies information about the IS-IS Link State database by TLV.

Job aid

The following table describes the fields for the **show isis lsdb** command.

Parameter	Description
LSP ID	Indicates the LSP ID assigned to external IS-IS routing devices.
LEVEL	Indicates the level of the external router: l1, l2, or l12.
LIFETIME	Indicates the maximum age of the LSP. If the max-lsp-gen-interval is set to 900 (default) then the lifetime value begins to count down from 1200 seconds and updates after 300 seconds if connectivity remains. If the timer counts down to zero, the counter adds on an additional 60 seconds, and then the LSP for that router is lost. This situation happens because of the zero age lifetime, which is detailed in the RFC standards.
SEQNUM	Indicates the LSP sequence number. This number changes each time the LSP is updated.
CKSUM	Indicates the LSP checksum. The checksum is an error checking mechanism used to verify the validity of the IP packet.
HOST-NAME	Indicates the host-name.

Layer 2 VSN configuration using EDM

This section provides procedures to configure Layer 2 Virtual Services Networks (VSNs) using Enterprise Device Manager (EDM).

Configuring SPBM Layer 2 VSN

After you have configured the SPBM infrastructure, you can enable the SPBM Layer 2 Virtual Service Network (VSN) using the following procedure.

SPBM supports Layer 2 VSN functionality where customer VLANs (C-VLANs) are bridged over the SPBM core infrastructure.

At the BEBs, customer VLANs (C-VLAN) are mapped to I-SIDs based on the local service provisioning. Outgoing frames are encapsulated in a MAC-in-MAC header, and then forwarded across the core to the far-end BEB, which strips off the encapsulation and forwards the frame to the destination network based on the I-SID-to-C-VLAN provisioning.

Note:

Different from Ethernet Routing Switch 8800 SPBM implementation, on Virtual Services Platform 9000 when you map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID), you must include the IST MLT peer switches in the VLAN, which means all IST member ports must be members of this C-VLAN. In ERS 8800, you cannot configure IST MLT to be part

of the C-VLAN. For more information about this and other differences between the two products, see *Platform Migration Reference for Avaya Virtual Services Platform 9000,* NN46250-107.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Advanced** tab.
4. To map a C-VLAN to a Service instance identifier (I-SID), in the **I-sid** field, specify the I-SID to associate with the specified VLAN.
5. Click **Apply**.

Important:

- When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.
- Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

Displaying the remote MAC table for a C-VLAN

Use the following procedure to view a the remote MAC table for a C-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click **Basic** tab and highlight a C-VLAN.
4. Click the **Remote MAC** tab.

Remote MAC field descriptions

Use the data in the following table to use the **Remote MAC** tab.

Name	Description
VlanId	Indicates the VLAN ID for this MAC address.

Name	Description
Addr	Indicates the customer MAC address for which the bridge has forwarding and/or filtering information
DestAddr	Indicates the provider MAC address for which the bridge has forwarding and/or filtering information.
PrimaryBVlanId	Indicates the primary B-VLAN ID for this MAC address.
PrimaryDestSysName	Indicates the primary system name of the node where the MAC address entry comes from.
PrimaryPort	Either displays the value 0, or indicates the primary port on which a frame came from.
SecondaryBVlanId	Indicates the secondary B-VLAN ID for this MAC address
SecondaryDestSysName	Indicates the secondary system name of the node where the MAC address entry comes from.
SecondaryPort	Either displays the value 0, or indicates the secondary port on which a frame came from.
SmltRemote	Indicates the MAC address entry for the remote IST peer.
Status	Indicates the status of this entry: • other • invalid • learned • self • mgmt

Configuring IP multicast over SPBM on a Layer 2 VSN

Use this procedure to enable IP multicast over SPBM for a Layer 2 VSN. With Layer 2 VSN IP multicast over SPBM, multicast traffic remains in the same Layer 2 VSN across the SPBM cloud.

No explicit configuration exists for a Layer 2 VSN. After you configure IP IGMP snooping on a VLAN that has an I-SID configured, the device enables that VLAN for SPBM multicast services.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology.
- You must enable SPBM multicast globally.

About this task

SPBM supports enabling IGMP snooping on a C-VLAN, but it does not support enabling PIM on a C-VLAN. If you enable IGMP snooping on a C-VLAN, then its operating mode is Layer 2 VSN with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network.

In this release, the switch only supports IPv4 multicast traffic.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Select the **IGMP** tab.
7. Select the **SnoopEnable** check box.
8. **(Optional)** Select the **SsmSnoopEnable** check box, if you use IGMP version 3.
9. **(Optional)** Select the **ProxySnoopEnable** check box.
10. **(Optional)** If you want to enable IGMP version 3, select version3 in the **Version** check box.

    You must enable SSM snoop before you configure IGMP version 3 and both ssm-snoop and snooping must be enabled for IGMPv3.
11. If you want to enable IGMP version 2, select version2 in the **Version** check box.

    For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.
12. **(Optional)** If you want to enable snoop querier, select **SnoopQuerierEnable**.
13. **(Optional)** If you want to configure an address for IGMP queries, enter the IP address in **SnoopQuerierAddr**.

    Note:

    This step is not always required. The IGMP Querier on the BEB uses a source address 0.0.0.0 by default. When you do not configure this, a BEB sends IGMP queries on the UNI ports with 0.0.0.0 as the source IP address. Some Layer 2 edge switches do not support a 0.0.0.0 querier. You can use a fictitious IP address as the querier address, and use the same address on all BEBs in the network.
14. Click **Apply**.

Viewing the IGMP interface table

Use the Interface tab to view the IGMP interface table. When an interface does not use an IP address, it does not appear in the IGMP table.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IGMP**.
3. Click the **Interface** tab.

IGMP Interface field descriptions

Use the data in the following table to configure the **Interface** tab.

Name	Description
IfIndex	Displays the interface where IGMP is enabled.
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The default is 125.
Status	Displays the IGMP row status.
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 currently running on this interface.
Querier	Shows the address of the IGMP querier on the IP subnet to which this interface is attached.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255. The default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the QueryInterval.
WrongVersionQueries	Displays 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 queries are received with the wrong version, this value indicates a version mismatch.
Joins	Displays the number of times a group membership is added on this interface; that is, 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 a network is expected 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 one query for each query interval is dropped without the querier aging out.

Name	Description
LastMembQueryIntvl	Configures the maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable 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 second. Avaya recommends that you configure this parameter to values greater than 3. If a fast leave process is not required, Avaya recommends values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
OtherQuerierPresentTimeout	Specifies the timeout interval.
FlushAction	Configures the flush action to one of the following: • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	Displays if the router alert IP is enabled. When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set. **Important:** To maximize your network performance, Avaya recommends that you set this parameter according to the version of IGMP currently in use. • IGMPv1 — Disable • IGMPv2 — Enable • IGMPv3 — Enable
SsmSnoopEnable	Enables SSM snoop. The default is disabled.
SnoopQuerierEnable	Enables snoop querier.
SnoopQuerierAddr	Specifies the pseudo address of the IGMP snoop querier.
ExplicitHostTrackingEnable	Enables the IGMP protocol running in version 3 to track hosts for each channel or group. The default is false.

Name	Description
Mcast Mode	Specifies the multicast mode: • snoop • pim • snoopSpb • routedSpb • none The default is none.

Layer 2 VSN configuration examples

This section provides configuration examples to configure Layer 2 VSNs.

Layer 2 VSN configuration example

The following figure shows a sample Layer 2 VSN deployment.

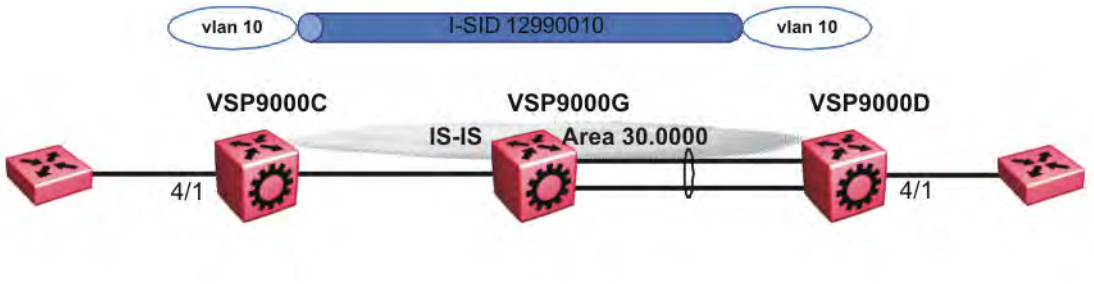

Figure 14: Layer 2 VSN

The following sections show the steps required to configure the Layer 2 VSN parameters in this example. You must first configure basic SPBM and IS-IS infrastructure. For more information, see: SPBM configuration examples on page 114.

VSP9000C

```
VLAN CONFIGURATION

vlan create 10 type port-mstprstp 1
vlan members 10 4/1 portmember
vlan i-sid 10 12990010
```

VSP9000D

```
VLAN CONFIGURATION

vlan create 10 type port-mstprstp 1
vlan members 10 4/1 portmember
vlan i-sid 10 12990010
```

Verifying Layer 2 VSN operation

The following sections show how to verify the Layer 2 VSN operation in this example.

VSP9000C

```
VSP9000C:1# show isis spbm i-sid all
================================================================================
                                SPBM ISID INFO
================================================================================
ISID   SOURCE NAME    VLAN   SYSID               TYPE                    HOST_NAME
--------------------------------------------------------------------------------
12990010 f.30.14        4000   0014.0da0.13df       discover        VSP9000D
12990010 f.30.13        4000   0015.e89f.e3df       config            VSP9000C

--------------------------------------------------------------------------------
 Total number of SPBM ISID entries configured: 1
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries discovered: 1
--------------------------------------------------------------------------------
 Total number of SPBM ISID entries: 2
--------------------------------------------------------------------------------

VSP9000C:1#show isis spbm multicast-fib

================================================================================

              SPBM MULTICAST FIB ENTRY INFO
================================================================================
MCAST DA           ISID    BVLAN SYSID          HOST-NAME OUTGOING-INTERFACES INCOMING
                                                                      INTERFACE
--------------------------------------------------------------------------------
0f3:30:14:c6:36:3a 12990010 4000 0014.0da0.13df VSP9000D  4/1                  4/10
f3:30:13:c6:36:3a  12990010 4000 0015.e89f.e3df VSP9000C  4/30,4/1             4/10


--------------------------------------------------------------------------------
 Total number of SPBM MULTICAST FIB entries 2
--------------------------------------------------------------------------------
```

VSP9000D

```
VSP9000D:1# show isis spbm i-sid all
================================================================================
                                SPBM ISID INFO
================================================================================
ISID   SOURCE NAME    VLAN   SYSID               TYPE                    HOST_NAME
--------------------------------------------------------------------------------
12990010 f.30.14        4000   0014.0da0.13df       config            VSP9000D
12990010 f.30.13        4000   0015.e89f.e3df       discover        VSP9000C

VSP9000D:1#show isis spbm multicast-fib

================================================================================

              SPBM MULTICAST FIB ENTRY INFO
================================================================================
MCAST DA           ISID    BVLAN SYSID          HOST-NAME OUTGOING-INTERFACES INCOMING
                                                                      INTERFACE
--------------------------------------------------------------------------------
f3:30:14:c6:36:3a  12990010 4000 0014.0da0.13df VSP9000D  MLT-1,4/1          4/10
f3:30:13:c6:36:3a  12990010 4000 0015.e89f.e3df VSP9000C  4/1                4/10
--------------------------------------------------------------------------------
 Total number of SPBM MULTICAST FIB entries 2
--------------------------------------------------------------------------------
```

VSP9000C — verifying with CFM

```
VSP9000C:1# l2 tracetree 4000 12990010

 Please wait for l2tracetree to complete or press any key to abort
```

```
l2tracetree to f3:30:13:c6:36:3a, vlan 4000 i-sid 12990010 nickname f.30.13 hops 64
1   VSP9000C            00:15:e8:9f:e3:df -> VSP9000G            00:0e:62:25:a3:df
2   VSP9000G            00:0e:62:25:a3:df -> VSP9000D            00:14:0d:a0:13:df
```

VSP9000D — verifying with CFM

```
VSP9000D:1# l2 tracetree 4000 12990010

 Please wait for l2tracetree to complete or press any key to abort

l2tracetree to f3:30:14:c6:36:3a, vlan 4000 i-sid 12990010 nickname f.30.14 hops 64
1   VSP9000D            00:14:0d:a0:13:df -> VSP9000G            00:0e:62:25:a3:df
2   VSP9000G            00:0e:62:25:a3:df -> VSP9000C            00:15:e8:9f:e3:df
```

VSP9000C — verifying FDB

```
VSP9000C:1# show vlan mac-address-entry 10
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                  SMLT
ID   STATUS     ADDRESS            INTERFACE         REMOTE  TUNNEL
--------------------------------------------------------------------------------
10   learned    00:00:00:00:00:01  Port-4/1          false    VSP9000D
10   learned    00:00:00:00:00:02  Port-4/1          false    VSP9000D

2 out of 4 entries in all fdb(s) displayed.

VSP9000C:1# show vlan remote-mac-table 10
================================================================================
                            Vlan Remote Mac Table
================================================================================
VLAN STATUS  MAC-ADDRESS        DEST-MAC           BVLAN    DEST-SYSNAME PORTS SMLTREMOTE
--------------------------------------------------------------------------------
10   learned 00:00:00:00:00:02  00:14:0d:a0:13:df  0014.0da0.13df VSP9000D 4/30 false
--------------------------------------------------------------------------------
 Total number of VLAN Remote MAC entries 1
 -------------------------------------------------------------------------------
```

VSP9000D — verifying FDB

```
VSP9000D:1# show vlan mac-address-entry 10
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                  SMLT
ID   STATUS     ADDRESS            INTERFACE         REMOTE  TUNNEL
--------------------------------------------------------------------------------
10   learned    00:00:00:00:00:01  Port-4/1          false   VSP9000C
10   learned    00:00:00:00:00:02  Port-4/1          false   VSP9000C

2 out of 4 entries in all fdb(s) displayed.

VSP9000D:1# show vlan remote-mac-table 10
================================================================================
                            Vlan Remote Mac Table
================================================================================
VLAN STATUS MAC-ADDRESS        DEST-MAC          DEST-SYSID DEST-SYSNAME PORTS SMLTREMOTE
--------------------------------------------------------------------------------
10   learned 00:00:00:00:00:01 00:15:e8:9f:e3:df 0015.e89f.e3df VSP9000C MLT-1    false
--------------------------------------------------------------------------------
 Total number of VLAN Remote MAC entries 1
 -------------------------------------------------------------------------------
```

Layer 2 VSN example with VLAN ID translation

The following figure shows a sample Layer 2 VSN deployment where the C- VLAN IDs are different at each end. You must first configure basic SPBM and IS-IS infrastructure. For more information, see SPBM configuration examples on page 114.

Figure 15: Layer 2 VSN with different VLAN IDs

The following sections show the steps required to configure the Layer 2 VSN parameters in this example.

VSP9000C

```
VLAN CONFIGURATION

vlan create 9 type port 1
vlan members 9 4/1 portmember
vlan i-sid 9 129900009
```

VSP9000D

```
VLAN CONFIGURATION

vlan create 19 type port 1
vlan members 19 4/1 portmember
vlan i-sid 19 129900009
```

Layer 2 VSN with IP multicast over SPBM configuration example

The example below shows the configuration steps to enable IP multicast over SPBM support on C-VLAN 1001 that is part of a Layer 2 VSN, including the querier address.

```
enable
configure terminal

ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable

VLAN CONFIGURATION

interface vlan 9
ip igmp snooping
ip igmp snoop-querier-addr 192.0.2.201
exit
```

When using IGMPv3, the configuration is:

```
enable
configure terminal

ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable


VLAN CONFIGURATION

interface vlan 19
ip igmp snooping
ip igmp version 3
ip igmp ssm-snoop
ip igmp snoop-querier-addr 192.0.2.201
exit
```

Note:

You must enable SSM snoop before you configure IGMP version to version 3, and you must enable both **`ssm-snoop`** and **`snooping`** for IGMPv3.

Note:

You must configure basic SPBM and IS-IS infrastructure. For more information, see SPBM configuration examples on page 114.

IP Shortcuts configuration

This section provides concepts and procedures to configure IP Shortcuts.

IP Shortcuts configuration fundamentals

This section provides fundamental concepts for IP Shortcuts.

SPBM IP shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

Unlike Layer 2 VSN, with SPBM IP shortcuts, no I-SID configuration is required. Instead, SPBM nodes propagate Layer 3 reachability as “leaf” information in the IS-IS LSPs using Extended IP reachability TLVs (TLV 135), which contain routing information such as neighbors and locally configured subnets. SPBM nodes receiving the reachability information can use this information to populate the routes to the announcing nodes. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes.

The following figure shows a network running SPBM IP shortcuts.

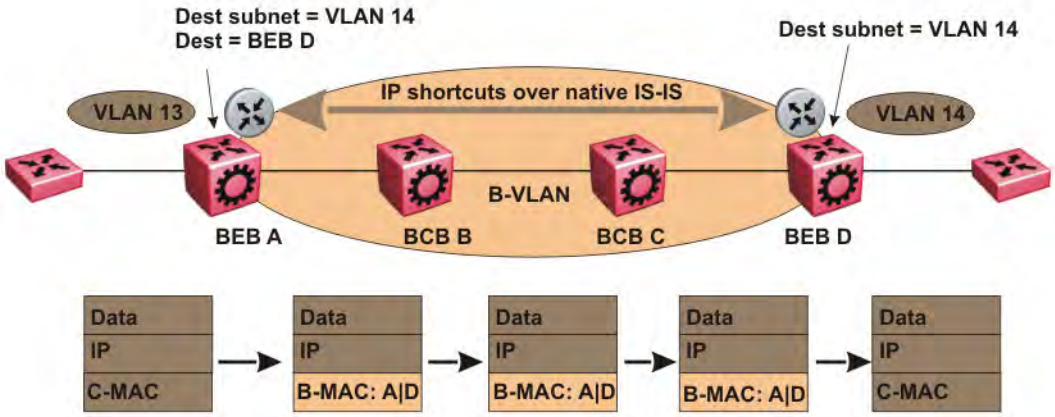

Figure 16: SPBM IP Shortcuts

In this example, BEB A receives a packet with a destination IP address in the subnet of VLAN 14 and knows to forward the packet to BEB D based on the IP route propagation within IS-IS. After a route lookup, BEB A knows that BEB D is the destination for the subnet and constructs a new B-MAC header with destination B-MAC: D. BCBs B and C need only perform normal Ethernet switching to forward the packet to BEB D. A route lookup is only required once, at the source BEB, to identify BEB D as the node that is closest to the destination subnet.

In contrast to IP routing or Multiprotocol Label Switching (MPLS), SPBM IP shortcuts provide a simpler method of forwarding IP packets in an Ethernet network using the preestablished Ethernet FIBs on the BEBs. SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet switching based on the existing SPT. This allows for end to end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

In the above example, the SPBM nodes in the core that are not enabled with IP shortcuts can be involved in the forwarding of IP traffic. Since SPBM nodes only forward on the MAC addresses that comprise the B-MAC header, and since unknown TLVs in IS-IS are relayed to the next hop but ignored locally, SPBM nodes need not be aware of IP subnets to forward IP traffic.

With IP shortcuts, there is only one IP routing hop, as the SPBM backbone acts as a virtualized switching backplane.

The following figure shows a sample campus network implementing SPBM IP shortcuts.

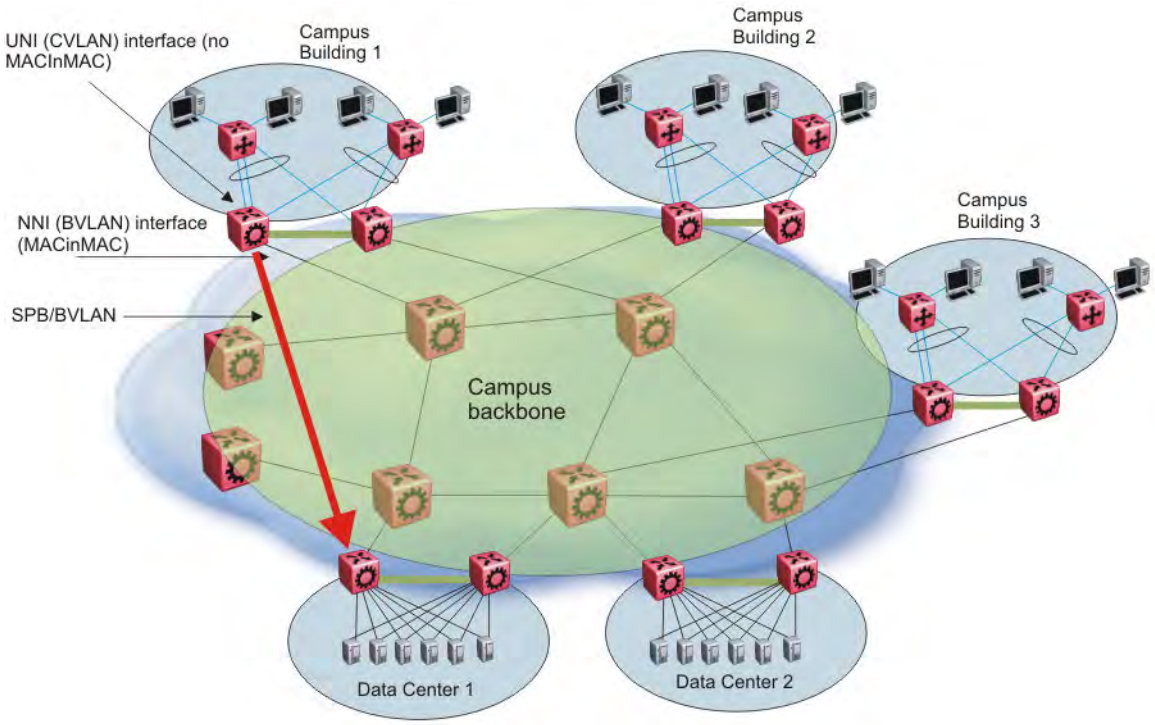

Figure 17: SPBM IP shortcuts in a campus

To enable IP shortcuts on the BEBs, you must configure a circuitless IP address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct, static, OSPF, RIP, or BGP routes into IS-IS.

ECMP within ISIS routes

Equal Cost Multipath (ECMP) allows the device to determine up to eight equal cost paths to the same destination prefix. If the device learns the same route from multiple sources, the information is ECMP only if the routes:

- are from the same VSN
- have the same SPBM cost
- have the same prefix cost
- have the same IP route preference

Multiple BEBs can announce the same route, either because the Layer 2 LAN connects to multiple BEBs for redundancy, or because segments of the LAN are Layer 2 bridged. If the device has to tie-break between the multiple sources, the device uses the following precedence rules to tie-break. In the following order, the device prefers:

1. Local routes over Inter-VSN routes.
2. Routes with the lowest route preference.

   You can change this with route-map within the IS-IS accept policy.

3. Routes with the lowest SPBM cost.
4. Routes with lowest prefix cost.

   You can change this with route-map on the remote advertising node with the **`redistribute`** command, or with route-map on the local node with the IS-IS accept policy.
5. Routes from the VSN with the lower Layer 3 VSN I-SID.

   The device considers the Global Routing Table (GRT) to have an I-SID equal to zero.

When you use multiple B-VLANs in the SPBM core, multiple paths exist to reach a particular SPBM node, one on each B-VLAN; therefore, any IP prefix the device receives from a BEB results in multiple ECMP paths. These paths may or may not be physically diverse. SPBM currently supports up to two B-VLANs; a primary B-VLAN and a secondary B-VLAN.

By default, when ECMP is enabled up to eight equal paths can exist to a destination, but you can change this number to a value from 1 to 8.

If more ECMP paths are available than the configured number of paths, then the device adds the routes using the following order. The device selects all routes from the primary B-VLAN and orders the routes learned through that B-VLAN from lowest system ID to the highest IS-IS system ID, then device moves on to select all routes from the secondary B-VLAN, ordering those routes from lowest IS-IS system ID to the highest IS-IS system ID until you reach the number of equal paths configured.

For example, if the SPB core is configured with two B-VLANs (primary B-VLAN 1000 and secondary B-VLAN 2000), and the device learns routes from two BEBs called BEB-A (with a lower IS-IS system ID) and BEB-B (with a higher IS-IS system ID), the order in which the next-hop is chosen for that route is as follows:

1. BEB-A B-VLAN 1000
2. BEB-B B-VLAN 1000
3. BEB-A B-VLAN 2000
4. BEB-B B-VLAN 2000

If ECMP is disabled, then the route the device adds is from the lowest system ID with the primary B-VLAN. In this example, the device adds BEB-A B-VLAN 1000.

IS-IS IP redistribution policies

When you connect an SPBM core using IP shortcuts to existing networks running a routing protocol such as OSPF or RIP, a redundant configuration requires two Virtual Services Platform 9000s:

- Both routers redistribute IP routes from Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) into IS-IS (IP) and redistribute IS-IS (IP) routes into RIP or OSPF. This can create a routing loop, special precaution need to be taken to prevent this.

The following figure illustrates this configuration.

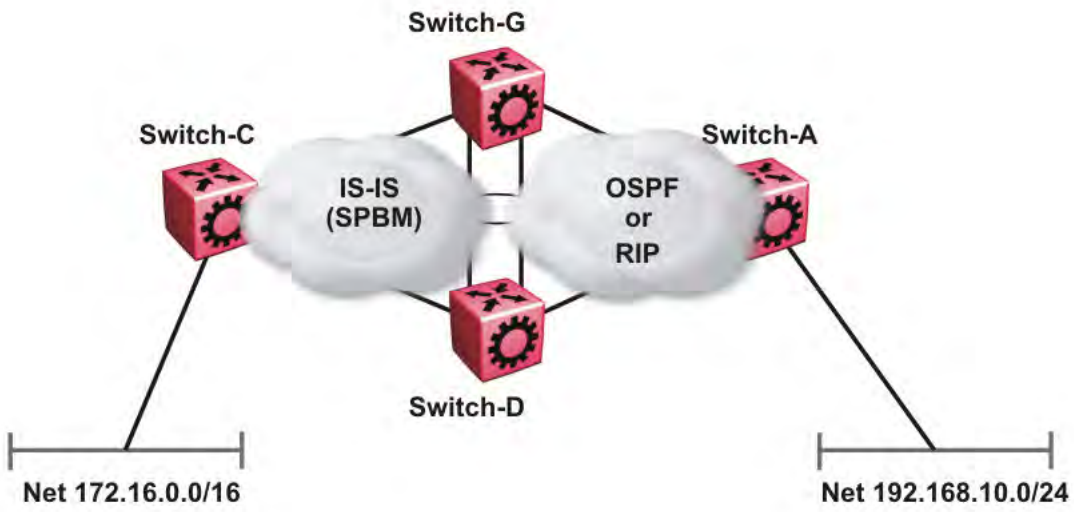

Figure 18: Redundant OSPF or RIP network

In this scenario it is necessary to take extra care when redistributing through both VSP 9000s. By default the preference value for IP routes generated by SPBM-IP (IS-IS) is 7. This is a higher preference than OSPF (20 for intra-area, 25 for inter-area, 120 for ext type1, 125 for ext type2) or RIP (100).

Important:

The lower numerical value determines the higher preference.

In the diagram above both nodes (VSP 9000G and VSP 9000D) have an OSPF or a RIP route to 192.168.10.0/24 with the next-hop to VSP 9000A.

As soon as the VSP 9000G node redistributes that IP route into IS-IS, the VSP 9000D node learns the same route through IS-IS from VSP 9000G. (The VSP9000G node already has the route through OSPF or RIP). Because IS-IS has a higher preference, VSP 9000D replaces its 192.168.10.0 OSPF route with an IS-IS one that points at VSP 9000G as the next-hop. The following figure illustrates this scenario.

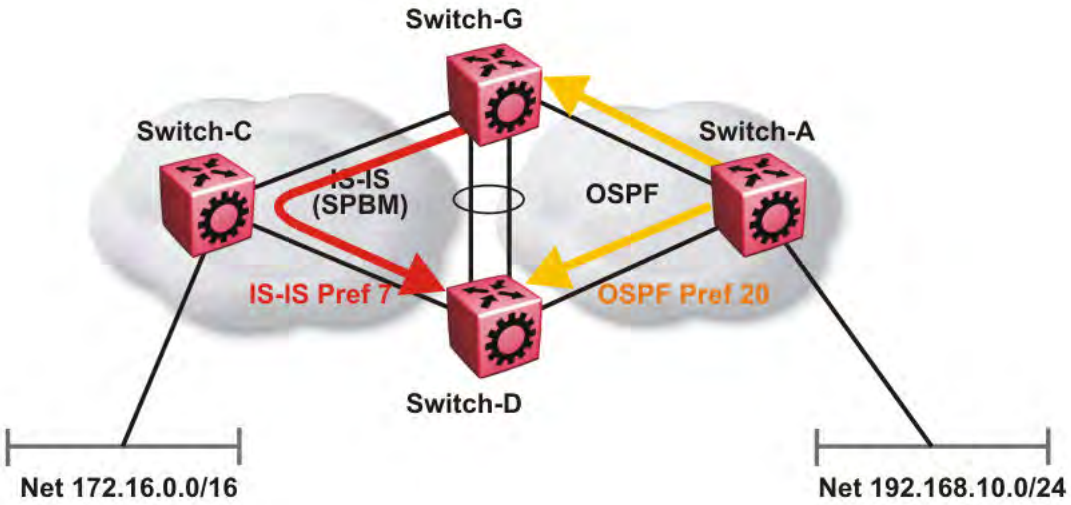

Figure 19: Redistributing routes into IS-IS

Clearly this is undesirable and care needs to be taken to ensure that the two redistributing nodes (VSP 9000G and VSP 9000D) do not accept redistributed routes from each other. With IS-IS accept policies, you can associate an IS-IS accept policy on VSP 9000D to reject all redistributed IP routes received from VSP 9000G, and VSP 9000G to reject all redistribute IP routes from VSP 9000D.

An alternate way to solve the preceding problem with existing functionality is to reverse the problem by lowering the SPBM-IP (IS-IS) preference by configuring it to a value greater than RIP (100) or OSPF (20,25,120,125). For example, log on to Global Configuration mode and use the following command to configure a preference of 130:

```
ip route preference protocol spbm-level1 130
```

Now that the OSPF or RIP routes have a higher preference than SPBM-IP (IS-IS), the above problem is temporarily solved. However, the same issue resurfaces when the IS-IS IP routes are redistributed into OSPF or RIP in the reverse direction as shown in the following figure for OSPF:

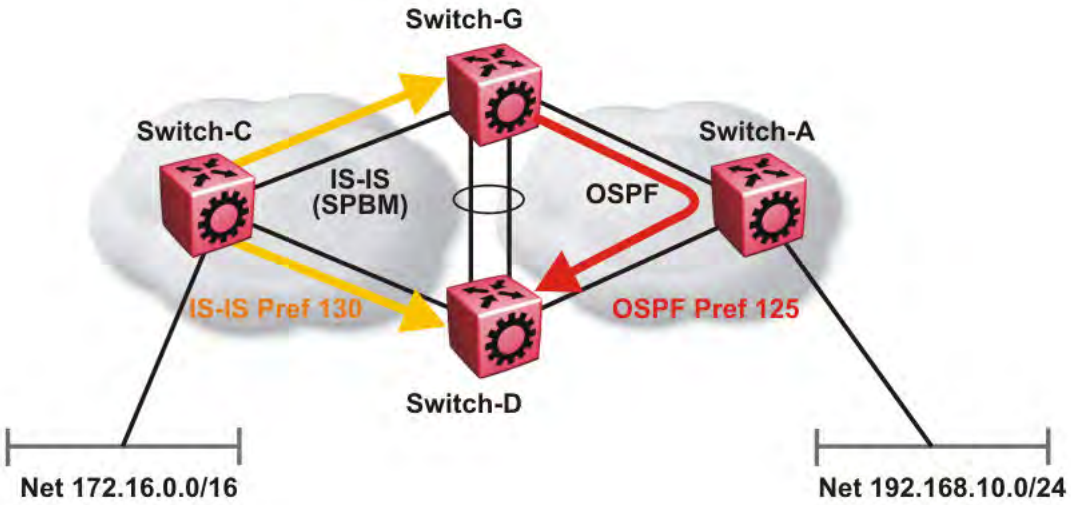

Figure 20: Redistributing routes into OSPF

In the preceding figure, both VSP 9000G and VSP 9000D have an IS-IS IP route for 172.16.0.0/16 with the next hop as VSP 9000C. As soon as the VSP 9000G redistributes the IS-IS route into OSPF, the VSP 9000D node learns that same route through OSPF from VSP 9000G. (The VSP 9000G node already has the route through IS-IS).

Because OSPF has a higher preference, VSP 9000D replaces its 172.16.0.0/16 IS-IS route with an OSPF one. (Note that the 172.16.0.0/16 route will be redistributed into OSPF as an AS external route, hence with preference 120 or 125 depending on whether type1 or type2 was used). In this case, however, you can leverage OSPF Accept policies, which can be configured to prevent VSP 9000D from accepting any AS External (LSA5) routes from VSP 9000G and prevent VSP 9000G from accepting any AS External (LSA5) routes from VSP 9000D. The following is a sample configuration:

```
enable
configure terminal
route-map

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "reject" 1
no permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

OSPF CONFIGURATION - GlobalRouter

router ospf enable

OSPF ACCEPT CONFIGURATION - GlobalRouter

router ospf
accept adv-rtr {A.B.C.D}
```

```
accept adv-rtr {A.B.C.D} enable route-policy "reject"
exit
```

Note:

Avaya recommends you disable alternative routes by issuing the command **no ip alternative-route** to avoid routing loops on the SMLT Backbone Edge Bridges (BEBs).

In the preceding figure, if VSP 9000A advertises 25000 OSPF routes to VSP 9000G and VSP 9000D, then both VSP 9000G and VSP 9000D install the 25000 routes as OSPF routes. Since VSP 9000D and VSP 9000G have OSPF to IS-IS redistribution enabled, they also learn these 25000 routes as IS-IS routes. IS-IS route preference is configured with a higher numerical value (130) than the OSPF route preference (125), so VSP 9000D and VSP 9000G keep IS-IS learned routes as alternative routes.

If VSP 9000A withdraws its 25000 OSPF routes, VSP 9000G and VSP 9000D remove the OSPF routes. While the OSPF routes are removed the routing tables of VSP 9000G and VSP 9000D activate the alternative IS-IS routes for the same prefix. Since VSP 9000G and VSP 9000D have IS-IS to OSPF redistribution enabled, VSP 9000A learns these routes as OSPF and this causes a routing loop. Use the **no ip alternative-route** command to disable alternative routes on VSP 9000G and VSP 9000D to avoid routing loops.

In the preceding figure, you leveraged OSPF Accept policies, which can be configured to prevent VSP 9000D from accepting any AS External (LSA5) routes from VSP 9000G and prevent VSP 9000G from accepting any AS External (LSA5) routes from VSP 9000D. In the case of a RIP access network, the preceding solution is not possible because RIP has no concept of external routes and no equivalent of accept policies. However, if you assume that a RIP network acts as an access network to an SPBM core, then it is sufficient to ensure that when IS-IS IP routes are redistributed into RIP they are aggregated into a single default route at the same time. The following figure and sample configuration example illustrates this scenario:

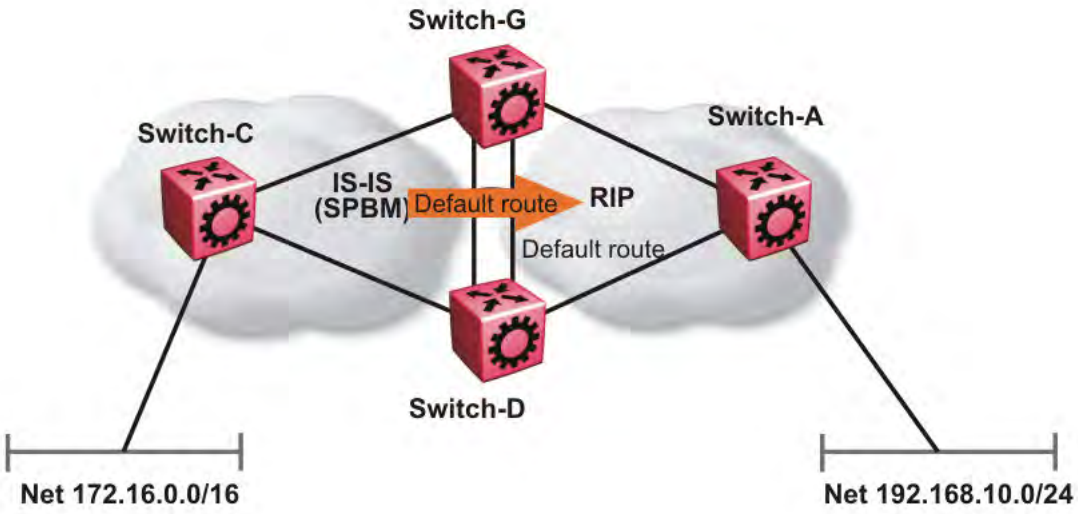

Figure 21: Redistributing routes into RIP

VSP 9000G

```
IP PREFIX LIST CONFIGURATION - GlobalRouter

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

RIP PORT CONFIGURATION

interface gigabitethernet 3/11
ip rip default-supply enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter
```

```
router rip
redistribute isis
redistribute isis metric 1
redistribute isis route-map "inject-default"
redistribute isis enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
```

VSP 9000A

```
RIP PORT CONFIGURATION

interface gigabitethernet 3/2
ip rip default-listen enable
exit
interface gigabitethernet 3/3
ip rip default-listen enable
exit
```

VSP 9000D

```
IP PREFIX LIST CONFIGURATION - GlobalRouter

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

RIP PORT CONFIGURATION

interface gigabitethernet 4/11
ip rip default-supply enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router rip
redistribute isis
redistribute isis metric 1
redistribute isis route-map "inject-default"
redistribute isis enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
```

You can control the propagation of the default route on the RIP network so that both VSP 9000G and VSP 9000D supply the default route on their relevant interfaces, and not accept it on the same interfaces. Likewise, VSP 9000A will accept the default route on its interfaces to both VSP9000G and VSP9000D but it will not supply the default route back to them. This will prevent the default route advertised by VSP9000G from being installed by VSP9000D, and vice-versa.

The above example where IS-IS IP routes are aggregated into a single default route when redistributed into the RIP network can also be applied when redistributing IS-IS IP routes into OSPF if that OSPF network is an access network to an SPBM core. In this case use the following redistribution policy configuration as an example for injecting IS-IS IP routes into OSPF:

```
IP PREFIX LIST CONFIGURATION - GlobalRouter

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit


OSPF CONFIGURATION - GlobalRouter

router ospf enable
router ospf
as-boundary-router enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router ospf
redistribute isis
redistribute isis route-policy "inject-default"
redistribute isis enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip ospf apply redistribute isis
```

IS-IS accept policies

You can use Intermediate-System-to-Intermediate-System (IS-IS) accept policies to filter incoming IS-IS routes over the SPBM cloud and apply route policies to the incoming IS-IS routes. IS-IS accept policies enable the device to determine whether to add an incoming route to the routing table.

IS-IS accept policy filters

You can filter traffic with IS-IS accept policies by:

- advertising BEB
- I-SID or I-SID list
- route-map

You can use IS-IS accept policies to apply at a global default level for all advertising Backbone Edge Bridges (BEBs) or for a specific advertising BEB.

IS-IS accept policies also allow you to use either a service instance identifier (I-SID) or an I-SID list to filter routes. The switch uses I-SIDs to define Virtual Services Networks (VSNs). I-SIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. IS-IS accept policies can use I-SIDs or I-SID lists to filter the incoming virtualized traffic.

IS-IS accept policies can also apply route policies to determine what incoming traffic to accept into the routing table. With route policies the device can determine which routes to accept into the routing table based on the criteria you configure. You can match on the network or the route metric.

For more information on configuring route policies, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

The following table describes IS-IS accept policy filters.

Filters into	Filter	Description
Global Routing Table (GRT)	accept route-map *WORD<1-64>*	By default, the device accepts all routes into the GRT and VRF routing table. This is the default accept policy.
	accept adv-rtr <*x.xx.xx*>route-map *WORD<1-64>*	The device filters based on the specific advertising BEB defined by the SPBM nickname.
	accept i-sid <*1-16777215*>route-map *WORD<1-64>*	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN.
	accept adv-rtr<*x.xx.xx*> i-sid <*1-16777215*>route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN.
	accept isid-list *WORD<1-32>*route-map *WORD<1-64>*	The device filters based on the list of I-SIDs.
	accept adv-rtr <*x.xx.xx*> isid-list *WORD<1-32>*route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT).
	accept adv-rtr <*x.xx.xx*>route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the route policy.

Filters into	Filter	Description
Virtual Routing and Forwarding (VRF) routing table	isis accept adv-rtr *<x.xx.xx>*	The device filters based on the specific advertising BEB defined by the SPBM nickname.
	isis accept i-sid *<0-16777215>*	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT).
	isis accept adv-rtr *<x.xx.xx>*i-sid *<0-16777215>*	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT).
	isis accept isid-list *WORD<1-32>*	The device filters based on the list of I-SIDs to which the IS-IS accept policy applies. The number 0 represents the Global Routing Table (GRT).
	isis accept adv-rtr *<x.xx.xx>* isid-list *WORD<1-32>*	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT).
	isis accept route-map *WORD<1-64>*	The device filters based on the route policy.
	isis accept adv-rtr *<x.xx.xx>* route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the route policy.

IS-IS accept policies for the GRT and VRFs

You can create an IS-IS accept policy for incoming routes for the Global Routing Table (GRT), which accepts routes into the routing table, or for a Virtual Routing and Forwarding (VRF) instance, which accepts incoming routes to the routing table of the VRF.

If you create an IS-IS accept policy on the switch for either the GRT or a VRF that operates at a global default level, the accept policy applies to all routes for all BEBs in the GRT or VRF.

If you create an IS-IS accept policy on the switch for a specific advertising BEB for either the GRT or a VRF, the IS-IS accept policy instance applies for that specific advertising BEB. If you use a more specific filter, the system gives preference to the specific filter over the global default level.

IS-IS accept policies for inter-VRF route redistribution

You can also use the filter mechanism for IS-IS accept policies to redistribute routes between different VRFs, or between a VRF and the GRT. For inter-VRF route redistribution, you match the filter based on the I-SID, which represents the Layer 3 VSN context.

You can apply the filter at the global default level, where the IS-IS accept policy applies to all routes for that I-SID from all BEBs, or at a specific advertising BEB level, where the filter only applies to a specific advertising BEB. The device gives preference to a specific filter for a specific advertising BEB over the global default filter.

For inter-VRF route redistribution, an I-SID value of 0 represents the GRT. For inter-VRF route redistribution between VRFs, the I-SID is the source VRF (or remote VRF).

IS-IS accept policy considerations

Consider the following when you configure IS-IS accept policies:

- The switch does not support IS-IS accept policies for IPv6 addresses for the current release.
- If a VRF uses a different protocol to redistribute routes from another VRF, the IS-IS accept policy feature cannot be used. You can only use the IS-IS accept policy for inter-VSN route redistribution between VRFs.

Precedence rules in the same VSN

The following precedence rules apply for IS-IS accept policies used in the same VSN:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply either a default filter for all advertising BEBs or a filter for a specific advertising BEB.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device prefers the **`accept adv-rtr`** filter, which filters based on a specific advertising BEB, over the default filter for all advertising BEBs.
- The device accepts all routes within the same VSN by default. You can apply a route policy to filter or change the characteristics of the route by metric or preference.
- The **`i-sid`** or **`isid-list`** filters are not valid for routes within the same VSN.

Precedence rules for inter-VSN route redistribution

The following precedence rules apply for IS-IS accept policies used for inter-VSN route redistribution:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply filters at a global default level for all BEBs for a specific I-SID or I-SID list, or you can apply filters for a specific advertising BEB for a specific I-SID or I-SID list.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device requires a specific filter to redistribute routes between VSNs through the use of the **`i-sid`** or **`isid-list`** filters.
- The **`i-sid`** filter takes precedence over the **`isid-list`** filter.
- The **`adv-rtr`** filter for a specific advertising BEB takes precedence over a filter with the same **`i-sid`** filter without the **`adv-rtr`** filter.
- The **`i-sid`** or **`isid-list`** filters only apply to routes for inter-VSN route redistribution.
- If multiple **`isid-list`** filters have the same I-SID within the list, the first on the list alphabetically has the higher precedence.

Route preference

The relative value of the route preference among different protocols determines which protocol the device prefers. If multiple protocols are in the routing table, the device prefers the route with the lower value. You can change the value at the protocol level, and you can also change the preference of incoming ISIS routes using the route-map with the ISIS Accept policy filter.

Route metric

Use route-map to change the metric of a route when you accept a remote IS-IS route with IS-IS accept policies.

You can use route-map to change the metric of a route when you redistribute the route from another protocol to IS-IS through the route redistribution mechanism.

You can also configure the route metric with the base **`redistribute`** command without the use of route-map.

For more information on configuration of route-map, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505

IP shortcuts with IP multicast over SPBM

With IP multicast over SPBM for IP Shortcuts, all or a subset of VLANs can exchange IP multicast traffic with the Global Routing Table (GRT). Applications that can use IP Shortcuts with IP multicast over SPBM include: Video surveillance, TV/Video/Ticker/Image distribution, VX-LAN.

With IP multicast over SPBM for IP Shortcuts, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP multicast over SPBM enabled. When you enable IP multicast over SPBM on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must enable **`ip spb-multicast`** on each of the VLANs within the GRT that need to support IP multicast traffic. Enable IP multicast over SPBM on all VLANs to which IP multicast senders and receivers attach. IP multicast over SPBM is typically configured only on BEBs.

Note:

- If you only want to use IP multicast over SPBM, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP multicast over SPBM routing does not depend on unicast routing, which allows you to more easily migrate from a PIM environment to IP multicast over SPBM. You can migrate a PIM environment to IP multicast over SPBM first, and then migrate unicast separately or not at all.
- If no IP interface exists on the VLAN, then you create one. (The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN).
- You do not need to enable IP Shortcuts to support multicast routing in the Layer 3 VSN using SPBM. IPVPN creation and I-SID assignment for the IPVPN is required, but you do not need to enable IPVPN.

I-SIDs

Unlike IP Shortcuts with unicast, a data I-SID is required for IP Shortcuts using IP multicast over SPBM. When the multicast stream reaches the BEB, the BEB assigns a data I-SID to the stream.

The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID.

Unlike Layer 2 VSNs and Layer 3 VSNs, for IP Shortcuts, no scope I-SID value exists to determine the scope of the multicast traffic. Instead the scope is the Global Routing Table, with the information carried in TLV 144.

TLVs

The scope and data I-SID information is propagated through the SPBM cloud using IS-IS Link State Packets (LSPs), which carry TLV updates, and result in the multicast tree creation for that stream. For IP Shortcuts with IP multicast over SPBM, the LSPs carry I-SID information and information about where IP multicast stream senders and receivers exist using TLV 144 and TLV 186.

IGMP

After you configure **`ip spb-multicast enable`** for IP Shortcuts, you cannot enable IGMP, IGMP Snooping, or IGMP proxy on the interface. If you try to enable IGMP Snooping or proxy on any interface where SPBM multicast is enabled, an error message appears for EDM and ACLI.

After you configure **`ip spb-multicast enable`** for IP Shortcuts on each of the VLANs within the GRT that need to support IP multicast traffic, any IGMP functions required for IP multicast over SPBM for IP Shortcuts are automatically enabled. You do not need to configure anything IGMP related.

IP Shortcuts configuration using ACLI

This section provides procedures to configure IP Shortcuts using ACLI.

Configuring SPBM IP shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet switching based on the existing shortest path tree. This allows for end to end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

To enable IP shortcuts on the BEBs, you must configure a circuitless IP address (loopback address) and specify this adress as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct, static, OSPF, RIP, or BGP routes into IS-IS.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, you must create the Customer VLANs, add slots/ports, and add the IP addresses and network masks.

Procedure

1. Enter Loopback Interface Configuration mode

   ```
   enable
   configure terminal
   interface Loopback <1-256>
   ```

2. Configure a CLIP interface to use as the source address for SPBM IP shortcuts:

   ```
   ip address [<1-256>] <A.B.C.D/X>
   ```

3. Exit to Global Configuration mode:

   ```
   exit
   ```

4. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

5. Specify the CLIP interface as the source address for SPBM IP shortcuts:

   ```
   ip-source-address <A.B.C.D>
   ```

6. Configure SPBM IP shortcuts:

   ```
   spbm <1-100> ip enable
   ```

7. Display the status of SPBM IP shortcuts on the switch:

   ```
   show isis spbm
   ```

8. Identify routes on the local switch to be announced into the SPBM network:

   ```
   redistribute {bgp | direct | ospf | rip | static}
   ```

9. Enable routes to be announced into the SPBM network

   ```
   redistribute {bgp | direct | ospf | rip | static} enable
   ```

10. If you want to delete the configuration, use the no option:

    ```
    no redistribute {bgp | direct | ospf | rip | static}
    no redistribute {bgp | direct | ospf | rip | static} enable
    ```

11. Exit to Global Configuration mode:

    ```
    exit
    ```

12. Apply the configured redistribution:

    ```
    isis apply redistribute {direct | bgp | ospf | rip | static}
    ```

13. Redistribute IS-IS routes into other routing protocols using the following steps.
14. Enter Router Configuration mode for the routing protocol:

```
router { bgp | ospf | rip}
```

15. Configure redistribution of IS-IS routes:

```
redistribute isis
```

16. Enable redistribution of IS-IS routes

```
redistribute isis enable
```

17. Exit to Global Configuration mode:

```
exit
```

18. Apply the configured redistribution:

```
ip {bgp | ospf | rip} apply redistribute isis
```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# interface loopback 1
VSP-9012:1(config-if)# ip address 10.0.0.2/8
VSP-9012:1(config-if)# exit
VSP-9012:1(config)# router isis
VSP-9012:1(config-isis)# ip-source-address 10.0.0.2
VSP-9012:1(config-isis)# spbm 1 ip enable
VSP-9012:1(config-isis)# show isis spbm
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM       B-VID      PRIMARY    NICK     LSDB     IP       MULTICAST
INSTANCE              VLAN       NAME     TRAP
--------------------------------------------------------------------------------
1         1000,2000   1000       1.11.16  disable  enable   disable

================================================================================
                              ISIS SPBM SMLT Info
================================================================================
SPBM       SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
--------------------------------------------------------------------------------
1          secondary            00:14:c7:e1:33:e0     0018.b0bb.b3df


--------------------------------------------------------------------------------
 Total Num of SPBM instances: 1
--------------------------------------------------------------------------------
VSP-9012:1(config-isis)# redistribute rip
VSP-9012:1(config-isis)# redistribute rip enable
VSP-9012:1(config-isis)# exit
```

```
VSP-9012:1(config)# isis apply redistribute rip
```

Variable definitions

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

Variable	Value
<1–256>	Specifies an interface ID value. This value is optional.
<A.B.C.D/X>	Specifies an IP address and subnet mask. Use the no option to delete the specified IP address.
<A.B.C.D>	Specifies an IP address. Use the no option to delete the specified IP address.

Use the data in the following table to use the **ip-source-address** command.

Variable	Value
<A.B.C.D>	Specifies the CLIP interface to use as the source address for SPBM IP shortcuts.

Use the data in the following table to use the **spbm** command.

Variable	Value
<1–100> ip enable	Enables or disables SPBM IP shortcut state. The default is disabled. Use the no or default options to disable SPBM IP shortcuts.

Use the data in the following table to use the **redistribute** command.

Variable	Value
{bgp | direct | ospf | rip | static}	Specifies the protocol.
enable	Enables the redistribution of the specified protocol into the SPBM network. The default is disabled. Use the no option to disable the redistribution.
metric *<0–65535>*	Configures the metric (cost) to apply to redistributed routes. The default is 1.
metric-type {external|internal}	Configures the type of route to import into the protocol. The default is internal.
route-map *WORD<1-64>*	Configures the route policy to apply to redistributed routes. Type a name between 0 to 64 characters in length.
subnets {allow|suppress}	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the

Variable	Value
	subnets individually with the learned or configured mask of the subnet. The default is allow.

Use the data in the following table to use the **isis apply redistribute** command.

Variable	Value
{direct | bgp | ospf | rip | static}	Specifies the protocol.

Configuring IS-IS accept policies

Use the following procedure to create and enable IS-IS accept policies to apply to routes from all Backbone Edge Bridges (BEBs) or to all routes from a specific BEB.

Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

IS-IS accept policies are disabled by default.

Note:

- The **isis apply accept [vrf WORD<1-16>]** command can disrupt traffic and cause temporary traffic loss. After you apply **isis apply accept [vrf <1-16>]**, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. Avaya recommends you make all the relevant accept policy changes, and then apply **isis apply accept [vrf WORD<1-16>]** at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless the IS-IS accept policy was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. **(Optional)** If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

   ```
   ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
   ```

3. **(Optional)** Delete an I-SID list:

   ```
   no ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
   ```

4. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

5. Create an IS-IS accept policy instance to apply to all BEBs for a specific I-SID or I-SID list:

   ```
   accept [i-sid <1-16777215>][isid-list WORD <1-32>]
   ```

6. Create an IS-IS accept policy instance to apply to a specific advertising BEB:

   ```
   accept adv-rtr <x.xx.xx> [i-sid <1-16777215>][isid-list WORD <1-32>]
   ```

7. **(Optional)** Delete an IS-IS accept policy instance:

   ```
   no accept [adv-rtr <x.xx.xx>][i-sid <1-16777215>][isid-list WORD
   <1-32>]
   ```

8. Specify an IS-IS route policy to apply to routes from all BEBs:

   ```
   accept route-map WORD<1-64>
   ```

9. Specify an IS-IS route policy to apply to a specific advertising BEB:

   ```
   accept adv-rtr <x.xx.xx>[route-map WORD<1-64>]
   ```

10. **(Optional)** Delete an IS-IS route policy:

    ```
    no accept [adv-rtr <x.xx.xx>] [route-map]
    ```

11. Enable an IS-IS route accept instance:

    ```
    accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list
    WORD<1-32>]
    ```

12. **(Optional)** Disable an IS-IS route accept instance:

    ```
    no accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list
    WORD<1-32>]
    ```

13. Exit IS-IS Router Configuration mode:

    ```
    exit
    ```

    You are in Global Configuration mode.

14. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

    ```
    isis apply accept [vrf WORD <1-16>]
    ```

Example

Configure an IS-IS accept policy:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 101
Switch:1(config-isis)#accept i-sid 101 route-map test
Switch:1(config-isis)#accept i-sid 101 enable
```

```
Switch:1#exit
Switch:1(config)#isis apply accept
```

Variable definitions

Use the data in the following table to use the **ip isid-list** command.

Variable	Value
WORD<1-32>	Creates a name for your I-SID list.
<1-16777215>	Specifies an I-SID number.
list *WORD<1-1024>*	Specifies a list of I-SID values. For example, in the format 1,3,5,8-10.

Use the data in the following table to use the **accept** command.

Variable	Value
adv-rtr *<x.xx.xx>*	Specifies the SPBM nickname for each advertising BEB to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.
enable	Enables an IS-IS accept policy.
i-sid *<1-16777215>*	Specifies an I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies. Use the parameter to apply a filter for routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system can redistribute the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
isid-list *WORD<1-32>*	Specifies the I-SID list name that represents the local or remote Layer 3 VSNs to which the IS-IS accept policy applies. Use the parameter to apply a default filter for all routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
route-map *WORD<1-64>*	Specifies a route policy by name. You must configure the route policy earlier in a separate procedure.

Use the data in the following table to use the **isis apply accept** command.

Variable	Value
vrf *WORD<1-16>*	Specifies a specific VRF instance.

Configuring inter-VRF accept policies on VRFs

Configure IS-IS accept policies on a VRF to use inter-VRF accept policies in the SPB cloud. You can use IS-IS accept policies to redistribute routes between different VRFs, including the global routing table (GRT). First you apply the filter, and then you match the filter based on the I-SID, which represents the Layer 3 VSN context.

Note:

- The **isis apply accept [vrf WORD<1-16>]** command can disrupt traffic and cause temporary traffic loss. After you apply **isis apply accept [vrf<1-16>]**, the command reapplies the accept policies, which deletes all of the IS-IS route,s and adds the IS-IS routes again. Avaya recommends you make all the relevant accept policy changes, and then apply **isis apply accept [vrf WORD<1-16>]** at the end.
- If you use the **accept** command for inter-VRF routes based on the remote I-SID, the device only accepts routes coming from remote BEBs. For instance, if a local Layer 3 VSN exists with the same I-SID, the device does not add the local routes. The assumption is that the device uses existent methods, either through use of another protocol or static configuration, to obtain those routes.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure that a route policy exists.
- Ensure that the VRFs exist.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

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

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

2. **(Optional)** If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

   ```
   ip isid-list WORD<1-32> [<0-16777215>][list WORD<1-1024>]
   ```

3. Create an IS-IS accept policy instance to apply to routes from all Backbone Edge Bridges (BEBs):

```
isis accept [i-sid <0-16777215>][isid-list WORD<1-32>]
```

4. Create an IS-IS accept policy instance to apply to routes for a specific BEB:

   ```
   isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
   WORD<1-32>]
   ```

5. **(Optional)** Delete an IS-IS accept policy instance:

   ```
   no isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
   WORD<1-32>]
   ```

6. Specify an IS-IS route policy to apply to routes from all BEBs:

   ```
   isis accept route-map WORD<1-64>
   ```

7. Specify an IS-IS route policy to apply for a specific BEB:

   ```
   isis accept adv-rtr <x.xx.xx> route-map WORD<1-64>
   ```

8. **(Optional)** Delete an IS-IS route policy:

   ```
   no isis accept [adv-rtr <x.xx.xx>] [route-map]
   ```

9. Enable a configured IS-IS accept policy instance:

   ```
   isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
   WORD<1-32>] [enable]
   ```

10. **(Optional)** Disable a configured IS-IS accept policy instance:

    ```
    no isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
    WORD<1-32>] [enable]
    ```

11. Exit VRF Router Configuration mode:

    ```
    exit
    ```

    You are in Global Configuration mode.

12. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

    ```
    isis apply accept [vrf WORD<1-16>]
    ```

Example

Configure Inter-VRF accept policies on a VRF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf green
Switch:1(router-vrf)#isis accept i-sid 100
Switch:1(router-vrf)#isis accept i-sid 100 enable
Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept vrf green
```

Variable definitions

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

Variable	Value
WORD<1-32>	Creates a name for your I-SID list.
<0-16777215>	Specifies an I-SID value.
list *WORD<1-1024>*	Specifies a list of I-SID values. For example, in the format 1,3,5,8-10.

Use the data in the following table to use the **`isis accept`** command.

Variable	Value
adv-rtr *<x.xx.xx>*	Specifies a specific advertising BEB in which to apply the IS-IS accept policy to routes for a specific advertising BEB. *x.xx.xx* specifies an SPBM nickname. The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter. The system requires an explicit filter to redistribute routes from a particular VSN. If the default global filter or the filter for a specific advertising BEB does not exist, the system does not redistribute the routes from the remote VSN.
enable	Enables the IS-IS accept policy.
i-sid *<0-16777215>*	Configures the I-SID to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).
isid-list *WORD<1-32>*	Configures a list of I-SIDs to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).
route-map *WORD <1-64>*	Specifies a route policy. You must configure a route policy earlier in a separate procedure.

Use the data in the following table to use the **`isis apply accept`** command.

Variable	Value
vrf *WORD<1-16>*	Specifies a specific VRF instance.

Viewing IS-IS accept policy information

Use the following procedure to view IS-IS accept policy information on the switch.

Procedure

1. Display IS-IS accept policy information:

```
show ip isis accept [vrf WORD<1-16>][vrfids WORD<0-512>]
```

2. Display I-SID list information:

   ```
   show ip isid-list [vrf WORD<1-16>][vrfids WORD<0-512>][WORD<1-32>]
   ```

3. Display route information:

   ```
   show ip route [vrf WORD<1-16>]
   ```

   The NH VRF/ISID column displays the I-SID for inter-Virtual Services Network (VSN) routes redistributed with IS-IS accept policies, only if the I-SID redistributed does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays. If the I-SID is 0, the column represents and displays as the GlobalRouter.

   The existing IS-IS routes for Layer 3 VSNs continue to display as the VRF name of the IP VSN.

4. Display the SPBM IP unicast Forwarding Information Base (FIB):

   ```
   show isis spbm ip-unicast-fib [all][id <1-16777215>][spbm-nh-as-mac]
   ```

Example

View IS-IS accept policy information:

Note:

The following example uses slot/port 10/7, which is appropriate for a VSP 9000 configuration. The slot/port configuration for your product may be different.

```
Switch:1#show ip route vrf test
================================================================================

                     IP Route - VRF test
================================================================================
                                      NH                INTER
DST       MASK            NEXT      VRF/ISID    COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------
1.1.1.5   255.255.255.255 1.1.1.5   GlobalRouter 0   0    ISIS 0   IB   200
1.1.1.13  255.255.255.255 VSP13     GRT          10  1000 ISIS 0   IBSV 7
1.1.1.200 255.255.255.255 VSP200    GRT          10  1000 ISIS 0   IBSV 7
5.7.1.0   255.255.255.0   5.7.1.1   -            1   7    LOC  0   DB   0
13.7.1.0  255.255.255.0   VSP13     GlobalRouter 10  1000 ISIS 0   IBSV 7
100.0.0.0 255.255.255.0   100.0.0.1 GlobalRouter 0   100  ISIS 0   IB   200
111.1.1.0 255.255.255.0   111.1.1.1 hub          0   111  ISIS 0   IB   200

Switch:1(config)#show isis spbm ip-unicast-fib
================================================================================================

            SPBM IP-UNICAST FIB ENTRY INFO

================================================================================================
     VRF     DEST                                  OUTGOING  SPBM    PREFIX  IP ROUTE
VRF  ISID    ISID    Destination  NH BEB  VLAN INTERFACE COST    COST    PREFERENCE
------------------------------------------------------------------------------------------------
GRT   -      101     1.1.1.13/32  VSP13   1000  10/7      10     44       7
GRT   -      101     1.1.1.13/32  VSP13   1001  10/7      10     44       7

------------------------------------------------------------------------------------------------

 Total number of SPBM IP-UNICAST FIB entries 2
------------------------------------------------------------------------------------------------
```

```
Switch:1(config)#show ip isid-list test
================================================================================
                                IP ISID LIST
================================================================================
List Name                       I-SID                VRF
--------------------------------------------------------------------------------
test                            1                    GlobalRouter
                                3                    GlobalRouter
                                4                    GlobalRouter
                                5                    GlobalRouter
                                10                   GlobalRouter
                                22                   GlobalRouter

All 6 out of 6 Total Num of Isid Lists displayed

Switch:1(router-vrf)#show ip isid-list vrf red
================================================================================
                                IP ISID LIST red
================================================================================
List Name                       I-SID                VRF
--------------------------------------------------------------------------------
test1                           11                   1
                                12                   1
                                13                   1
                                14                   1
                                15                   1
```

Variable definitions

Use the data in the following table to use the **show ip isis accept** command.

Variable	Value
vrf *WORD<1-16>*	Displays I-SID list information for a particular VRF by name.
vrfids *WORD<0-512>*	Displays I-SID list information for a particular VRF ID.

Use the data in the following table to use the **show ip isid-list** command.

Variable	Value
vrf *WORD<1-16>*	Displays I-SID list information for a particular VRF by name.
vrfids *WORD<0-512>*	Displays I-SID list information for a particular VRF ID.
WORD<1-32>	Displays I-SID list information for a particular I-SID list name.

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

Variable	Value
vrf *WORD<1-16>*	Displays I-SID list information for a particular VRF by name.

Use the data in the following table to use the **show isis spbm ip-unicast-fib** command.

Variable	Value
all	Displays all IS-IS SPBM IP unicast Fowarding Information Base (FIB) information.
id *<1-16777215>*	Displays IS-IS SPBM IP unicast FIB information by I-SID ID.
spbm-nh-as-mac	Displays the next hop B-MAC of the IP unicast FIB entry.

Configuring IP Shortcuts with IP multicast over SPBM

Use this procedure to configure IP multicast over SPBM for IP Shortcuts. The default is disabled.

Note:

- You do not have to enable IP Shortcuts to support IP multicast routing in the GRT using SPBM.
- You cannot enable IP PIM when IP multicast over SPBM is enabled on the VLAN.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable IP multicast over SPBM globally.
- If no IP interface exists on the VLAN, then you create one. (The IP interface must be the same subnet as the IGMP hosts that connect to the VLAN).

About this task

With IP multicast over SPBM for IP Shortcuts, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP multicast over SPBM enabled. When you enable IP multicast over SPBM on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must configure **`ip spb-multicast enable`** on each of the VLANs within GRT that need to support IP multicast traffic. The default is disabled. After you enable SPBM multicast on each of the VLANs within the GRT that need to support IP multicast traffic, any IGMP functions required for IP multicast over SPBM for IP Shortcuts are automatically enabled. You do not need to configure anything IGMP related.

If you only want to use IP multicast over SPBM, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP multicast over SPBM routing does not depend on unicast routing, which allows for you to more easily migrate from a PIM environment to IP multicast over SPBM. You can migrate a PIM environment to IP multicast over SPBM first and then migrate unicast separately or not at all.

The switch only supports IPv4 address with IP multicast over SPBM.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   ```

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

2. Create an IP interface on the VLAN:

   `ip address <A.B.C.D/X>`

3. Enable IP multicast over SPBM:

   `ip spb-multicast enable`

   Note:

   After you configure **`ip spb-multicast enable`** for IP Shortcuts, you cannot enable IGMP, IGMP Snooping, or IGMP proxy on the interface. If you try to enable IGMP Snooping or proxy on any interface where SPBM multicast is enabled, an error message appears for EDM and ACLI.

4. **(Optional)** Disable IP multicast over SPBM:

   `no ip spb-multicast enable`

   `default ip spb-multicast enable`

5. Ensure IP multicast over SPBM for IP Shortcuts is configured properly:

   `show ip igmp interface`

   If routed-spb appears under mode, IP multicast over SPBM for IP Shortcuts is properly enabled on the VLAN.

Example

Enable IP multicast over SPBM for IP Shortcuts:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 500
Switch:1(config-if)#ip address 192.0.2.1 255.255.255.0
Switch:1(config-if)#ip spb-multicast enable
Switch: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
------------------------------------------------------------------------------
V150  125    active 2     2   0.0.0.0    100     0     0     2      10    routed-spb
V2000 125    inact  2     2   0.0.0.0    100     0     0     2      10
```

Variable definitions

Use the data in the following table to use the **`interface vlan`** command.

Variable	Value
<1-4084>	Specifies the VLAN ID.

Use the data in the following table to use the **`interface GigabitEthernet`** command.

Variable	Value
{slot/port[-slot/port][,...]}	Specifies the port number using slot/port notation.

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

Variable	Value
<A.B.C.D/X>	Specifies the address and mask.

Configuring the VRF timeout value

Use this procedure to configure the VRF timeout value. The timeout value ages out the sender when there is no multicast stream on the VRF. The default is 210 seconds.

Note:

You can use this procedure for Layer 3 VSN with IP multicast over SPBM services and IP multicast over SPBM for IP Shortcuts.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable SPBM multicast globally.

Procedure

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

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

2. Configure the timeout value on the VRF:

   ```
   mvpn fwd-cache-timeout(seconds) <10-86400>
   ```

3. **(Optional)** Configure the timeout value to the default value of 210 seconds:

   ```
   no mvpn fwd-cache-timeout
   default mvpn fwd-cache-timeout(seconds)
   ```

Example

Configure the timeout value on the VRF to 500 seconds:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf green
Switch:1(router-vrf)#mvpn fwd-cache-timeout(seconds) 500
```

Variable definitions

Use the data in the following table to use the **router vrf** command.

Variable	Value
WORD<*1–16*>	Specifies the VRF name.

Use the data in the following table to use the **`mvpn fwd-cache-timeout(seconds)`** command.

Variable	Value
<*10–86400*>	Specifies the timeout value. The default is 210 seconds.

Configuring the Global Routing Table timeout value

Use this procedure to configure the timeout value in the GRT. The timeout value ages out the sender when there are no multicast streams coming from the sender for a specified period of time in seconds. The default timeout value is 210 seconds.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable SPBM multicast globally.

Procedure

1. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

2. Configure the SPBM multicast forward-cache timeout:

   ```
   spbm <1-100> multicast fwd-cache-timeout(seconds) <10-86400>
   ```

3. **(Optional)** Configure the SPBM multicast forward-cache timeout to the default value of 210 seconds:

   ```
   default spbm <1-100> multicast fwd-cache-timeout(seconds)
   no spbm <1-100> multicast fwd-cache-timeout(seconds)
   ```

Example

Configure the SPBM multicast forward-cache timeout to 300:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 multicast 1 fwd-cache-timeout(seconds) 300
```

Variable definitions

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

Variable	Value
<1–100>	Specifies the SPBM instance. **Note:** In this release, the switch only supports one instance.
<10–86400>	Specifies the SPBM multicast forward-cache timeout in seconds. The default is 210 seconds.

Viewing IP Shortcuts with IP multicast over SPBM information

Use the following options to display IP Shortcuts with IP multicast over SPBM information to confirm proper configuration.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display all IP multicast over SPBM route information:

   ```
   show isis spbm ip-multicast-route [all]
   ```

3. Display detailed IP multicast over SPBM route information:

   ```
   show isis spbm ip-multicast-route [detail]
   ```

4. Display the IP multicast over SPBM multicast group and source address information:

   ```
   show isis spbm ip-multicast-route [group {A.B.C.D}] [source
   {A.B.C.D}][source-beb WORD<0-255>]
   ```

5. Display summary information for each S, G, V tuple with the corresponding scope, data I-SID, and the host name of the source:

   ```
   show isis spb-mcast-summary [host-name WORD<0-255>][lspid
   <xxxx.xxxx.xxxx.xx-xx>]
   ```

Example

Display IP Shortcuts with IP multicast over SPBM information:

```
Switch:1#show isis spbm ip-multicast-route all
================================================================================
                            SPBM IP-multicast ROUTE INFO ALL
================================================================================
Type VrfName  Vlan Source  Group    VSN-ISID   Data ISID BVLAN Source-BEB

Id
--------------------------------------------------------------------------------
routed GRT    501 192.0.2.1 233.252.0.1  5010  16300001   10    el2
routed GRT    501 192.0.2.1 233.252.0.2  5010  16300002   20    el2
routed GRT    501 192.0.2.1 233.252.0.3  5010  16300003   10    el2
routed GRT    501 192.0.2.1 233.252.0.4  5010  16300004   20    el2
routed GRT    501 192.0.2.1 233.252.0.5  5010  16300005   10    el2
routed GRT    501 192.0.2.1 233.252.0.6  5010  16300006   20    el2
routed GRT    501 192.0.2.1 233.252.0.7  5010  16300007   10    el2
routed GRT    501 192.0.2.1 233.252.0.8  5010  16300008   20    el2
```

```
routed GRT    501 192.0.2.1 233.252.0.9  5010  16300009   10    el2
routed GRT    501 192.0.2.1 233.252.0.10 5010  16300010   20    el2

--------------------------------------------------------------------------------
 Total Number of SPBM IP multicast ROUTE Entries: 10
--------------------------------------------------------------------------------

Switch:1#show isis spbm ip-multicast-route detail
================================================================================
                 SPBM IP-MULTICAST ROUTE INFO
================================================================================
Source        Group     Data ISID BVLAN NNI Rcvrs UNI Rcvrs Source-BEB
--------------------------------------------------------------------------------
192.0.2.10 233.252.0.1     16300001  10     4/3      V604:9/38  el2
192.0.2.10 233.252.0.2     16300002  20     4/2,4/3  V604:9/38  el2
192.0.2.10 233.252.0.3     16300003  10     4/3      V604:9/38  el2
192.0.2.10 233.252.0.4     16300004  20     4/2,4/3  V604:9/38  el2
192.0.2.10 233.252.0.5     16300005  10     4/3      V604:9/38  el2
192.0.2.10 233.252.0.6     16300006  20     4/2,4/3  V604:9/38  el2
192.0.2.10 233.252.0.7     16300007  10     4/3      V604:9/38  el2
192.0.2.10 233.252.0.8     16300008  20     4/2,4/3  V604:9/38  el2
192.0.2.10 233.252.0.9     16300009  10     4/3      V604:9/38  el2
192.0.2.10 233.252.0.10 16300010  20     4/2,4/3  V604:9/38  el2
--------------------------------------------------------------------------------
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
--------------------------------------------------------------------------------
Switch:1# show isis spb-mcast-summary

================================================================================
                            SPB multicast - Summary
================================================================================
SCOPE    SOURCE           GROUP            DATA            LSP  HOST
I-SID    ADDRESS          ADDRESS          I-SID     BVID  FRAG NAME
--------------------------------------------------------------------------------
GRT     192.0.2.1       233.252.0.1     16300001  10    0x0  el2
GRT     192.0.2.1       233.252.0.3     16300003  10    0x0  el2
GRT     192.0.2.1       233.252.0.5     16300005  10    0x0  el2
GRT     192.0.2.1       233.252.0.7     16300007  10    0x0  el2
GRT     192.0.2.1       233.252.0.9     16300009  10    0x0  el2
GRT     192.0.2.1       233.252.0.2     16300002  20    0x0  el2
GRT     192.0.2.1       233.252.0.4     16300004  20    0x0  el2
GRT     192.0.2.1       233.252.0.6     16300006  20    0x0  el2
GRT     192.0.2.1       233.252.0.8     16300008  20    0x0  el2
GRT     192.0.2.1       233.252.0.10    16300010  20    0x0  el2
```

Variable definitions

Use the data in the following table to use the **show isis spbm ip-multicast-route** command.

Variable	Value
all	Displays all IP multicast over SPBM route information.
detail	Displays detailed IP multicast over SPBM route information.

Variable	Value
group *{A.B.C.D}* source *{A.B.C.D}* [source-beb *WORD<0–255>*]	Displays information on the group IP address for the IP multicast over SPBM route. If you select source it will also display the source IP address. Specifies the source BEB name.
vlan	Displays IP multicast over SPBM route information by VLAN.
vrf	Displays IP multicast over SPBM route information by VRF.
vsn-isid	Displays IP multicast over SPBM route information by I-SID.

Use the data in the following table to use the **`show isis spb-mcast-summary`** command.

Variable	Value
host-name *WORD<0–255>*	Displays the SPBM multicast summary for a given host-name.
lspid *<xxxx.xxxx.xxxx.xx-xx>*	Displays the SPBM multicast summary for a given LSP ID.

Job aid

The following table describes the fields for the **`show isis spbm ip-multicast-route all`** command.

Parameter	Description
Type	Specifies the type of interface. The options include: • routed— For IP Shortcuts and Layer 3 VSN. • snoop— For Layer 2 VSN.
VrfName	Specifies the VRF name of the interface.
Vlan Id	Specifies the VLAN ID of the interface.
Source	Specifies the group IP address for the IP multicast over SPBM route.
Group	Specifies the group IP address for the IP multicast over SPBM route.
VSN-ISID	Specifies the GRT for IP Shortcuts with IP multicast over SPBM because IP Shortcuts with IP multicast over SPBM does not use a VSN I-SID.
Data ISID	Specifies the data I-SID for the IP multicast over SPBM route. After a BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is

Parameter	Description
	received on. The data I-SID is a child of the scope or VSN I-SID.
BVLAN	Specifies the B-VLAN for the IP multicast over SPBM route.
Source-BEB	Specifies the source BEB for the IP multicast over SPBM route.

The following table describes the fields for the **`show isis spbm ip-multicast-route detail`** command.

Parameter	Description
Source	Specifies the group IP address for the IP multicast over SPBM route.
Group	Specifies the group IP address for the IP multicast over SPBM route.
Data ISID	Specifies the data I-SID for the IP multicast over SPBM route. After a BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVLAN	Specifies the B-VLAN for the IP multicast over SPBM route.
NNI Rcvrs	Specifies the NNI receivers.
UNI Rcvrs	Specifies the UNI receivers.
Source-BEB	Specifies the source BEB for the IP multicast over SPBM route.

The following table describes the fields for the **`show isis spb-mcast-summary`** command.

Parameter	Description
SCOPE I-SID	Specifies the scope I-SID. Layer 2 VSN and Layer 3 VSN each require a scope I-SID.
SOURCE ADDRESS	Specifies the group IP address for the IP multicast over SPBM route.
GROUP ADDRESS	Specifies the group IP address for the IP multicast over SPBM route.
DATA I-SID	Specifies the data I-SID for the IP multicast route. After a BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16, 512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is

Parameter	Description
	received on. The data I-SID is a child of the scope or VSN I-SID.
BVID	Specifies the Backbone VLAN ID associated with the SPBM instance.
LSP FRAG	Specifies the LSP fragment number.
HOST NAME	Specifies the host name listed in the LSP, or the system name if the host is not configured.

Viewing IGMP information for IP Shortcuts multicast

Use the following commands to display IGMP information.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display information about the interfaces where IGMP is enabled:

   ```
   show ip igmp interface [gigabitethernet {slot/port[-slot/port][,...]}][vlan <1-4084>][vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

   Ensure that the ouput displays `routed-spb` under MODE.

3. Display information about the IGMP cache:

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

4. Display information about the IGMP group:

   ```
   show ip igmp group [count][group {A.B.C.D}][member-subnet default|{A.B.C.D/X}][vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

5. Display information about the IGMP sender:

   ```
   show ip igmp sender [count][group {A.B.C.D}][member-subnet default|{A.B.C.D/X}][vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

Example

Display IGMP information for IP Shortcuts with IP multicast over SPBM:

```
Switch:#enable
Switch:1#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

Switch:1(config)#show ip igmp interface vlan 1
```

```
================================================================================
                                  Vlan Ip Igmp
================================================================================
VLAN QUERY QUERY ROBUST VERSION LAST  PROXY  SNOOP  SSM    FAST   FAST
ID   INTVL MAX                  MEMB  SNOOP  ENABLE SNOOP  LEAVE  LEAVE
           RESP                 QUERY ENABLE        ENABLE ENABLE PORTS
--------------------------------------------------------------------------------
1    125   100   2      2       10    false  false  false  false

VLAN SNOOP   SNOOP          DYNAMIC   COMPATIBILITY EXPLICIT
ID   QUERIER QUERIER        DOWNGRADE MODE          HOST
     ENABLE  ADDRESS        VERSION                 TRACKING
--------------------------------------------------------------------------------
1    false   0.0.0.0         enable    disable       disable

Switch:1# show ip igmp sender

================================================================================
                           IGMP Sender - GlobalRouter
================================================================================
                                          PORT/
GRPADDR         IFINDEX    MEMBER         MLT        STATE
--------------------------------------------------------------------------------
233.252.0.1    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.2    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.3    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.4    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.5    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.6    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.7    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.8    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.9    Vlan 501   192.2.0.1      9/16        NOTFILTERED
233.252.0.10   Vlan 501   192.2.0.1      9/16        NOTFILTERED

10 out of 10 entries displayed
```

```
Switch:1# show ip igmp group

================================================================================
                           IGMP Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
233.252.0.1    V501-9/16       192.2.0.1      204        Dynamic
233.252.0.2    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.3    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.4    V501-9/16       192.2.0.1      207        Dynamic
233.252.0.5    V501-9/16       192.2.0.1      204        Dynamic
233.252.0.6    V501-9/16       192.2.0.1      209        Dynamic
233.252.0.7    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.8    V501-9/16       192.2.0.1      206        Dynamic
233.252.0.9    V501-9/16       192.2.0.1      211        Dynamic
233.252.0.10   V501-9/16       192.2.0.1      207        Dynamic

10 out of 10 group Receivers displayed

Total number of unique groups 10
```

Variable definitions

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

Variable	Value
gigabitethernet *{slot/port [-slot/port][,...]}*	Specifies the GigabitEthernet interface. Use <slot/port [-slot/port][,...]> to specify the slot and port.
vlan *<1-4084>*	Specifies the VLAN.
vrf *WORD<1–16>*	Specifies the VRF by name.
vrfids *WORD<0–512>*	Specifies the VRF by VRF ID.

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

Variable	Value
vrf *WORD<1–16>*	Specifies the VRF by name.
vrfids *WORD<0–512>*	Specifies the VRF by VRF ID.

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

Variable	Value
count	Specifies the number of entries.
group *{A.B.C.D}*	Specifies the group address.
member-subnet *{A.B.C.D/X}*	Specifies the IP address and network mask.
vrf *WORD<1–16>*	Displays the multicast route configuration for a particular VRF by name.
vrfids *WORD<0–512>*	Displays the multicast route configuration for a particular VRF by VRF ID.

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

Variable	Value
count	Specifies the number of entries.
group *{A.B.C.D}*	Specifies the group address.
member-subnet *{A.B.C.D/X}*	Specifies the IP address and network mask.
vrf *WORD<1–16>*	Displays the multicast route configuration for a particular VRF by name.
vrfids *WORD<0–512>*	Displays the multicast route configuration for a particular VRF by VRF ID.

Job aid

The following table describes the fields for the **`show ip igmp interface`** command.

Parameter	Description
IF	Indicates the interface where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.

Parameter	Description
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.
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. If routed-spb displays in the MODE column, then IP multicast over SPBM is enabled on the Layer 3 VSN or for IP shortcuts. If snoop-spb displays in the MODE column, then IGMP is enabled on a VLAN with an associated I-SID (Layer 2 VSN).

The following table describes the fields for the **`show ip igmp interface vlan`** command.

Parameter	Description
VLAN ID	Identifies the VLAN where IGMP is configured.

Parameter	Description
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.
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	Indicates the set of ports that are enabled for fast leave.
VLAN ID	Identifies the VLAN where IGMP is configured.
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 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.

The following table describes the fields for the **`show ip igmp cache`** command.

Parameter	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.
EXPIRATION	Indicates the minimum amount of time that remains before this entry ages out.
V1HOSTIMER	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.

The following table describes the fields for the **`show ip igmp group`** command.

Parameter	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 minimum amount of time that remains before this entry ages out.
TYPE	Indicates whether the entry is learned dynamically or is added statically.

The following table describes the fields for the **`show ip igmp sender`** command.

Parameter	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.

Parameter	Description
STATE	Indicates if a sender exists because of an IGMP access filter. Options include filtered and nonfiltered.

Viewing TLV information for IP Shortcuts multicast

Use the following commands to check TLV information.

For IP Shortcuts multicast, TLV 186 on the BEB where the source is located displays the multicast source and group addresses and have the Tx bit set. Each multicast group has its own unique data I-SID with a value between 16,000,000 to 16,512,000. TLV 144 on the BEB bridge, where the sender is located, has the Tx bit set while on all BEB bridges, where a receiver exists, has the Rx bit set.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display IS-IS Link State Database information by TLV:

   ```
   show isis lsdb tlv <1-186> [sub-tlv <1-3>][detail]
   ```

3. Display IS-IS Link State Database information by Link State Protocol ID:

   ```
   show isis lsdb lspid <xxxx.xxxx.xxxx.xx-xx> tlv <1-186> [sub-tlv <1-3>] [detail]
   ```

Example

Display TLV information:

```
Switch:1# show isis lsdb tlv 186 detail
================================================================================
ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------
Level-1 LspID: 000c.f803.83df.00-06 SeqNum: 0x000002eb Lifetime: 1113
Chksum: 0x7e3b PDU Length: 556
Host_name: Switch
Attributes: IS-Type 1
TLV:186 SPBM IP Multicast:
        GRT ISID
        Metric:0
        IP Source Address: 192.2.0.10
        Group Address : 233.252.0.1
        Data ISID : 16300012
        BVID : 20
        TX : 1
        Route Type : Internal
        GRT ISID
        Metric:0
        IP Source Address: 192.2.0.10
        Group Address : 233.252.0.2
        Data ISID : 16300013
        BVID : 10
        TX : 1
        Route Type : Internal
        GRT ISID
        Metric:0
        IP Source Address: 192.2.0.10
```

```
        Group Address : 233.252.0.3
        Data ISID : 16300014
        BVID : 20
        TX : 1
        Route Type : Internal
        GRT ISID
        Metric:0
        IP Source Address: 192.2.0.10
        Group Address : 233.252.0.4
        Data ISID : 16300015
        BVID : 10
        TX : 1
        Route Type : Internal
        GRT ISID
        Metric:0
        IP Source Address: 192.2.0.10
        Group Address : 233.252.0.5
        Data ISID : 16300016
        BVID : 20
        TX : 1
        Route Type : Internal
        GRT ISID
        Metric:0
        IP Source Address: 192.2.0.10
        Group Address : 233.252.0.6
        Data ISID : 16300017
        BVID : 10
        TX : 1
        Route Type : Internal
        GRT ISID
        Metric:0
        IP Source Address: 192.2.0.10
        Group Address : 233.252.0.7
        Data ISID : 16300018
        BVID : 20
        TX : 1
        Route Type : Internal
```

Variable definitions

Use the data in the following table to use the **`show isis lsdb`** command.

Variable	Value
detail	Displays detailed information about the IS-IS Link State database.
level *{l1, l2, l12}*	Displays information on the IS-IS level. The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled in the current release.
local	Displays information on the local LSDB.
lspid <*xxxx.xxxx.xxxx.xx-xx*>	Specifies information about the IS-IS Link State database by LSP ID.
sub-tlv <*1–3*>	Specifies information about the IS-IS Link State database by sub-TLV.
sysid <*xxxx.xxxx.xxxx*>	Specifies information about the IS-IS Link State database by System ID.

Variable	Value
tlv <*1–186*>	Specifies information about the IS-IS Link State database by TLV.

Job aid

The following table describes the fields for the **`show isis lsdb tlv`** command.

Parameter	Description
LSP ID	Indicates the LSP ID assigned to external IS-IS routing devices.
LEVEL	Indicates the level of the external router: l1, l2, or l12.
LIFETIME	Indicates the maximum age of the LSP. If the max-lsp-gen-interval is set to 900 (default) then the lifetime value begins to count down from 1200 seconds and updates after 300 seconds if connectivity remains. If the timer counts down to zero, the counter adds on an additional 60 seconds, then the LSP for that router is lost.
SEQNUM	Indicates the LSP sequence number. This number changes each time the LSP is updated.
CKSUM	Indicates the LSP checksum. This is an error checking mechanism used to verify the validity of the IP packet.
HOST-NAME	Specifies the host-name.

IP Shortcuts configuration using EDM

This section provides procedures to configure IP Shortcuts using Enterprise Device Manager (EDM).

Configuring SPBM IP shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet switching based on the existing shortest path tree. This allows for end to end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

To enable IP shortcuts on the BEBs, you must configure a circuitless IP address (loopback address) and specify this adress as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct, static, OSPF, RIP, or BGP routes into IS-IS.

After you have configured the SPBM infrastructure, you can enable SPBM IP shortcuts to advertise IP routes across the SPBM network using the following procedure.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, you must create the Customer VLANs, add slots/ports, and add the IP addresses and network masks.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Circuitless IP** tab.
4. Click **Insert**.
5. In the **Interface** box, assign a CLIP interface number.
6. In the **Ip Address** box, type the IP address.
7. In the **Net Mask** box, type the network mask address.
8. Click **Insert**.
9. In the navigation tree, expand the following folders:**Configuration** > **IS-IS**.
10. Click **IS-IS**.
11. From the **Globals** tab, in the **IpSourceAddress** field, specify the CLIP interface to use as the source address for SPBM IP shortcuts.
12. Click **Apply**.
13. In the navigation tree, expand the following folders:**Configuration** > **IS-IS** > **SPBM**.
14. Click the **SPBM** tab.
15. In the **IpShortcut** field double-click and select **enable**.
16. Click **Apply**.
17. In the navigation tree, expand the following folders:**Configuration** > **IP**.
18. Click **Policy**.
19. Click the **Route Redistribution** tab.
20. Click **Insert** to identify routes on the local switch to be announced into the SPBM network.
21. Using the fields provided, specify the source protocols to redistribute into IS-IS. In the **Protocol** field, ensure to specify **isis** as the destination protocol.
22. Click **Insert**.

Configuring IS-IS redistribution

Use this procedure to configure IS-IS redistribution. In the Virtual Routing and Forwarding (VRF), just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IP VPN IP reachability information, the routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IS-IS**.
3. Click the **Redistribute** tab.
4. Click **Insert**.
5. Complete the fields as required.
6. Click **Insert**.

IS-IS Redistribute field descriptions

Use the data in the following table to configure the **IS-IS Redistribute** tab.

Name	Description
DstVrfId	Specifies the destination Virtual Routing and Forwarding (VRF) ID used in the redistribution.
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrfId	Specifies the source VRF ID used in the redistribution. For IS-IS, the source VRF ID must be the same as the destination VRF ID.
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables or disables a redistribution entry. The default is disable.
RoutePolicy	Specifies the route policy to be used for the detailed redistribution of external routes from a specified source into the IS-IS domain.
Metric	Specifies the metric for the redistributed route. The value can be a range between 0 to 65535. The default value is 0. Avaya recommends that you use a value that is consistent with the destination protocol.
MetricType	Specifies the metric type. Specifies a type1 or a type2 metric. For metric type1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type2, the cost of the external routes is equal to the external cost alone. The default is type2.

Name	Description
Subnets	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.

Applying IS-IS accept policies globally

Apply IS-IS accept policies globally. Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

After you apply the IS-IS accept filters, the device removes and re-adds all routes with updated filters.

IS-IS accept policies are disabled by default.

Note:

- After you apply IS-IS accept policies globally the application can disrupt traffic and cause temporary traffic loss. After you configure the IS-IS accept policies value to **Apply**, the device reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. Avaya recommends you make all the relevant accept policy changes, and then apply IS-IS accept policies globally at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Ensure the IP IS-IS filter exists.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Global** tab.
4. Select a name from the list or enter name in the **DefaultPolicyName** field to specify the route policy name for the default filter.
5. Select **Apply** to apply the default policy.

Accept Global field descriptions

Use the data in the following table to configure the **Accept Global** tab.

Name	Description
DefaultPolicyName	Specifies the route policy name for the default filter.

Name	Description
Apply	Applies the default policy when you configure the field to apply. The device only activates the default policy if the route map (the default policy name) has a value. If you do not select apply, the device takes no action. The GRT always returns no action.

Configuring an IS-IS accept policy for a specific advertising BEB

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB). Specify the SPBM nickname and the IS-IS accept policy name to allow you to apply the IS-IS accept policy.

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Nick Name** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Nick Name field descriptions

Use the data in the following table to configure the **Accept Nick Name** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific

Name	Description
	advertising BEB is present the device applies the specific filter.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy. You must configure a policy earlier in a separate procedure.

Configuring an IS-IS accept policy to apply for a specific I-SID

Configure an IS-IS accept policy for a specific I-SID number to represent a local or remote Layer 3 VSN, which allows the system to redistribute the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Isid** tab.
4. Click **Insert**.
5. In the **Isid** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Isid field descriptions

Use the data in the following table to configure the **Accept Isid** tab.

Name	Description
Isid	Configures a specific I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies.

Name	Description
	Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
Enable	Enables or disables the I-SID entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies the route map name. You must configure a policy earlier in a separate procedure.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID

Configures a specific advertising Backbone Edge Bridge (BEB) with a specific I-SID to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Isid** field, specify an I-SID number.
7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick-Name Isid descriptions

Use the date in the following table to configure the **Accept Nick-Name Isid** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB.
Isid	Specifies an I-SID used to filter. The value 0 is used for the Global Router.
Enable	Enables or disables the I-SID entry. The default is disabled.
PolicyName	Specifies the route policy name. You must configure a policy earlier in a separate procedure.

Configuring an I-SID list for an IS-IS accept policy

Configures a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. After you create the list of I-SID numbers, you must then create, configure, and enable the IS-IS accept policy.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify a name for the I-SID list.
6. Select **Isid** or **Isid-List**.
7. Specify an I-SID number or a list of I-SID numbers.
8. Click **Insert**.

Isid-List field descriptions

Use the data in the following table to configure the **Isid-List** tab.

Name	Description
Name	Specifies the name of the I-SID list.
Isid or **Isid-List**	Specifies that you either want to add a particular I-SID or a list of I-SID numbers.
Isid	Specifies a particular I-SID number or a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies.

Name	Description
	An I-SID value of 0 represents the global routing table (GRT).

Configuring an IS-IS accept policy for a specific I-SID list

Configure an IS-IS accept policy for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify the I-SID list name.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Isid–List field descriptions

Use the data in the following table to configure **Accept Isid-List** tab.

Name	Description
Name	Specifies the name of I-SID list.
Enable	Enables or disables the I-SID list entry. The value must be enabled to filter. The default is disabled.
PolicyName	Specifies the route policy name.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB) for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

About this task

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid-List** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Name** field, specify an I-SID list name.
7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick–Name Isid-List field descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid-List** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter is present the device applies the specific filter.
Name	Specifies the name of the I-SID list used to filter.

Name	Description
Enable	Enables or disables the SPBM nickanme advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy name.

Configuring IP multicast over SPBM on a VLAN

Use this procedure to enable IP multicast over SPBM on each of the VLANs within the GRT that need to support IP multicast traffic. The default is disabled.

To configure a VRF with IP multicast over SPBM, see Configuring IP multicast over SPBM on a Layer 3 VSN on page 275.

Note:

- You do not have to enable IP Shortcuts to support IP multicast routing in the GRT using SPBM.
- You cannot enable IP PIM when IP multicast over SPBM is enabled on the VLAN.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable IP multicast over SPBM globally.
- If there is no IP interface on the VLAN, then you create one. (The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN).

About this task

With IP multicast over SPBM for IP Shortcuts, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP multicast over SPBM enabled. When you enable IP multicast over SPBM on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must enable IP multicast over SPBM on each of the VLANs within the GRT that need to support IP multicast traffic. After you enable IP multicast over SPBM on each of the VLANs within the GRT that need to support IP multicast traffic, any IGMP functions required for IP multicast over SPBM for IP Shortcuts are automatically enabled. You do not need to configure anything IGMP related.

If you only want to use IP multicast over SPBM, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP multicast over SPBM routing does not depend on unicast routing (for IP Shortcuts using the GRT). This allows for you to more easily migrate from a PIM environment to IP multicast over SPBM. You can migrate a PIM environment to IP multicast over SPBM first and then migrate unicast separately or not at all.

The switch only supports IPv4 address with IP multicast over SPBM.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.

3. Choose a VLAN, and then click the **IP** button.
4. Click the **SPB Multicast** tab.

   Note:

   After you enable SPBM multicast, you cannot enable IGMP, IGMP Snooping, or IGMP proxy on the interface. If you try to enable IGMP Snooping or proxy on any interface where SPBM multicast is enabled, an error message appears for EDM and ACLI.
5. Click **Enable**.
6. Click **Apply**.

Configuring IP multicast over SPBM on a brouter port for IP Shortcuts

Use this procedure to enable IP multicast over SPBM on a brouter port IP interface. The default is enabled.

To configure a brouter port for a VRF with IP multicast over SPBM, see Configuring IP multicast over SPBM on a brouter port for L3 VSN on page 276.

Note:

- You do not have to enable IP Shortcuts to support IP multicast routing in the GRT using SPBM.
- You cannot enable IP PIM when IP multicast over SPBM is enabled on the VLAN.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable SPBM multicast globally.
- If there is no IP interface on the VLAN, then you create one. (The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN).

About this task

With IP multicast over SPBM for IP Shortcuts, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP multicast over SPBM enabled. When you enable IP multicast over SPBM on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must enable IP multicast over SPBM on each of the VLANs within the GRT that need to support IP multicast traffic. After you enable IP multicast over SPBM on each of the VLANs within the GRT that need to support IP multicast traffic, any IGMP functions required for IP multicast over SPBM for IP Shortcuts are automatically enabled.

If you only want to use IP multicast over SPBM, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP multicast over SPBM routing does not depend on unicast routing, which allows for you to more easily migrate from a PIM environment to multicast over SPBM. You can migrate a PIM environment to IP multicast over SPBM first, and then migrate unicast separately or not at all.

The switch only supports IPv4 address with IP multicast over SPBM.

Procedure

1. Select an enabled port on the Physical Device View.
2. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **SPB Multicast** tab.
5. Click **Enable**.
6. Click **Apply**.

Viewing the IGMP interface table

Use the Interface tab to view the IGMP interface table. When an interface does not use an IP address, it does not appear in the IGMP table.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IGMP**.
3. Click the **Interface** tab.

Configuring the Global Routing Table timeout value

Use this procedure to configure the timeout value in the GRT. The timeout value ages out the sender when there are no multicast streams coming from the sender for a specified period of time. The default timeout value is 210 seconds.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable IP multicast over SPBM globally.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **IS-IS** > **SPBM**.
2. Click the **SPBM** tab.
3. Modify the **McastFwdCacheTimeout** value.
4. Click **Apply**.

IP Shortcuts configuration example

This section provides a configuration example for IP Shortcuts.

IP Shortcuts SPBM configuration example

The following figure shows a sample IP Shortcuts over SPBM deployment.

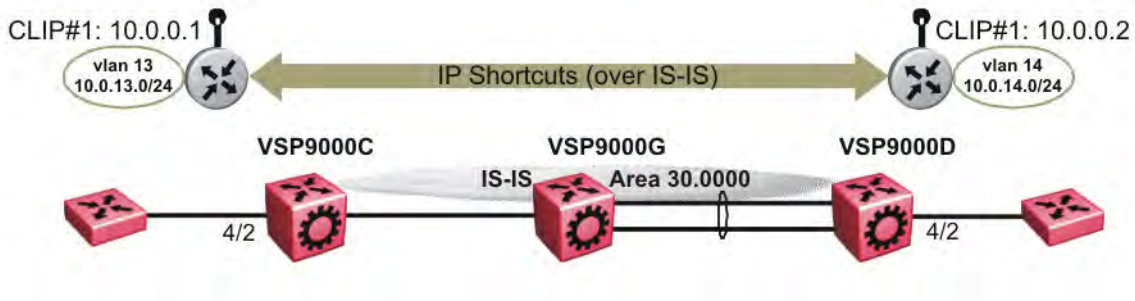

Figure 22: SPBM IP Shortcuts

The following sections show the steps required to configure the SPBM IP Shortcuts parameters in this example. You must first configure basic SPBM and IS-IS infrastructure. For more information, see SPBM configuration examples on page 114.

Note the following:

- IP IS-IS redistribution needs to be configured to inject IP shortcuts routes into IS-IS. The one exception is the circuitless IP address configured as the IS-IS ip-source-address. This address is automatically advertised without the need for a redistribution rule.
- In the displayed configuration, only direct routes are injected (the same configuration is possible for RIP, OSPF, BGP, and static routes).
- No IP address needs to be configured on VSP9000G.

The following sections show the steps required to configure the SPBM IP Shortcuts parameters in this example.

VSP9000C

```
CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter

interface loopback 1
ip address 1 10.0.0.1/255.255.255.0
exit

ISIS CONFIGURATION

router isis
ip-source-address 10.0.0.1

ISIS SPBM CONFIGURATION

spbm 1 ip enable
exit

VLAN CONFIGURATION

vlan create 13 type port-mstprstp 1
vlan members 13 4/2 portmember
interface Vlan 13
ip address 10.0.13.1 255.255.255.0
exit
```

```
IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router isis
redistribute direct
redistribute direct metric 1
redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct
```

VSP9000D

```
CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter

interface loopback 1
ip address 1 10.0.0.2/255.255.255.0
exit

ISIS CONFIGURATION

router isis
ip-source-address 10.0.0.2

ISIS SPBM CONFIGURATION

spbm 1 ip enable
exit

VLAN CONFIGURATION

vlan create 14 type port-mstprstp 1
vlan member add 14 4/2
interface Vlan 14
ip address 10.0.14.1 255.255.255.0
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router isis
redistribute direct
redistribute direct metric 1
redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct
```

Verifying operation — VSP9000C

```
VSP9000C:1# show isis spbm ip-unicast-fib
================================================================================================
        VRF     DEST                                OUTGOING  SPBM      PREFIX IP ROUTE
VRF     ISID    ISID    Destination  NH BEB  VLAN INTERFACE COST      COST   PREFERENCE
------------------------------------------------------------------------------------------------
GRT      -       -     10.0.0.2/32  VSP9000D  4000     4/30     20     1      7
GRT      -       -     10.0.14.1/24 VSP9000D  4000     4/30     20     1      7
------------------------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 2
--------------------------------------------------------------------------------
```

```
VSP9000C:1# show ip route
================================================================================
                        IP Route - GlobalRouter
================================================================================
                                      NH               INTER
DST         MASK           NEXT       VRF/ISID    COST  FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.0.0.1    255.255.255.255 10.0.0.1      -          1     0   LOC   0  DB   0
10.0.0.2    255.255.255.255 VSP9000D   GlobalRouter 20    4000 ISIS  0  IBS  7
10.0.13.1   255.255.255.0   10.0.13.1     -          1    13   LOC   0  DB   0
10.0.14.1   255.255.255.0   VSP9000D   GlobalRouter 20    4000 ISIS  0  IBS  7

4 out of 4 Total Num of Route Entries, 4 Total Num of Dest Networks displayed.
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Rout
e,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

Verifying operation — VSP9000D

```
VSP9000D:1# show isis spbm ip-unicast-fib
================================================================================
     VRF     DEST                              OUTGOING  SPBM      PREFIX IP ROUTE
VRF  ISID    ISID    Destination  NH BEB  VLAN INTERFACE COST      COST   PREFERENCE
--------------------------------------------------------------------------------
GRT   -      -      10.0.0.1/32  VSP9000C  4000   4/20       20     1      7
GRT   -      -      10.0.13.1/24 VSP9000C  4000   4/20       20     1      7
--------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 2
--------------------------------------------------------------------------------

VSP9000D:1# show ip route
================================================================================
                           IP Route - GlobalRouter
================================================================================
                                     NH              INTER
DST         MASK           NEXT      VRF/ISID   COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.0.0.1    255.255.255.255 VSP9000C  GlobalRouter  20   4000 ISIS 0   IBS  7
10.0.0.2    255.255.255.255 10.0.0.2   -            1    0    LOC  0   DB   0
10.0.13.1   255.255.255.0   VSP9000C  GlobalRouter  20   4000 ISIS 0   IBS  7
10.0.14.1   255.255.255.0   10.0.14.1  -            1    14   LOC  0   DB   0

4 out of 4 Total Num of Route Entries, 4 Total Num of Dest Networks displayed.
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Rout
e,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

IP Shortcuts with IP multicast over SPBM configuration example

The example below shows the configuration steps to enable IP multicast over SPBM support on VLANs 10 and 11 that are part of the GRT:

```
ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable

VLAN CONFIGURATION - PHASE I
```

```
interface vlan 500
ip address 192.0.2.1 255.255.255.0
ip spb-multicast enable
exit

interface vlan 501
ip address 192.0.2.2 255.255.255.0
ip spb-multicast enable
exit
```

Layer 3 VSN configuration

This section provides concepts and procedures to configure Layer 3 Virtual Services Network (VSNs).

Layer 3 VSN configuration fundamentals

This section provides fundamental concepts on Layer 3 VSN.

SPBM Layer 3 VSN

The SPBM Layer 3 VSN feature is a mechanism to provide IP connectivity over SPBM for VRFs. SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF.

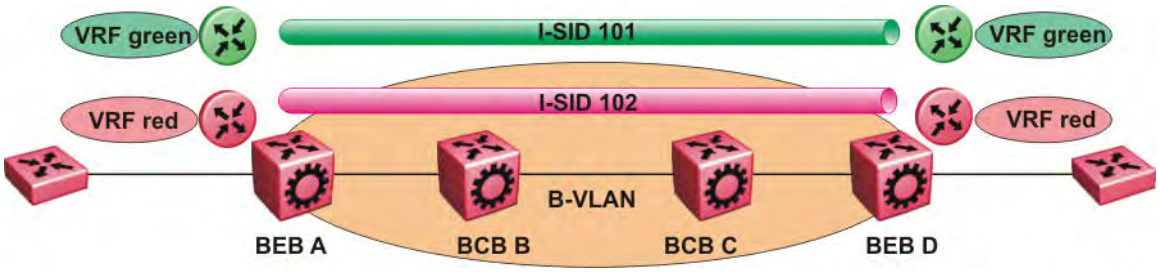

Figure 23: SPBM Layer 3 VSN

In the preceding figure, the BEBs are connected over the SPBM cloud running IS-IS. VRF red and green are configured on the BEBs. VRF red on BEB A has to send and receive routes from VRF red on BEB D. Similar operations are required for VRF green on BEB A and BEB D.

IS-IS TLV 184 is used to advertise SPBM Layer 3 VSN route information across the SPBM cloud. To associate advertised routes with the appropriate VRF, each VRF is associated with an I-SID. All VRFs in the network that share the same I-SID participate in the same VSN.

In this example, I-SID 101 is associated with VRF green and I-SID 102 is associated with VRF red. The I-SID is used to tie the advertised routes to a particular VRF. This identifier has to be the same on all edge nodes for a particular VRF, and has to be unique across all the VRFs on the same node

When IS-IS receives an update from an edge node, it looks for the Layer 3 VSN TLV, and if one exists, it looks at the I-SID identifier. If that identifier is mapped to a local VRF, it extracts the IP routes and adds them to the RTM of that VRF; otherwise the TLV is ignored.

With SPBM Layer 3 VSN, the packet forwarding works in a similar fashion as the IP Shortcuts on the Global Router, with the difference that the encapsulation includes the I-SID to identify the VRF that the packet belongs to. The following figure shows the packet forwarding for VRF red.

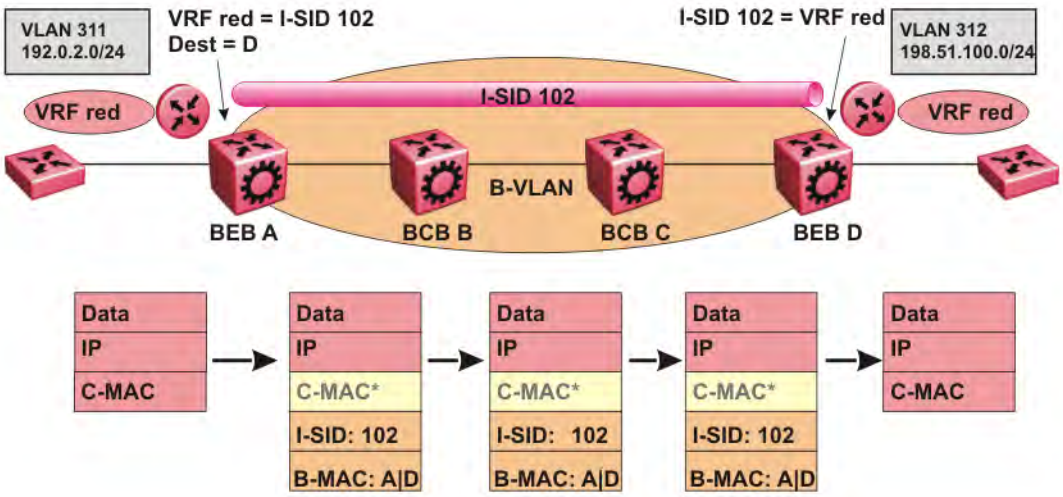

Figure 24: Packet forwarding in SPBM Layer 3 VSN

When BEB A receives traffic from VRF red that must be forwarded to the far-end location, it performs a lookup and determines that VRF red is associated with I-SID 102 and that BEB D is the destination for I-SID 102. BEB A then encapsulates the IP data into a new B-MAC header, using destination B-MAC: D.

Note:

With SPBM Layer 3 VSN, the CMAC header is all null. This header does not have any significance in the backbone. It is included to maintain the same 802.1ah format for ease of implementation.

At BEB D, the node strips off the B-MAC encapsulation, and performs a lookup to determine the destination for traffic with I-SID 102. After identifying the destination as VRF red, the node forwards the packet to the destination VRF.

IS-IS redistribution policies

In the VRF, just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IP VPN IP reachability information, the routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

For each VRF, the next-hop for the installed VPN routes is the node from which the LSPs that carry the IP VPN routes with the same I-SID as the VRF are received. For the IP VPN, the next hop IP

address is the internally generated IP address that corresponds to the nodal BMAC of the next hop that creates the virtual ARP for the node MAC address.

To make IS-IS retrieve the routes from the routing table of a specific VRF for which you enable IP VPN, and advertise the routes to IS-IS peers, use route redistribution and route policies. If you only need to advertise a subset of routes from a specific route type, use route policies, but under the specific VRF context.

The following example shows the configuration to export routes from directly connected interfaces into IS-IS from the SPBM cloud:

```
IP REDISTRIBUTION CONFIGURATION - VRF

router vrf blue
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
```

The following example shows the configuration to distribute IS-IS learned routes into BGP in a VRF context:

```
BGP CONFIGURATION - VRF

router vrf green
ip bgp
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf green
ip bgp redistribute isis
exit

ip bgp redistribute isis enable
```

Interconnection with OSPF or RIP networks

When you connect an SPBM core using Layer 3 VSNs to existing networks that run a routing protocol such as OSPF or RIP, a redundant configuration requires two switches:

- Both routers redistribute IP routes from Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) into IS-IS (IP) and redistribute IS-IS (IP) routes into RIP or OSPF. This can create a routing loop, special precaution need to be taken to prevent this.

The following figure illustrates this configuration.

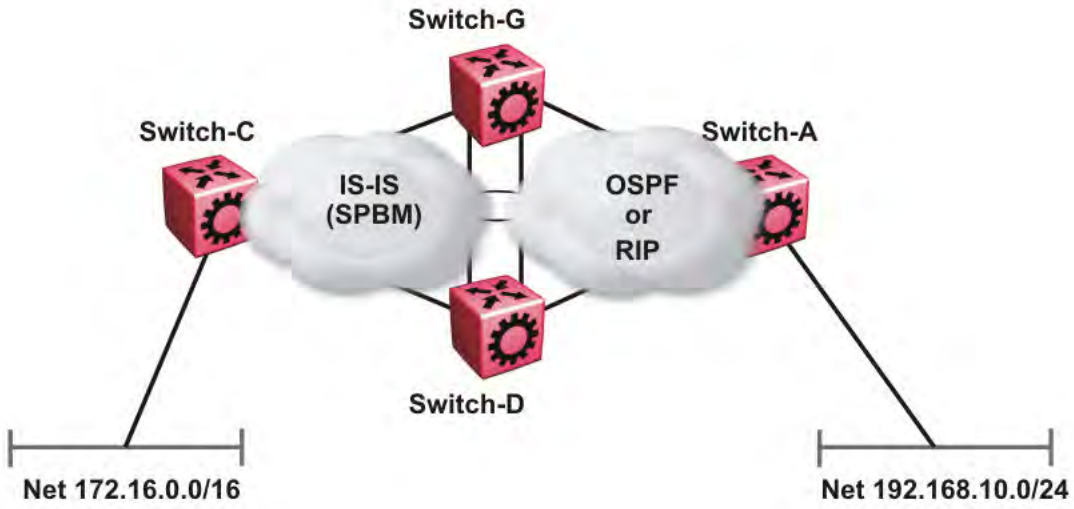

Figure 25: Redundant OSPF or RIP network

In this scenario you must take extra care when redistributing through both switches. By default the preference value for IP routes generated by SPBM-IP (IS-IS) is 7. This is a higher preference than OSPF (20 for intra-area, 25 for inter-area, 120 for ext type1, 125 for ext type2) or RIP (100).

Important:

The lower numerical value determines the higher preference.

In the preceding diagram both nodes (Switch-G and Switch-D) have an OSPF or a RIP route to 192.168.10.0/24 with the next-hop to Switch-A.

As soon as the Switch-G node redistributes that IP route into IS-IS, the Switch-D node learns the same route through IS-IS from Switch-G. (The Switch-G node already has the route through OSPF or RIP). Because IS-IS has a higher preference, Switch-D replaces its 192.168.10.0 OSPF route with an IS-IS one that points at Switch-G as the next-hop. The following figure illustrates this scenario.

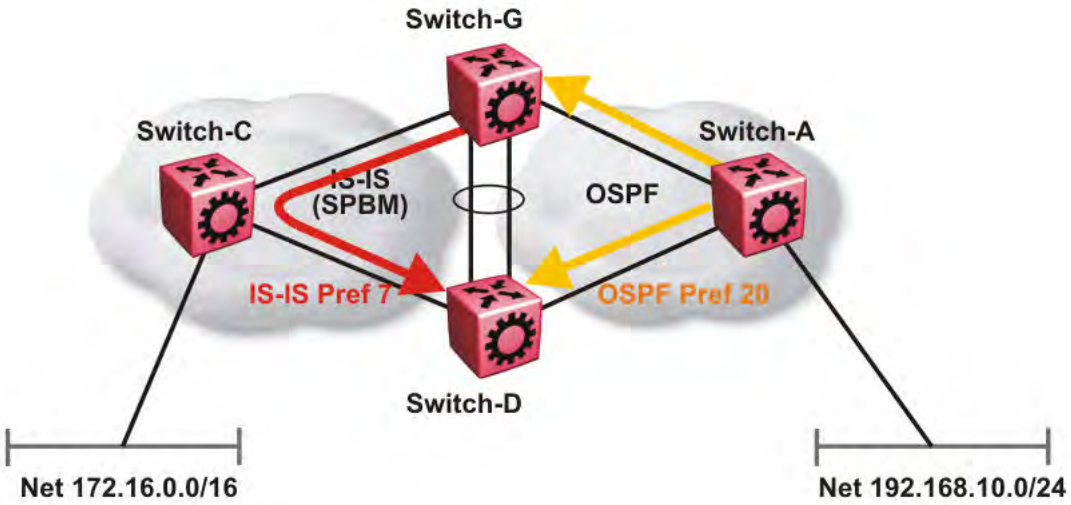

Figure 26: Redistributing routes into IS-IS

This situation is undesirable and you must ensure that the two redistributing nodes (Switch-G and Switch-D) do not accept redistributed routes from each other. With IS-IS accept policies, you can associate an IS-IS accept policy on Switch-D to reject all redistributed IP routes received from Switch-G, and Switch-G to reject all redistribute IP routes from Switch-D.

ISIS Accept configuration used on Switch-G

```
router isis
   redistribute ospf
   redistribute ospf enable
exit
isis apply redistribute ospf

router ospf
   as-boundary-router enable
   redistribute isis
   redistribute isis enable
exit
ip ospf apply redistribute isis

route-map "reject" 1
   no permit
   enable
exit
router isis
   accept adv-rtr <SPB nickname of Switch-D>
   accept adv-rtr <SPB nickname of Switch-D> route-map "reject"
   accept adv-rtr <SPB nickname of Switch-D> enable
exit
isis apply accept
```

ISIS Accept configuration used on Switch-D

```
router isis
```

```
   redistribute ospf
   redistribute ospf enable
exit
isis apply redistribute ospf

router ospf
   as-boundary-router enable
   redistribute isis
   redistribute isis enable
exit
ip ospf apply redistribute isis

route-map "reject" 1
   no permit
   enable
exit
router isis
   accept adv-rtr <SPB nickname of Switch-G>
   accept adv-rtr <SPB nickname of Switch-G> route-map "reject"
   accept adv-rtr <SPB nickname of Switch-G> enable
exit
isis apply accept
```

Note:

Avaya recommends you disable alternative routes by issuing the command **no ip alternative-route** to avoid routing loops on the SMLT Backbone Edge Bridges (BEBs).

In the preceding figure, if Switch-C advertises 25000 IS-IS routes to Switch-G and Switch-D, then both Switch-G and Switch-D install the 25000 routes as IS-IS routes. Since Switch-D and Switch-G have IS-IS to OSPF redistribution enabled, they also learn these 25000 routes as OSPF routes from each other. The OSPF route preference for external (Type1 or Type2) routes normally has a higher numerical value (120 or 125) than the default IS-IS route preference (7), so Switch-D and Switch-G keep the OSPF learned routes as alternative routes.

If Switch-C withdraws its 25000 IS-IS routes, Switch-G and Switch-D remove the IS-IS routes. While the IS-IS routes are removed the routing tables of Switch-G and Switch-D activate the alternative OSPF routes for the same prefix. Since Switch-G and Switch-D also have OSPF to IS-IS redistribution enabled, Switch-C will briefly learn these routes as IS-IS from both Switch-G and Switch-D and this causes a temporary, transient routing loop. This is because the alternative OSPF routes existed because they were redistributed from IS-IS in the first place, before the IS-IS route was withdrawn by Switch-B. To avoid these issues, it is better to simply disable alternative routes on redundant routers which are redistributing the same routes between two different routing protocols. To do this use the **no ip alternative-route** command to disable alternative routes on Switch-G and Switch-D to avoid routing loops.

```
no ip alternative-route
```

The following example demonstrates how to redistribute a default route, instead of all individual IS-IS routes, into an access OSPF or RIP network. In this example a RIP network example is used first then with OSPF. The following figure and sample configuration example illustrates this scenario.

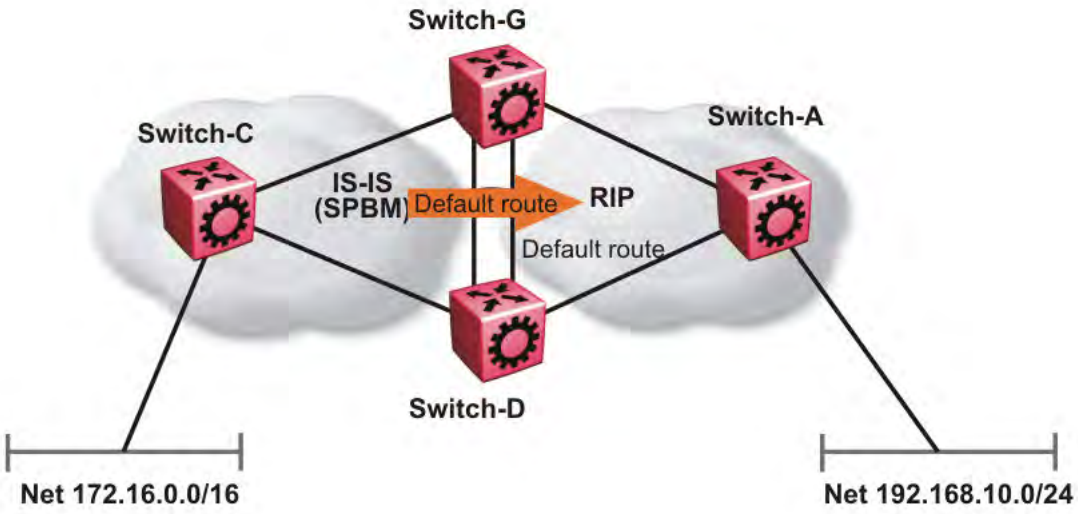

Figure 27: Redistributing routes into RIP

Note:

The following example uses slot/port numbers that are applicable to the VSP 9000, such as 3/2, 3/3, 3/11, and 4/11. If you are using a different platform, such as the VSP 4000 or VSP 8200 the slot/port numbers will be different.

Switch-G

```
enable
configure terminal

IP PREFIX LIST CONFIGURATION
ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION
route-map "inject-default" 1
permit
    set injectlist "default"
enable
exit


IP REDISTRIBUTION CONFIGURATION

router rip
   redistribute isis
   redistribute isis route-map "inject-default"
   redistribute isis enable
exit

RIP PORT CONFIGURATION

interface GigabitEthernet 3/11
```

```
ip rip default-supply enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
```

Switch-A

```
RIP PORT CONFIGURATION

interface gigabitethernet 3/2
ip rip default-listen enable
exit

interface gigabitethernet 3/3
ip rip default-listen enable
exit
```

Switch-D

```
enable
configure terminal

IP PREFIX LIST CONFIGURATION

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION

route-map "inject-default" 1
permit
    set injectlist "default"
enable
exit



IP REDISTRIBUTION CONFIGURATION

router rip
redistribute isis
redistribute isis route-map "inject-default"
redistribute isis enable
exit

RIP PORT CONFIGURATION

interface GigabitEthernet 4/11
ip rip default-supply enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
```

You can control the propagation of the default route on the RIP network so that both Switch-G and Switch-D supply the default route on their relevant interfaces, and not accept it on the same interfaces. Likewise, Switch-A will accept the default route on its interfaces to both Switch-G and Switch-D but it will not supply the default route back to them.

The preceding example where IS-IS IP routes are aggregated into a single default route when redistributed into the RIP network also applies to redistributing IS-IS IP routes into OSPF if that

OSPF network is an access network to an SPBM core. In this case use the following redistribution policy configuration as an example for injecting IS-IS IP routes into OSPF:

```
enable
configure terminal

IP PREFIX LIST CONFIGURATION

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION

route-map "inject-default" 1
permit
set injectlist "default"
enable
exit

IP REDISTRIBUTION CONFIGURATION

router ospf
redistribute isis
redistribute isis route-policy "inject-default"
redistribute isis enable
exit

OSPF CONFIGURATION

router ospf
ip ospf as-boundary-router enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip ospf apply redistribute isis
```

IS-IS accept policies

You can use Intermediate-System-to-Intermediate-System (IS-IS) accept policies to filter incoming IS-IS routes over the SPBM cloud and apply route policies to the incoming IS-IS routes. IS-IS accept policies enable the device to determine whether to add an incoming route to the routing table.

IS-IS accept policy filters

You can filter traffic with IS-IS accept policies by:

- advertising BEB
- I-SID or I-SID list
- route-map

You can use IS-IS accept policies to apply at a global default level for all advertising Backbone Edge Bridges (BEBs) or for a specific advertising BEB.

IS-IS accept policies also allow you to use either a service instance identifier (I-SID) or an I-SID list to filter routes. The switch uses I-SIDs to define Virtual Services Networks (VSNs). I-SIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. IS-IS accept policies can use I-SIDs or I-SID lists to filter the incoming virtualized traffic.

IS-IS accept policies can also apply route policies to determine what incoming traffic to accept into the routing table. With route policies the device can determine which routes to accept into the routing table based on the criteria you configure. You can match on the network or the route metric.

For more information on configuring route policies, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

The following table describes IS-IS accept policy filters.

Filters into	Filter	Description
Global Routing Table (GRT)	accept route-map *WORD<1-64>*	By default, the device accepts all routes into the GRT and VRF routing table. This is the default accept policy.
	accept adv-rtr <*x.xx.xx*>route-map *WORD<1-64>*	The device filters based on the specific advertising BEB defined by the SPBM nickname.
	accept i-sid <*1-16777215*>route-map *WORD<1-64>*	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN.
	accept adv-rtr<*x.xx.xx*> i-sid <*1-16777215*>route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN.
	accept isid-list *WORD<1-32>*route-map *WORD<1-64>*	The device filters based on the list of I-SIDs.
	accept adv-rtr <*x.xx.xx*> isid-list *WORD<1-32>*route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT).
	accept adv-rtr <*x.xx.xx*>route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the route policy.
Virtual Routing and Forwarding (VRF) routing table	isis accept adv-rtr <*x.xx.xx*>	The device filters based on the specific advertising BEB defined by the SPBM nickname.
	isis accept i-sid <*0-16777215*>	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT).
	isis accept adv-rtr <*x.xx.xx*>i-sid <*0-16777215*>	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT).
	isis accept isid-list *WORD<1-32>*	The device filters based on the list of I-SIDs to which the IS-IS accept policy applies. The number 0

Filters into	Filter	Description
		represents the Global Routing Table (GRT).
	isis accept adv-rtr <*x.xx.xx*> isid-list *WORD<1-32>*	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT).
	isis accept route-map *WORD<1-64>*	The device filters based on the route policy.
	isis accept adv-rtr <*x.xx.xx*> route-map *WORD<1-64>*	The device filters based on the specific advertising BEB and the route policy.

IS-IS accept policies for the GRT and VRFs

You can create an IS-IS accept policy for incoming routes for the Global Routing Table (GRT), which accepts routes into the routing table, or for a Virtual Routing and Forwarding (VRF) instance, which accepts incoming routes to the routing table of the VRF.

If you create an IS-IS accept policy on the switch for either the GRT or a VRF that operates at a global default level, the accept policy applies to all routes for all BEBs in the GRT or VRF.

If you create an IS-IS accept policy on the switch for a specific advertising BEB for either the GRT or a VRF, the IS-IS accept policy instance applies for that specific advertising BEB. If you use a more specific filter, the system gives preference to the specific filter over the global default level.

IS-IS accept policies for inter-VRF route redistribution

You can also use the filter mechanism for IS-IS accept policies to redistribute routes between different VRFs, or between a VRF and the GRT. For inter-VRF route redistribution, you match the filter based on the I-SID, which represents the Layer 3 VSN context.

You can apply the filter at the global default level, where the IS-IS accept policy applies to all routes for that I-SID from all BEBs, or at a specific advertising BEB level, where the filter only applies to a specific advertising BEB. The device gives preference to a specific filter for a specific advertising BEB over the global default filter.

For inter-VRF route redistribution, an I-SID value of 0 represents the GRT. For inter-VRF route redistribution between VRFs, the I-SID is the source VRF (or remote VRF).

IS-IS accept policy considerations

Consider the following when you configure IS-IS accept policies:

- The switch does not support IS-IS accept policies for IPv6 addresses for the current release.
- If a VRF uses a different protocol to redistribute routes from another VRF, the IS-IS accept policy feature cannot be used. You can only use the IS-IS accept policy for inter-VSN route redistribution between VRFs.

Precedence rules in the same VSN

The following precedence rules apply for IS-IS accept policies used in the same VSN:

- You can only apply one configured IS-IS accept policy for each route.

- You can apply either a default filter for all advertising BEBs or a filter for a specific advertising BEB.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device prefers the **accept adv-rtr** filter, which filters based on a specific advertising BEB, over the default filter for all advertising BEBs.
- The device accepts all routes within the same VSN by default. You can apply a route policy to filter or change the characteristics of the route by metric or preference.
- The **i-sid** or **isid-list** filters are not valid for routes within the same VSN.

Precedence rules for inter-VSN route redistribution

The following precedence rules apply for IS-IS accept policies used for inter-VSN route redistribution:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply filters at a global default level for all BEBs for a specific I-SID or I-SID list, or you can apply filters for a specific advertising BEB for a specific I-SID or I-SID list.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device requires a specific filter to redistribute routes between VSNs through the use of the **i-sid** or **isid-list** filters.
- The **i-sid** filter takes precedence over the **isid-list** filter.
- The **adv-rtr** filter for a specific advertising BEB takes precedence over a filter with the same **i-sid** filter without the **adv-rtr** filter.
- The **i-sid** or **isid-list** filters only apply to routes for inter-VSN route redistribution.
- If multiple **isid-list** filters have the same I-SID within the list, the first on the list alphabetically has the higher precedence.

Route preference

The relative value of the route preference among different protocols determines which protocol the device prefers. If multiple protocols are in the routing table, the device prefers the route with the lower value. You can change the value at the protocol level, and you can also change the preference of incoming ISIS routes using the route-map with the ISIS Accept policy filter.

Route metric

Use route-map to change the metric of a route when you accept a remote IS-IS route with IS-IS accept policies.

You can use route-map to change the metric of a route when you redistribute the route from another protocol to IS-IS through the route redistribution mechanism.

You can also configure the route metric with the base **redistribute** command without the use of route-map.

For more information on configuration of route-map, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505

Layer 3 VSN with IP multicast over SPBM

IP multicast over SPBM supports Layer 3 VSN functionality where multicast traffic is bridged over the SPBM core infrastructure. Layer 3 VSN using IP multicast over SPBM is helpful when you need complete security and total isolation of data. No one outside of the Layer 3 VSN can join or even see the Layer 3 VSN. Applications that can use Layer 3 VSN with IP multicast over SPBM include: Video surveillance, TV/Video/Ticker/Image Distribution, VX-LAN, Multi-tenant IP multicast.

Configure the Layer 3 VSN (VRF) as a multicast VPN, and then enable IP multicast over SPBM on VRF VLANs to which IP multicast senders and receivers attach. This configuration automatically enables IGMP snooping and proxy on those VLANs. IGMPv2 at the VLAN level is the default setting, with no other configuration required. If you want to use IGMPv3, you must configure IGMPv3.

IP multicast over SPBM is only configured on BEBs.

Note:

- If you only want to use IP multicast over SPBM, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP multicast over SPBM routing does not depend on unicast routing for Layer 3 VSNs using VRFs, which allows you to more easily migrate from a PIM environment to IP multicast over SPBM. You can migrate a PIM environment to IP multicast over SPBM first and then migrate unicast separately or not at all.
- If no IP interface exists on the VLAN, then you create one. (The IP interface must be the same subnet as the IGMP hosts that connect to the VLAN).

With Layer 3 VSN with IP multicast over SPBM, multicast traffic remains in the same Layer 3 VSN across the SPBM cloud. For a Layer 3 VSN, traffic can cross VLAN boundaries but remains confined to the subset of VLANs within the VRF that has IP multicast over SPBM enabled. If a sender transmits a multicast stream to a BEB on a Layer 3 VSN with IP multicast over SPBM enabled, only receivers that are part of the same Layer 3 VSN can receive that stream.

I-SIDs

After a BEB receives IP multicast data from a sender, the BEB allocates a data service instance identifier (I-SID) in the range of 16,000,000 to 16,512,000 for the multicast stream. The stream is identified by the S, G, V tuple, which is the source IP address, the group IP address and the local VLAN the multicast stream is received on. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID.

In the context of Layer 3 VSNs with IP multicast over SPBM, the scope is the I-SID value of the Layer 3 VSN associated with the local VLAN that the IP multicast data was received on.

TLVs

This information is propagated through the SPBM cloud using IS-IS Link State Packets (LSPs), which carry TLV updates, that result in the multicast tree creation for that stream. For Layer 3 VSNs, the LSPs carry I-SID information and information about where IP multicast stream senders and receivers exist using TLV 144 and TLV 185.

IS-IS acts dynamically using the TLV information received from BEBs that connect to the sender and the receivers to create a multicast tree between them.

IGMP

After a BEB receives an IGMP join message from a receiver, the BEB queries the IS-IS database to check if a sender exists for the requested stream within the scope of the receiver. If the requested stream does not exist, the IGMP information is kept, but no further action is taken. If the requested stream exists, the BEB sends an IS-IS TLV update to its neighbors to inform them of the presence of a receiver and this information is propagated through the SPBM cloud.

Layer 3 VSN configuration using ACLI

This section provides a procedure to configure Layer 3 VSNs using Avaya Command Line Interface (ACLI).

Configuring SPBM Layer 3 VSN

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 VSN to advertise IP routes across the SPBM network from one VRF to another using the following procedure.

SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF. In the VRF, just like in the Global Router (VRF 0), the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF on the switch. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.
- You must create the Customer VLANs and add slots/ports.

Procedure

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

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

2. Create an IP VPN instance on the VRF:

   ```
   ipvpn
   ```

3. Configure SPBM Layer 3 VSN:

   ```
   i-sid <0-16777215>
   ```

4. Enable IP VPN on the VRF:

   ```
   ipvpn enable
   ```

   By default, a new IP VPN instance is disabled.

5. Display all IP VPNs:

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

6. Identify routes on the local switch to be announced into the SPBM network:

   ```
   isis redistribute {direct | bgp | ospf | rip | static}
   ```

7. Enable routes on the local switch to be announced into the SPBM network:

   ```
   isis redistribute {direct | bgp | ospf | rip | static} enable
   ```

8. If you want to delete or disable the configuration, use the no option:

   ```
   no isis redistribute {direct | bgp | ospf | rip | static}
   ```

   ```
   no isis redistribute {direct | bgp | ospf | rip | static} enable
   ```

9. Identify other routing protocols to which to redistribute IS-IS routes:

   ```
   ip {bgp | ospf | rip} redistribute isis
   ```

10. Enable IS-IS redistribution to other routing protocols::

    ```
    ip {bgp | ospf | rip} redistribute isis enable
    ```

11. Exit Privileged EXEC mode:

    ```
    exit
    ```

12. Apply the configured redistribution:

    ```
    isis apply redistribute {direct | bgp | ospf | rip | static} vrf
    WORD<1-16>
    ```

    ```
    ip bgp apply redistribute isis vrf WORD<1-16>
    ```

    ```
    ip ospf apply redistribute isis vrf WORD<1-16>
    ```

    ```
    ip rip apply redistribute isis vrf WORD<1-16>
    ```

13. Display the redistribution configuration:

    ```
    show ip isis redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
    ```

Example

Create the IP VPN instance:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf vrfred
Switch:1(config)#ipvpn
Switch:1(config)#i-sid 100
Switch:1(config)#ipvpn enable
Switch:1(config)#show ip ipvpn
          VRF Name             : vrfred
          Ipvpn-state          : enabled
          I-sid                : 100
Switch:1(config)#isis redistribute ospf
Switch:1(config)#isis redistribute ospf enable
Switch:1(config)#isis redistribute ospf enable
Switch:1(config)#end
```

```
Switch:1(config)#isis apply redistribute ospf vrf vrfred
Switch:1(config)#show ip isis redistribute vrf vrfred
================================================================================
                          ISIS Redistribute List - VRF vrfred
================================================================================

SOURCE MET MTYPE      SUBNET   ENABLE LEVEL  RPOLICY
--------------------------------------------------------------------------------
LOC    1   internal   allow    FALSE  l1
```

Variable definitions

Use the data in the following table to configure the **`show ip ipvpn`** command.

Table 3: Variable definitions

Variable	Value
vrf *WORD<1–16>*	Specifies the VRF name.
vrfids *WORD<0–512>*	Specifies the VRF ID.

Use the data in the following table to configure the **`i-sid`** command.

Variable	Value
<0–16777215>	Assigns an I-SID to the VRF being configured. Use the no or default option to remove the I-SID to VRF allocation for this VRF.

Use the data in the following table to configure the **`isis redistribute`** command.

Variable	Value
{direct | bgp | ospf | rip | static}	Specifies the protocol.
enable	Enables the redistribution of the specified protocol into the SPBM network. The default is disabled. Use the no or default options to disable the redistribution.
metric *<0–65535>*	Configures the metric (cost) to apply to redistributed routes. The default is 1.
metric-type {external|internal}	Configures the type of route to import into the protocol. The default is internal.
route-map *WORD<0–64>*	Configures the route policy to apply to redistributed routes. Specifies a name.
subnets {allow|suppress}	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.

Use the data in the following table to configure the **`isis apply redistribute`** command.

Variable	Value
{direct | bgp | ospf | rip | static}	Specifies the protocol.
vrf *WORD<1–16>*	Applies IS-IS redistribute for a particular VRF. Specifies the VRF name.

Configuring IS-IS accept policies

Use the following procedure to create and enable IS-IS accept policies to apply to routes from all Backbone Edge Bridges (BEBs) or to all routes from a specific BEB.

Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

IS-IS accept policies are disabled by default.

Note:

- The **`isis apply accept [vrf WORD<1-16>]`** command can disrupt traffic and cause temporary traffic loss. After you apply **`isis apply accept [vrf <1-16>]`**, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. Avaya recommends you make all the relevant accept policy changes, and then apply **`isis apply accept [vrf WORD<1-16>]`** at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless the IS-IS accept policy was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. **(Optional)** If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

   ```
   ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
   ```

3. **(Optional)** Delete an I-SID list:

   ```
   no ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
   ```

4. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   ```

```
router isis
```

5. Create an IS-IS accept policy instance to apply to all BEBs for a specific I-SID or I-SID list:

   ```
   accept [i-sid <1-16777215>][isid-list WORD <1-32>]
   ```

6. Create an IS-IS accept policy instance to apply to a specific advertising BEB:

   ```
   accept adv-rtr <x.xx.xx> [i-sid <1-16777215>][isid-list WORD <1-32>]
   ```

7. **(Optional)** Delete an IS-IS accept policy instance:

   ```
   no accept [adv-rtr <x.xx.xx>][i-sid <1-16777215>][isid-list WORD
   <1-32>]
   ```

8. Specify an IS-IS route policy to apply to routes from all BEBs:

   ```
   accept route-map WORD<1-64>
   ```

9. Specify an IS-IS route policy to apply to a specific advertising BEB:

   ```
   accept adv-rtr <x.xx.xx>[route-map WORD<1-64>]
   ```

10. **(Optional)** Delete an IS-IS route policy:

    ```
    no accept [adv-rtr <x.xx.xx>] [route-map]
    ```

11. Enable an IS-IS route accept instance:

    ```
    accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list
    WORD<1-32>]
    ```

12. **(Optional)** Disable an IS-IS route accept instance:

    ```
    no accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list
    WORD<1-32>]
    ```

13. Exit IS-IS Router Configuration mode:

    ```
    exit
    ```

    You are in Global Configuration mode.

14. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

    ```
    isis apply accept [vrf WORD <1-16>]
    ```

Example

Configure an IS-IS accept policy:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 101
Switch:1(config-isis)#accept i-sid 101 route-map test
Switch:1(config-isis)#accept i-sid 101 enable
Switch:1#exit
Switch:1(config)#isis apply accept
```

Variable definitions

Use the data in the following table to use the **ip isid-list** command.

Variable	Value
WORD<1-32>	Creates a name for your I-SID list.
<1-16777215>	Specifies an I-SID number.
list *WORD<1-1024>*	Specifies a list of I-SID values. For example, in the format 1,3,5,8-10.

Use the data in the following table to use the **accept** command.

Variable	Value
adv-rtr *<x.xx.xx>*	Specifies the SPBM nickname for each advertising BEB to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.
enable	Enables an IS-IS accept policy.
i-sid *<1-16777215>*	Specifies an I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies. Use the parameter to apply a filter for routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system can redistribute the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
isid-list *WORD<1-32>*	Specifies the I-SID list name that represents the local or remote Layer 3 VSNs to which the IS-IS accept policy applies. Use the parameter to apply a default filter for all routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
route-map *WORD<1-64>*	Specifies a route policy by name. You must configure the route policy earlier in a separate procedure.

Use the data in the following table to use the **isis apply accept** command.

Variable	Value
vrf *WORD<1-16>*	Specifies a specific VRF instance.

Configuring inter-VRF accept policies on VRFs

Configure IS-IS accept policies on a VRF to use inter-VRF accept policies in the SPB cloud. You can use IS-IS accept policies to redistribute routes between different VRFs, including the global routing table (GRT). First you apply the filter, and then you match the filter based on the I-SID, which represents the Layer 3 VSN context.

Note:

- The **`isis apply accept [vrf WORD<1-16>]`** command can disrupt traffic and cause temporary traffic loss. After you apply **`isis apply accept [vrf<1-16>]`**, the command reapplies the accept policies, which deletes all of the IS-IS route,s and adds the IS-IS routes again. Avaya recommends you make all the relevant accept policy changes, and then apply **`isis apply accept [vrf WORD<1-16>]`** at the end.
- If you use the **`accept`** command for inter-VRF routes based on the remote I-SID, the device only accepts routes coming from remote BEBs. For instance, if a local Layer 3 VSN exists with the same I-SID, the device does not add the local routes. The assumption is that the device uses existent methods, either through use of another protocol or static configuration, to obtain those routes.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure that a route policy exists.
- Ensure that the VRFs exist.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

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

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

2. **(Optional)** If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

   ```
   ip isid-list WORD<1-32> [<0-16777215>][list WORD<1-1024>]
   ```

3. Create an IS-IS accept policy instance to apply to routes from all Backbone Edge Bridges (BEBs):

   ```
   isis accept [i-sid <0-16777215>][isid-list WORD<1-32>]
   ```

4. Create an IS-IS accept policy instance to apply to routes for a specific BEB:

   ```
   isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
   WORD<1-32>]
   ```

5. **(Optional)** Delete an IS-IS accept policy instance:

   ```
   no isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
   WORD<1-32>]
   ```

6. Specify an IS-IS route policy to apply to routes from all BEBs:

   ```
   isis accept route-map WORD<1-64>
   ```

7. Specify an IS-IS route policy to apply for a specific BEB:

   ```
   isis accept adv-rtr <x.xx.xx> route-map WORD<1-64>
   ```

8. **(Optional)** Delete an IS-IS route policy:

   ```
   no isis accept [adv-rtr <x.xx.xx>] [route-map]
   ```

9. Enable a configured IS-IS accept policy instance:

   ```
   isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
   WORD<1-32>] [enable]
   ```

10. **(Optional)** Disable a configured IS-IS accept policy instance:

    ```
    no isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list
    WORD<1-32>] [enable]
    ```

11. Exit VRF Router Configuration mode:

    ```
    exit
    ```

    You are in Global Configuration mode.

12. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

    ```
    isis apply accept [vrf WORD<1-16>]
    ```

Example

Configure Inter-VRF accept policies on a VRF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf green
Switch:1(router-vrf)#isis accept i-sid 100
Switch:1(router-vrf)#isis accept i-sid 100 enable
Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept vrf green
```

Variable definitions

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

Variable	Value
WORD<1-32>	Creates a name for your I-SID list.

Variable	Value
<0-16777215>	Specifies an I-SID value.
list *WORD<1-1024>*	Specifies a list of I-SID values. For example, in the format 1,3,5,8-10.

Use the data in the following table to use the **`isis accept`** command.

Variable	Value
adv-rtr *<x.xx.xx>*	Specifies a specific advertising BEB in which to apply the IS-IS accept policy to routes for a specific advertising BEB. *x.xx.xx* specifies an SPBM nickname. The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter. The system requires an explicit filter to redistribute routes from a particular VSN. If the default global filter or the filter for a specific advertising BEB does not exist, the system does not redistribute the routes from the remote VSN.
enable	Enables the IS-IS accept policy.
i-sid *<0-16777215>*	Configures the I-SID to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).
isid-list *WORD<1-32>*	Configures a list of I-SIDs to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).
route-map *WORD <1-64>*	Specifies a route policy. You must configure a route policy earlier in a separate procedure.

Use the data in the following table to use the **`isis apply accept`** command.

Variable	Value
vrf *WORD<1-16>*	Specifies a specific VRF instance.

Viewing IS-IS accept policy information

Use the following procedure to view IS-IS accept policy information on the switch.

Procedure

1. Display IS-IS accept policy information:

   ```
   show ip isis accept [vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

2. Display I-SID list information:

   ```
   show ip isid-list [vrf WORD<1-16>][vrfids WORD<0-512>][WORD<1-32>]
   ```

3. Display route information:

   ```
   show ip route [vrf WORD<1-16>]
   ```

   The NH VRF/ISID column displays the I-SID for inter-Virtual Services Network (VSN) routes redistributed with IS-IS accept policies, only if the I-SID redistributed does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays. If the I-SID is 0, the column represents and displays as the GlobalRouter.

   The existing IS-IS routes for Layer 3 VSNs continue to display as the VRF name of the IP VSN.

4. Display the SPBM IP unicast Forwarding Information Base (FIB):

   ```
   show isis spbm ip-unicast-fib [all][id <1-16777215>][spbm-nh-as-mac]
   ```

Example

View IS-IS accept policy information:

Note:

The following example uses slot/port 10/7, which is appropriate for a VSP 9000 configuration. The slot/port configuration for your product may be different.

```
Switch:1#show ip route vrf test
================================================================================

                     IP Route - VRF test
================================================================================
                                      NH              INTER
DST       MASK            NEXT      VRF/ISID    COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------
1.1.1.5   255.255.255.255 1.1.1.5   GlobalRouter 0   0    ISIS 0   IB   200
1.1.1.13  255.255.255.255 VSP13     GRT          10  1000 ISIS 0   IBSV 7
1.1.1.200 255.255.255.255 VSP200    GRT          10  1000 ISIS 0   IBSV 7
5.7.1.0   255.255.255.0   5.7.1.1   -            1   7    LOC  0   DB   0
13.7.1.0  255.255.255.0   VSP13     GlobalRouter 10  1000 ISIS 0   IBSV 7
100.0.0.0 255.255.255.0   100.0.0.1 GlobalRouter 0   100  ISIS 0   IB   200
111.1.1.0 255.255.255.0   111.1.1.1 hub          0   111  ISIS 0   IB   200

Switch:1(config)#show isis spbm ip-unicast-fib
================================================================================================

            SPBM IP-UNICAST FIB ENTRY INFO

================================================================================================
     VRF     DEST                                 OUTGOING  SPBM    PREFIX  IP ROUTE
VRF  ISID    ISID    Destination  NH BEB  VLAN INTERFACE COST    COST    PREFERENCE
------------------------------------------------------------------------------------------------
GRT   -      101     1.1.1.13/32  VSP13   1000  10/7      10     44       7
GRT   -      101     1.1.1.13/32  VSP13   1001  10/7      10     44       7

------------------------------------------------------------------------------------------------

 Total number of SPBM IP-UNICAST FIB entries 2
------------------------------------------------------------------------------------------------
Switch:1(config)#show ip isid-list test
================================================================================
```

```
                                IP ISID LIST
================================================================================
List Name                       I-SID               VRF
--------------------------------------------------------------------------------
test                            1                   GlobalRouter
                                3                   GlobalRouter
                                4                   GlobalRouter
                                5                   GlobalRouter
                                10                  GlobalRouter
                                22                  GlobalRouter

All 6 out of 6 Total Num of Isid Lists displayed

Switch:1(router-vrf)#show ip isid-list vrf red
================================================================================
                                IP ISID LIST red
================================================================================
List Name                       I-SID               VRF
--------------------------------------------------------------------------------
test1                           11                  1
                                12                  1
                                13                  1
                                14                  1
                                15                  1
```

Variable definitions

Use the data in the following table to use the **`show ip isis accept`** command.

Variable	Value
vrf *WORD<1-16>*	Displays I-SID list information for a particular VRF by name.
vrfids *WORD<0-512>*	Displays I-SID list information for a particular VRF ID.

Use the data in the following table to use the **`show ip isid-list`** command.

Variable	Value
vrf *WORD<1-16>*	Displays I-SID list information for a particular VRF by name.
vrfids *WORD<0-512>*	Displays I-SID list information for a particular VRF ID.
WORD<1-32>	Displays I-SID list information for a particular I-SID list name.

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

Variable	Value
vrf *WORD<1-16>*	Displays I-SID list information for a particular VRF by name.

Use the data in the following table to use the **`show isis spbm ip-unicast-fib`** command.

Variable	Value
all	Displays all IS-IS SPBM IP unicast Fowarding Information Base (FIB) information.
id *<1-16777215>*	Displays IS-IS SPBM IP unicast FIB information by I-SID ID.
spbm-nh-as-mac	Displays the next hop B-MAC of the IP unicast FIB entry.

Configuring Layer 3 VSN with IP multicast over SPBM

Use this procedure to configure IP multicast over SPBM for a Layer 3 VSN.

Configure the Layer 3 VSN (VRF) as a multicast VPN, and then enable IP multicast over SPBM on VRF VLANs to which IP multicast senders and receivers attach. After you enable IP multicast over SPBM on VRF VLANs, snooping and proxy on those VLANs is enabled. IGMPv2 at the VLAN level is the default setting. No configuration is required.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable SPBM multicast globally.
- You must assign an I-SID for the IPVPN.

About this task

With Layer 3 VSN IP multicast over SPBM, multicast traffic remains in the same Layer 3 VSN across the SPBM cloud.

For a Layer 3 VSN, traffic can cross VLAN boundaries but remains confined to the subset of VLANs within the VRF that have **`ip spbm-multicast`** enabled. The default is disabled.

All or a subset of VLANs within a Layer 3 VSN can exchange multicast traffic. The BEB only sends out traffic for a multicast stream on which IGMP joins and reports are received.

In this release, the switch only supports IPv4 multicast traffic.

Note:

You cannot enable IP PIM when IP multicast over SPBM is enabled on the VLAN.

The IP VPN does not need to be enabled for Layer 3 VSN multicast to function.

Procedure

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

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

2. Enable Layer 3 VSN IP multicast over SPBM for a particular VRF:

```
mvpn enable
```

The default is disabled.

3. **(Optional)** If you want to disable Layer 3 VSN IP multicast over SPBM, enter:

   ```
   no mvpn enable
   ```

   ```
   default mvpn enable
   ```

4. Exit to Global Configuration mode:

   ```
   exit
   ```

5. Enter Interface Configuration mode:

   ```
   enable
   ```

   ```
   configure terminal
   ```

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

6. Enable Layer 3 VSN IP multicast over SPBM for a particular VRF:

   ```
   ip spb-multicast enable
   ```

7. **(Optional)** Disable Layer 3 multicast on the VRF:

   ```
   no ip spb-multicast enable
   ```

8. **(Optional)** Enable IGMP version 3:

   ```
   ip igmp snooping
   ```

   ```
   ip igmp ssm-snoop
   ```

   ```
   ip igmp compatibility-mode
   ```

   ```
   ip igmp version 3
   ```

   Note:

   IGMPv2 at the VLAN level is the default setting, with no other configuration required. You only need to use these commands if you use IGMPv3. You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

   For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

9. **(Optional)** Enable the IGMP Layer 2 Querier address:

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

   Note:

   If the SPBM bridge connects to an edge switch, it can be necessary to add an IGMP query address. If you omit adding a query address, the SPB bridge sends IGMP queries

with a source address of 0.0.0.0. Some edge switch models do not accept a query with a source address of 0.0.0.0.

Example

Configure IP multicast over SPBM for a Layer 3 VSN:

```
Switch:>enable
Switch:#configure terminal
Switch:(config)# router vrf green
Switch:(config-vrf)#mvpn enable
Switch:(config)#exit
Switch:(config)#interface vlan 500
Switch:(config-if)#ip spb-multicast enable
```

Variable definitions

Use the data in the following table to use the **`router vrf`** command.

Variable	Value
WORD<1–16>	Specifies the name of the VRF.

Use the data in the following table to use the **`interface vlan`** command.

Variable	Value
<1-4084>	Specifies the VLAN ID in the range of 1-4084. VLAN IDs 1 to 4084 are configurable; VLAN IDs 4085 to 4094 are reserved for internal use. VLAN ID 1 is the default VLAN and cannot be created nor deleted.

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

Variable	Value
GigabitEthernet*{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 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.

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 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).
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).

Variable	Value
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.
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 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>*	Configures multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. The default is 4.

Variable	Value
stream-limit-group {slot/port[-slot/port][,...]} enable max-streams <*0-65535*>	Configures 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 the VRF timeout value

Use this procedure to configure the VRF timeout value. The timeout value ages out the sender when there is no multicast stream on the VRF. The default is 210 seconds.

Note:

You can use this procedure for Layer 3 VSN with IP multicast over SPBM services and IP multicast over SPBM for IP Shortcuts.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable SPBM multicast globally.

Procedure

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

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

2. Configure the timeout value on the VRF:

   ```
   mvpn fwd-cache-timeout(seconds) <10-86400>
   ```

3. **(Optional)** Configure the timeout value to the default value of 210 seconds:

   ```
   no mvpn fwd-cache-timeout
   default mvpn fwd-cache-timeout(seconds)
   ```

Example

Configure the timeout value on the VRF to 500 seconds:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf green
Switch:1(router-vrf)#mvpn fwd-cache-timeout(seconds) 500
```

Variable definitions

Use the data in the following table to use the **router vrf** command.

Variable	Value
WORD<*1–16*>	Specifies the VRF name.

Use the data in the following table to use the **mvpn fwd-cache-timeout(seconds)** command.

Variable	Value
<*10–86400*>	Specifies the timeout value. The default is 210 seconds.

Configuring the Global Routing Table timeout value

Use this procedure to configure the timeout value in the GRT. The timeout value ages out the sender when there are no multicast streams coming from the sender for a specified period of time in seconds. The default timeout value is 210 seconds.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable SPBM multicast globally.

Procedure

1. Enter IS-IS Router Configuration mode:

   ```
   enable
   configure terminal
   router isis
   ```

2. Configure the SPBM multicast forward-cache timeout:

   ```
   spbm <1-100> multicast fwd-cache-timeout(seconds) <10-86400>
   ```

3. **(Optional)** Configure the SPBM multicast forward-cache timeout to the default value of 210 seconds:

   ```
   default spbm <1-100> multicast fwd-cache-timeout(seconds)
   no spbm <1-100> multicast fwd-cache-timeout(seconds)
   ```

Example

Configure the SPBM multicast forward-cache timeout to 300:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 multicast 1 fwd-cache-timeout(seconds) 300
```

Variable definitions

Use the data in the following table to use the **spbm** command.

Variable	Value
<1–100>	Specifies the SPBM instance. **Note:** In this release, the switch only supports one instance.
<10–86400>	Specifies the SPBM multicast forward-cache timeout in seconds. The default is 210 seconds.

Viewing Layer 3 VSN with IP multicast over SPBM information

Use the following options to display Layer 3 VSN with IP multicast over SPBM information to confirm proper configuration.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display all the VRFs that have MVPN enabled and their corresponding forward cache timeout values:

   ```
   show ip vrf mvpn
   ```

3. Display IP multicast over SPBM route information:

   ```
   show isis spbm ip-multicast-route [all][detail]
   ```

4. Display IP multicast over SPBM by group and source address:

   ```
   show isis spbm ip-multicast-route [group {A.B.C.D}][detail][source
   {A.B.C.D}]
   ```

5. Display IP multicast over SPBM route information by VRF:

   ```
   show isis spbm ip-multicast-route [vrf WORD<1-16>] [group {A.B.C.D}]
   ```

6. Display IP multicast over SPBM route information by VLAN:

   ```
   show isis spbm ip-multicast-route [vlan <1-4084>][detail][group
   {A.B.C.D}]
   ```

7. Display IP multicast over SPBM information by VSN I-SID:

   ```
   show isis spbm ip-multicast-route [vsn-isid <1-16777215>][detail]
   [group {A.B.C.D}]
   ```

8. Display summary information for each S, G, V tuple with the corresponding scope, Data I-SID, and the host name of the source:

   ```
   show isis spb-mcast-summary [host-name WORD<0-255>][lspid
   <xxxx.xxxx.xxxx.xx-xx>]
   ```

Example

Display Layer 3 VSN with IP multicast over SPBM information:

```
Switch:1>enable
Switch:1#show ip vrf mvpn

                            Vrf name : green
                                mvpn : enable
                    fwd-cache-timeout : 210

                            Vrf name : 4
                                mvpn : enable
                    fwd-cache-timeout : 210

                            Vrf name : blue
                                mvpn : enable
                    fwd-cache-timeout : 210

Switch:1#show isis spbm ip-multicast-route all
================================================================================
              SPBM IP-multicast ROUTE INFO ALL
================================================================================
Type     VrfName Vlan Source    Group       VSN-ISID    Data ISID BVLAN Source-BEB
            Id
--------------------------------------------------------------------------------
routed    GRT    501  192.0.2.1  233.252.0.1    5010     16300001   10    el2
routed GRT    501  192.0.2.1  233.252.0.2    5010     16300002   20    el2
routed GRT    501  192.0.2.1  233.252.0.3    5010     16300003   10    el2
routed    GRT    501  192.0.2.1  233.252.0.4    5010     16300004   20    el2
routed    GRT    501  192.0.2.1  233.252.0.5    5010     16300005   10    el2
routed    GRT    501  192.0.2.1  233.252.0.6    5010     16300006   20    el2
routed    GRT    501  192.0.2.1  233.252.0.7    5010     16300007   10    el2
routed    GRT    501  192.0.2.1  233.252.0.8    5010     16300008   20    el2
routed    GRT    501  192.0.2.1  233.252.0.9    5010     16300009   10    el2
routed    GRT    501  192.0.2.1  233.252.0.10   5010     16300010   20    el2

--------------------------------------------------------------------------------
 Total Number of SPBM IP multicast ROUTE Entries: 10
--------------------------------------------------------------------------------

Switch:1#show isis spbm ip-multicast-route vrf green
================================================================================
                  SPBM IP-MULTICAST ROUTE INFO
================================================================================
Source         Group         Data ISID BVLAN Source-BEB
--------------------------------------------------------------------------------
192.0.2.10     233.252.0.1     16300001  10     el2
192.0.2.10     233.252.0.2     16300002  20     el2
192.0.2.10     233.252.0.3     16300003  10     el2
192.0.2.10     233.252.0.4     16300004  20     el2
192.0.2.10     233.252.0.5     16300005  10     el2
192.0.2.10     233.252.0.6     16300006  20     el2
192.0.2.10     233.252.0.7     16300007  10     el2
192.0.2.10     233.252.0.8     16300008  20     el2
192.0.2.10     233.252.0.9     16300009  10     el2
192.0.2.10 233.252.0.10     16300010  20     el2
--------------------------------------------------------------------------------
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
```

```
--------------------------------------------------------------------------------

Switch:1#show isis spbm ip-multicast-route vlan 501
================================================================================
             SPBM IP-multicast ROUTE INFO ALL
================================================================================
Type VrfName Vlan  Source    Group      VSN-ISID Data ISID  BVLAN Source-BEB

Id
--------------------------------------------------------------------------------
routed GRT   501  192.0.2.1  233.252.0.1  5010   16300001   10    el2
routed GRT   501  192.0.2.1  233.252.0.2  5010   16300002   20    el2
routed GRT   501  192.0.2.1  233.252.0.3  5010   16300003   10    el2
routed GRT   501  192.0.2.1  233.252.0.4  5010   16300004   20    el2
routed GRT   501  192.0.2.1  233.252.0.5  5010   16300005   10    el2
routed GRT   501  192.0.2.1  233.252.0.6  5010   16300006   20    el2
routed GRT   501  192.0.2.1  233.252.0.7  5010   16300007   10    el2
routed GRT   501  192.0.2.1  233.252.0.8  5010   16300008   20    el2
routed GRT   501  192.0.2.1  233.252.0.9  5010   16300009   10    el2
routed GRT   501  192.0.2.1  233.252.0.10 5010   16300010   20    el2

--------------------------------------------------------------------------------
 Total Number of SPBM IP multicast ROUTE Entries: 10
--------------------------------------------------------------------------------


Switch:1# show isis spbm ip-multicast-route vsn-isid 5010
================================================================================
SPBM IP-multicast ROUTE INFO - VLAN ID : 501, VSN-ISID : 5010
================================================================================
Source          Group           Data ISID  BVLAN Source-BEB
--------------------------------------------------------------------------------
192.0.2.1      233.252.0.2    16300002      20    el2
192.0.2.1      233.252.0.3    16300003      10    el2
192.0.2.1      233.252.0.4    16300004      20    el2
192.0.2.1      233.252.0.5    16300005      10    el2
192.0.2.1      233.252.0.6    16300006      20    el2
192.0.2.1      233.252.0.7    16300007      10    el2
192.0.2.1      233.252.0.8    16300008      20    el2
192.0.2.1      233.252.0.9    16300009      10    el2
192.0.2.1      233.252.0.10   16300010      20    el2

--------------------------------------------------------------------------------
Total Number of SPBM IP multicast ROUTE Entries: 10
--------------------------------------------------------------------------------

Switch:1# show isis spb-mcast-summary

================================================================================
              SPB multicast - Summary
================================================================================
SCOPE    SOURCE           GROUP            DATA            LSP  HOST
I-SID    ADDRESS          ADDRESS          I-SID     BVID  FRAG NAME
--------------------------------------------------------------------------------
5010     192.0.2.1       233.252.0.1     16300001  10    0x0  el2
5010     192.0.2.1       233.252.0.3     16300003  10    0x0  el2
5010     192.0.2.1       233.252.0.5     16300005  10    0x0  el2
5010     192.0.2.1       233.252.0.7     16300007  10    0x0  el2
5010     192.0.2.1       233.252.0.9     16300009  10    0x0  el2
```

```
5010     192.0.2.1       233.252.0.2     16300002  20    0x0   el2
5010     192.0.2.1       233.252.0.4     16300004  20    0x0   el2
5010     192.0.2.1       233.252.0.6     16300006  20    0x0   el2
5010     192.0.2.1       233.252.0.8     16300008  20    0x0   el2
5010     192.0.2.1       233.252.0.10    16300010  20    0x0   el2
```

Variable definitions

Use the data in the following table to use the **show isis spbm ip-multicast-route** command.

Variable	Value
all	Displays all IP multicast over SPBM route information.
detail	Displays detailed IP multicast over SPBM route information.
group*{A.B.C.D}*	Displays information on the group IP address for the IP multicast over SPBM route.
vlan*<1-4084>*	Specifies the VLAN ID in the range of 1-4084. VLAN IDs 1 to 4084 are configurable; VLAN IDs 4085 to 4094 are reserved for internal use. VLAN ID 1 is the default VLAN and cannot be created nor deleted.
vrf*WORD<1–16>*	Displays IP multicast over SPBM route information by VRF.
vsn-isid*<1–16777215>*	Displays IP multicast over SPBM route information by I-SID.

Use the data in the following table to use the **show isis spb-mcast-summary** command.

Variable	Value
host-name*WORD<0–255>*	Displays the SPBM multicast summary information by host-name.
lspid*<xxxx.xxxx.xxxx.xx-xx>*	Displays the SPBM multicast summary information by LSP ID.

Job aid

The following table describes the fields for the **show ip vrf mvpn** command.

Parameter	Description
Vrf name	Specifies the VRF name.
mvpn	Specifies if MVPN is enabled.
fwd-cache-timeout	Specifies the forward cache timeout for the VRF.

The following table describes the fields for the **show isis spbm ip-multicast-route** command.

Parameter	Description
Type	Specifies the type for the IP multicast over SPBM route.
VrfName	Specifies the VRF name.
Vlan Id	Specifies the ID for the C-VLAN.
Source	Specifies the group IP address for the IP multicast over SPBM route.
Group	Specifies the group IP address for the IP multicast over SPBM route.
VSN-ISID	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID. This is the scope I-SID.
Data ISID	Specifies the data I-SID for the IP multicast over SPBM route. After a BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16, 512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVLAN	Specifies the B-VLAN for the IP multicast over SPBM route.
Source-BEB	Specifies the source BEB for the IP multicast over SPBM route.

The following table describes the fields for the **`show isis spbm ip-multicast-route all`** command.

Parameter	Description
Source	Specifies the group IP address for the IP multicast over SPBM route.
Group	Specifies the group IP address for the IP multicast over SPBM route.
Data ISID	Specifies the data I-SID for the IP multicast route. After a BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVLAN	Specifies the B-VLAN for the IP multicast over SPBM route.
NNI Rcvrs	Specifies the NNI receivers.
UNI Rcvrs	Specifies the UNI receivers.

Parameter	Description
Source-BEB	Specifies the source BEB for the IP multicast over SPBM route.

The following table describes the fields for the **show isis spb-mcast-summary** command.

Parameter	Description
SCOPE I-SID	Specifies the scope I-SID. Layer 2 VSN and Layer 3 VSN each require a scope I-SID.
SOURCE ADDRESS	Specifies the group IP address for the IP multicast over SPBM route.
GROUP ADDRESS	Specifies the group IP address for the IP multicast over SPBM route.
DATA I-SID	Specifies the data I-SID for the IP multicast over SPBM route. After the BEB receives IP multicast data from a sender, the BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
BVID	Specifies the Backbone VLAN ID associated with the SPBM instance.
LSP FRAG	Specifies the LSP fragment number.
HOST NAME	Specifies the host name listed in the LSP, or the system name if the host is not configured.

Viewing IGMP information for Layer 3 VSN multicast

Use the following commands to check IGMP information.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display information about the interfaces where IGMP is enabled:

   ```
   show ip igmp interface [gigabitethernet {slot/port [-slot/port]
   [,...]][vlan <1-4084>[vrf WORD<1-16>][vrfids WORD<0-512>]
   ```

   Ensure that the output displays `routed-spb` under MODE.

3. Display information about the IGMP cache:

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

4. Display information about the IGMP group:

```
show ip igmp group [count][member-subnet default|{A.B.C.D/X}][vrf
WORD<1-16>][vrfids WORD<0-512>]
```

5. Display information about the IGMP sender:

```
show ip igmp sender [count][member-subnet default|{A.B.C.D/X}][vrf
WORD<1-16>][vrfids WORD<0-512>]
```

Example

Display IGMP information for Layer 3 VSN with IP multicast over SPBM:

```
Switch:#enable

Switch:1#show ip igmp interface vrf green

================================================================================
                         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

Switch:1(config)#show ip igmp interface vlan 501

================================================================================
                                  Vlan Ip Igmp
================================================================================
VLAN QUERY QUERY ROBUST VERSION LAST  PROXY  SNOOP  SSM    FAST   FAST
ID   INTVL MAX                  MEMB  SNOOP  ENABLE SNOOP  LEAVE  LEAVE
           RESP                 QUERY ENABLE        ENABLE ENABLE PORTS
--------------------------------------------------------------------------------
501   125   100   2      2       10    false  false  false  false

VLAN SNOOP   SNOOP          DYNAMIC   COMPATIBILITY EXPLICIT
ID   QUERIER QUERIER         DOWNGRADE MODE          HOST
     ENABLE  ADDRESS         VERSION                 TRACKING
--------------------------------------------------------------------------------
501  false   0.0.0.0         enable    disable       disable

Switch:1# show ip igmp sender vrf green

================================================================================
                              IGMP Sender - GlobalRouter
================================================================================
                                                                  PORT/
GRPADDR          IFINDEX    MEMBER          MLT        STATE
--------------------------------------------------------------------------------
233.252.0.1    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.2    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.3    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.4    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.5    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.6    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.7    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.8    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.9    Vlan 501   192.2.0.1      9/5        NOTFILTERED
233.252.0.10   Vlan 501   192.2.0.1      9/5        NOTFILTERED

10 out of 10 entries displayed
```

```
Switch:1# show ip igmp group vrf green

================================================================================
                            IGMP Group - GlobalRouter
================================================================================
GRPADDR          INPORT           MEMBER           EXPIRATION TYPE
--------------------------------------------------------------------------------
233.252.0.1     V501-9/16        192.2.0.1       204         Dynamic
233.252.0.2     V501-9/16        192.2.0.1       206         Dynamic
233.252.0.3     V501-9/16        192.2.0.1       206         Dynamic
233.252.0.4     V501-9/16        192.2.0.1       207         Dynamic
233.252.0.5     V501-9/16        192.2.0.1       204         Dynamic
233.252.0.6     V501-9/16        192.2.0.1       209         Dynamic
233.252.0.7     V501-9/16        192.2.0.1       206         Dynamic
233.252.0.8     V501-9/16        192.2.0.1       206         Dynamic
233.252.0.9     V501-9/16        192.2.0.1       211         Dynamic
233.252.0.10    V501-9/16        192.2.0.1       207         Dynamic

10 out of 10 group Receivers displayed

Total number of unique groups 10
```

Variable definitions

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

Variable	Value
gigabitethernet *{slot/port [-slot/port][,...]}*	Specifies the GigabitEthernet interface. Use <slot/port [-slot/port][,...]> to specify the slot and port.
vlan *<1-4084>*	Specifies the VLAN ID in the range of 1-4084. VLAN IDs 1 to 4084 are configurable; VLAN IDs 4085 to 4094 are reserved for internal use. VLAN ID 1 is the default VLAN and cannot be created nor deleted.
vrf *WORD<1–16>*	Specifies the VRF by name.
vrfids *WORD<0–512>*	Specifies the VRF by VRF ID.

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

Variable	Value
vrf *WORD<1–16>*	Specifies the VRF by name.
vrfids *WORD<0–512>*	Specifies the VRF by VRF ID.

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

Variable	Value
count	Specifies the number of entries.
group *{A.B.C.D}*	Specifies the group address.
member-subnet *{A.B.C.D/X}*	Specifies the IP address and network mask.
vrf *WORD<1–16>*	Displays the multicast route configuration for a particular VRF by name.

Variable	Value
vrfids *WORD<0–512>*	Displays the multicast route configuration for a particular VRF by VRF ID.

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

Variable	Value
count	Specifies the number of entries.
group *{A.B.C.D}*	Specifies the group address.
member-subnet *{A.B.C.D/X}*	Specifies the IP address and network mask.
vrf *WORD<1–16>*	Displays the multicast route configuration for a particular VRF by name.
vrfids *WORD<0–512>*	Displays the multicast route configuration for a particular VRF by VRF ID.

Job aid

The following table describes the fields for the **`show ip igmp interface`** command.

Parameter	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.
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.

Parameter	Description
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. If routed-spb displays in the MODE column, then IP multicast over SPBM is enabled on the Layer 3 VSN or for IP Shortcuts. If snoop-spb displays in the MODE column, then IGMP is enabled on a VLAN with an associated I-SID (Layer 2 VSN).

The following table describes the fields for the **`show ip igmp interface vlan`** command.

Parameter	Description
VLAN ID	Identifies the VLAN 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.
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.

Parameter	Description
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	Indicates the set of ports that are enabled for fast leave.
VLAN ID	Identifies the VLAN where IGMP is configured.
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 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.

The following table describes the fields for the **`show ip igmp cache`** command.

Parameter	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.
EXPIRATION	Indicates the minimum amount of time that remains before this entry ages out.
V1HOSTIMER	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.

The following table describes the fields for the **`show ip igmp group`** command.

Parameter	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 minimum amount of time that remains before this entry ages out.
TYPE	Indicates whether the entry is learned dynamically or is added statically.

The following table describes the fields for the **`show ip igmp sender`** command.

Parameter	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.

Viewing TLV information for a Layer 3 VSN with IP multicast over SPBM

Use the following commands to check TLV information.

For a Layer 3 VSN multicast, TLV 185 on the BEB where the source is located displays the multicast source and group addresses and have the Tx bit set. Each multicast group should have its own unique data I-SID with a value between 16,000,000 to 16,512,000. TLV 144 on the BEB bridge, where the sender is located, has the Tx bit set. All BEB bridges, where a receiver exists, have the Rx bit set.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display IS-IS Link State Database information by TLV:

   ```
   show isis lsdb tlv <1-186> [sub-tlv <1-3>][detail]
   ```

3. Display IS-IS Link State Database information by Link State Protocol ID:

   ```
   show isis lsdb lspid <xxxx.xxxx.xxxx.xx-xx> [tlv <1-186>] [sub-tlv
   <1-3>] [detail]
   ```

Example

Display TLV information for a Layer 3 VSN with IP multicast over SPBM:

```
Switch:1# show isis lsdb tlv 185 detail
================================================================================
            ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------
Level-1 LspID: 000c.f803.83df.00-04 SeqNum: 0x000002eb Lifetime: 1113
Chksum: 0x7e3b PDU Length: 556
Host_name: el2
Attributes: IS-Type 1
TLV:185 SPBM IPVPN :
    VSN ISID:5010
    BVID :10
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.1
        Data ISID : 16300011
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.3
        Data ISID : 16300013
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.5
        Data ISID : 16300015
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.7
        Data ISID : 16300017
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.9
        Data ISID : 16300019
        TX : 1
        VSN ISID:5010
        BVID :20
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.2
        Data ISID : 16300012
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.4
        Data ISID : 16300014
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.6
        Data ISID : 16300016
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
        Group Address : 233.252.0.8
        Data ISID : 16300018
        TX : 1
        Metric:0
        IP Source Address: 192.0.2.10
```

```
        Group Address : 233.252.0.10
        Data ISID : 16300020
        TX : 1
```

Variable definitions

Use the data in the following table to use the **show isis lsdb** command.

Variable	Value
detail	Displays detailed information about the IS-IS Link State database.
level *{l1, l2, l12}*	Displays information on the IS-IS level. The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled in the current release.
local	Displays information on the local LSDB.
lspid *<xxxx.xxxx.xxxx.xx-xx>*	Specifies information about the IS-IS Link State database by LSP ID.
sub-tlv *<1–3>*	Specifies information about the IS-IS Link State database by sub-TLV.
sysid *<xxxx.xxxx.xxxx>*	Specifies information about the IS-IS Link State database by System ID.
tlv *<1–186>*	Specifies information about the IS-IS Link State database by TLV.

Job aid

The following table describes the fields for the **show isis lsdb tlv** and the **show isis lsdb lspid** commands.

Parameter	Description
LSP ID	Indicates the LSP ID assigned to external IS-IS routing devices.
LEVEL	Indicates the level of the external router: l1, l2, or l12.
LIFETIME	Indicates the maximum age of the LSP. If the max-lsp-gen-interval is set to 900 (default), then the lifetime value begins to count down from 1200 seconds and updates after 300 seconds if connectivity remains. If the timer counts down to zero, the counter adds on an additional 60 seconds, then the LSP for that router is lost. This happens because of the zero age lifetime, which is detailed in the RFC standards.
SEQNUM	Indicates the LSP sequence number. This number changes each time the LSP is updated.
CKSUM	Indicates the LSP checksum. This is an error checking mechanism used to verify the validity of the IP packet.

Parameter	Description
HOST-NAME	Indicates the host-name.

Layer 3 VSN configuration using EDM

This section provides procedures to configure Layer 3 Virtual Services Networks (VSNs) using Enterprise Device Manager (EDM).

Configuring SPBM Layer 3 VSN

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 Virtual Services Network (VSN) to advertise IP routes across the SPBM network from one VRF to another using the following procedure.

SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF. In the VRF, just like in the Global Router (VRF 0), the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF and IP VPN instance on the switch. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.
- You must create the Customer VLANs and add slots/ports.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP-VPN**.
3. Click the **VPN** tab.
4. To create an IP VPN instance, click **Insert**.
5. Click the ellipsis button (**...**), select a VRF to associate with the IP VPN, and click **Ok**.
6. Click **Insert**.
7. In the **Enable** column, select **enable** to enable the IP VPN on the VRF.
8. In the **IsidNumber** column, specify an I-SID to associate with the VPN.
9. Click **Apply**.
10. In the navigation tree, expand the following folders: **Configuration** > **IP**.
11. Click **Policy**.
12. To identify routes on the local switch to be announced into the SPBM network, click the **Route Redistribution** tab.
13. Click **Insert**.

14. In the **DstVrfId** box, click the ellipsis button (**...**), select the destination VRF ID and click **Ok**.
15. In the **Protocol** box, click **isis** as the route destination.
16. In the **SrcVrfId** box, click (**...**) button, select the source VRF ID and click **Ok**.
17. In the **RouteSource** box, click the source protocol.
18. In the **Enable** box, click **enable**.
19. In the **RoutePolicy** box, click the ellipsis (**...**) button, choose the route policy to apply to the redistributed routes and click **Ok**.
20. Configure the other parameters as required.
21. Click **Insert**.
22. To apply the redistribution configuration, click the **Applying Policy** tab.
23. Select **RedistributeApply**, and then click **Apply**.

Configuring IS-IS redistribution

Use this procedure to configure IS-IS redistribution. In the Virtual Routing and Forwarding (VRF), just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IP VPN IP reachability information, the routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IS-IS**.
3. Click the **Redistribute** tab.
4. Click **Insert**.
5. Complete the fields as required.
6. Click **Insert**.

IS-IS Redistribute field descriptions

Use the data in the following table to configure the **IS-IS Redistribute** tab.

Name	Description
DstVrfId	Specifies the destination Virtual Routing and Forwarding (VRF) ID used in the redistribution.
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrfId	Specifies the source VRF ID used in the redistribution. For IS-IS, the source VRF ID must be the same as the destination VRF ID.

Name	Description
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables or disables a redistribution entry. The default is disable.
RoutePolicy	Specifies the route policy to be used for the detailed redistribution of external routes from a specified source into the IS-IS domain.
Metric	Specifies the metric for the redistributed route. The value can be a range between 0 to 65535. The default value is 0. Avaya recommends that you use a value that is consistent with the destination protocol.
MetricType	Specifies the metric type. Specifies a type1 or a type2 metric. For metric type1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type2, the cost of the external routes is equal to the external cost alone. The default is type2.
Subnets	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.

Applying IS-IS accept policies globally

Apply IS-IS accept policies globally. Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

After you apply the IS-IS accept filters, the device removes and re-adds all routes with updated filters.

IS-IS accept policies are disabled by default.

Note:

- After you apply IS-IS accept policies globally the application can disrupt traffic and cause temporary traffic loss. After you configure the IS-IS accept policies value to **Apply**, the device reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. Avaya recommends you make all the relevant accept policy changes, and then apply IS-IS accept policies globally at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Ensure the IP IS-IS filter exists.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Global** tab.
4. Select a name from the list or enter name in the **DefaultPolicyName** field to specify the route policy name for the default filter.
5. Select **Apply** to apply the default policy.

Accept Global field descriptions

Use the data in the following table to configure the **Accept Global** tab.

Name	Description
DefaultPolicyName	Specifies the route policy name for the default filter.
Apply	Applies the default policy when you configure the field to apply. The device only activates the default policy if the route map (the default policy name) has a value. If you do not select apply, the device takes no action. The GRT always returns no action.

Configuring an IS-IS accept policy for a specific advertising BEB

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB). Specify the SPBM nickname and the IS-IS accept policy name to allow you to apply the IS-IS accept policy.

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Nick Name** tab.
4. Click **Insert**.

5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Nick Name field descriptions

Use the data in the following table to configure the **Accept Nick Name** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy. You must configure a policy earlier in a separate procedure.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID

Configures a specific advertising Backbone Edge Bridge (BEB) with a specific I-SID to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.

6. In the **Isid** field, specify an I-SID number.
7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick-Name Isid descriptions

Use the date in the following table to configure the **Accept Nick-Name Isid** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB.
Isid	Specifies an I-SID used to filter. The value 0 is used for the Global Router.
Enable	Enables or disables the I-SID entry. The default is disabled.
PolicyName	Specifies the route policy name. You must configure a policy earlier in a separate procedure.

Configuring an I-SID list for an IS-IS accept policy

Configures a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. After you create the list of I-SID numbers, you must then create, configure, and enable the IS-IS accept policy.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify a name for the I-SID list.
6. Select **Isid** or **Isid-List**.
7. Specify an I-SID number or a list of I-SID numbers.
8. Click **Insert**.

Isid-List field descriptions

Use the data in the following table to configure the **Isid-List** tab.

Name	Description
Name	Specifies the name of the I-SID list.
Isid or **Isid-List**	Specifies that you either want to add a particular I-SID or a list of I-SID numbers.
Isid	Specifies a particular I-SID number or a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).

Configuring an IS-IS accept policy for a specific I-SID list

Configure an IS-IS accept policy for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify the I-SID list name.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Isid–List field descriptions

Use the data in the following table to configure **Accept Isid-List** tab.

Name	Description
Name	Specifies the name of I-SID list.
Enable	Enables or disables the I-SID list entry. The value must be enabled to filter. The default is disabled.

Name	Description
PolicyName	Specifies the route policy name.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB) for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

About this task

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid-List** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Name** field, specify an I-SID list name.
7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick–Name Isid-List field descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid-List** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default

Name	Description
	filter, but if a more specific filter is present the device applies the specific filter.
Name	Specifies the name of the I-SID list used to filter.
Enable	Enables or disables the SPBM nickanme advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy name.

Enabling MVPN for a VRF

Use this procedure to enable MVPN for a particular VRF. IP multicast over SPBM, constrains multicast streams of senders to all receivers in the same Layer 3 VSN. MVPN functionality is disabled by default.

Note:

VLAN level configuration is also required to turn on the service on each VLAN within the VRF on which this services is required. You can turn it on under the VLAN context or the brouter context.

Before you begin

- You must enable SPBM multicast globally.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IP-MVPN**.
3. Click the **MVPN** tab.
4. Double-click in the **Enable** field in the table.
5. Select **Enable** from the drop down menu.
6. Double-click in the **FwdCacheTimeout** field in the table, and then type the VRF timeout value.
7. Click **Apply**.

MVPN field descriptions

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

Name	Description
VrfId	Specifies the VRF ID.
Enable	Enables Layer 3 VSN IP multicast over SPBM services for a particular VRF. The default is disabled.
FwdCacheTimeout	Specifies the VRF timeout value. The timeout value ages out the sender when there is no multicast stream on the VRF. The default is 210 seconds..

Configuring IP multicast over SPBM on a VLAN for Layer 3

Use this procedure to enable IP multicast over SPBM for a Layer 3 VSN. The default is disabled.

To configure a VLAN for IP Shortcuts with IP multicast over SPBM, see Configuring IP multicast over SPBM on a VLAN on page 215.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must configure a VRF and an IP VPN instance with an I-SID configured under it on the switch. The IP VPN does not need to be enabled for Layer 3 VSN multicast to function.
- You must enable SPBM multicast globally.
- If there is no IP interface on the VLAN, then you create one. (The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN).
- You must enable MVPN for the particular VRF.

About this task

You must configure VLANs to turn on the service on each VLAN with in the VRF on which the service is required. You can turn it on under the VLAN context or the brouter context.

If you only want to use IP multicast over SPBM, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP multicast over SPBM routing does not depend on unicast routing (for Layer 3 VSN). This allows for you to more easily migrate from a PIM environment to IP multicast over SPBM. You can migrate a PIM environment to IP multicast over SPBM first and then migrate unicast separately or not at all.

The switch only supports IPv4 address with IP multicast over SPBM.

Note:

You cannot enable IP PIM when IP multicast over SPBM is enabled on the VLAN.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **VRF Context View**.
2. Click **Set VRF Context View**.
3. Choose a VRF name.
4. Click **Launch VRF Context View**.
5. Select an enabled port on the Physical Device View.
6. From the navigation tree, expand the following folders: **Configuration** > **VLAN**
7. Click **VLANs**.
8. Choose a VLAN, and then click the **IP** from under the tab bar.
9. Click the **SPB Multicast** tab.
10. Check the **Enable** box.
11. Click **Apply**.

Configuring IP multicast over SPBM on a brouter port for a Layer 3 VSN

Use this procedure to enable IP multicast over SPBM on a brouter port. The default is disabled.

To configure a brouter port for IP Shortcuts with IP multicast over SPBM, see Configuring IP multicast over SPBM on a brouter port for IP Shortcuts on page 216.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must configure a VRF and an IP VPN instance with an I-SID configured under it on the switch. The IP VPN does not need to be enabled for Layer 2 VSN multicast to function.
- You must enable SPBM multicast globally.
- If there is no IP interface on the VLAN, then you create one. (The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN).
- You must enable MVPN for the particular VRF.

About this task

You must enable IP multicast over SPBM on each of the VLANs that need to support IP multicast traffic.

If you only want to use IP multicast over SPBM, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP multicast over SPBM routing does not depend on unicast routing, which allows for you to more easily migrate from a PIM environment to multicast over SPBM. You can migrate a PIM environment to IP multicast over SPBM first, and then migrate unicast separately or not at all.

The switch only supports IPv4 address with IP multicast over SPBM.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **VRF Context View**.
2. Click **Set VRF Context View**.
3. Choose a VRF name.
4. Click **Launch VRF Context View**.
5. Select an enabled port on the Physical Device View.
6. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
7. Click **IP**.
8. Click the **SPB Multicast** tab.
9. Click **Enable**.
10. Click **Apply**.

Configuring IGMP on a VLAN interface for a Layer 3 VRF

Use this procedure to configure IGMP for each VLAN interface to enable the interface to perform multicast operations.

IGMPv2 at the VLAN level is the default setting, with no other configuration required. You only need to enable IGMPv3. You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

Note:

You cannot enable IP PIM when IP multicast over SPBM is enabled on the VLAN.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF and IP VPN instance with an I-SID on the switch.
- You must create the C-VLANs and add slots/ports.
- You must enable IP multicast over SPBM for a Layer 3 VSN.

About this task

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **VRF Context View**.
2. Click **Set VRF Context View**.
3. Choose a VRF name.
4. Click **Launch VRF Context View**.
5. From the navigation tree, expand the following folders: **Configuration** > **VLAN**.
6. Click **VLANs**.
7. Select the desired VLAN from the listing.
8. Click the **IP** button.
9. Click the **IGMP** tab.
10. **(Optional)** If you want to enable SsmSnoopEnable, select the **SsmSnoopEnable** box.
11. **(Optional)** If you want to enable Snoop, select the **SnoopEnable** box.
12. **(Optional)** In the **Version** box, select the correct IGMP version.

    You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.
13. **(Optional)** Select **SnoopQuerierEnable**, to enable Snoop Querier. Only select this option, if you want to configure an address for the IGMP queries.
14. **(Optional)** In the **SnoopQuerierAddr** box, type an IP address, if you want to configure a snoop querier address.

Note:

If the SPBM bridge connects to an edge switch, it can be necessary to add an IGMP query address. If you omit adding a query address, the SPB bridge sends IGMP queries with a source address of 0.0.0. Some edge switch models do not accept a query with a source address of 0.0.0.0.

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. 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 SSM 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.

Name	Description
	For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.
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 switch 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.
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. If the SPBM bridge connects to an edge switch, it can be necessary to add an IGMP query address. If you omit adding a query address, the SPBM bridge sends IGMP queries with a source address of 0.0.0.0. Some edge switch models do not accept a query with a source address of 0.0.0.0.

Configuring the Global Routing Table timeout value

Use this procedure to configure the timeout value in the GRT. The timeout value ages out the sender when there are no multicast streams coming from the sender for a specified period of time. The default timeout value is 210 seconds.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the C-VLANs and add slots/ports.
- You must enable IP multicast over SPBM globally.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **IS-IS** > **SPBM**.
2. Click the **SPBM** tab.
3. Modify the **McastFwdCacheTimeout** value.
4. Click **Apply**.

Viewing the IGMP interface table

Use the Interface tab to view the IGMP interface table. When an interface does not use an IP address, it does not appear in the IGMP table.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**
2. Click **IGMP**.
3. Click the **Interface** tab.

Layer 3 VSN configuration example

The following figure shows a sample Layer 3 VSN deployment.

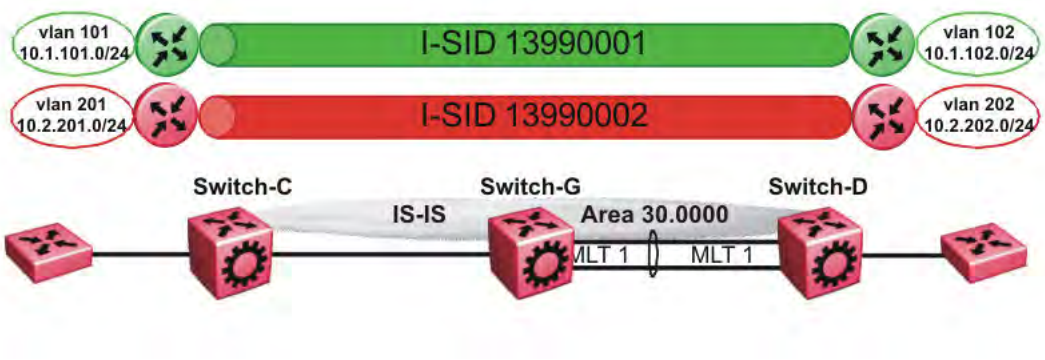

Figure 28: Layer 3 VSN

The following sections show the steps required to configure the Layer 3 VSN parameters in this example.

Note that IP IS-IS redistribution needs to be configured to inject the VRF routes into IS-IS.

You must first configure basic SPBM and IS-IS infrastructure.

For more information, see SPBM configuration examples on page 114.

VRF green configuration

The following figure shows the green VRF in this Layer 3 VSN example.

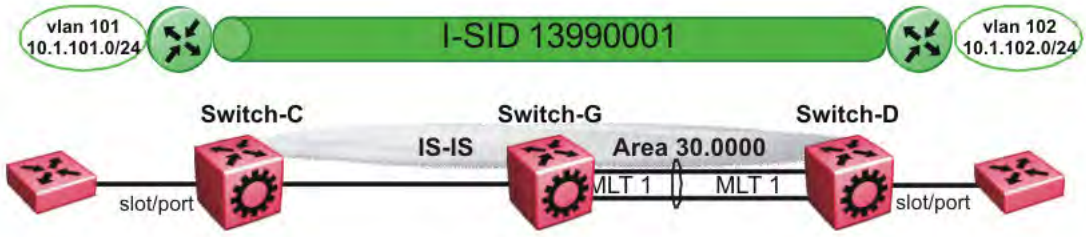

Figure 29: Layer 3 VSN — VRF green

The following sections show the steps required to configure the green VRF parameters in this example.

Note:

The following steps use slot/port 4/2, which is appropriate for a VSP 9000 configuration. The slot/port configuration for your product may be different.

VRF green – Switch-C

```
VRF CONFIGURATION

ip vrf green vrfid 1
```

```
VLAN CONFIGURATION

vlan create 101 type port-mstprstp 1
vlan mlt 101 1
vlan members 101 4/2 portmember
interface Vlan 101
vrf green
ip address 10.1.101.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 13990001
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf green
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct vrf green
```

VRF green – Switch-D

```
VRF CONFIGURATION

ip vrf green vrfid 1

VLAN CONFIGURATION

vlan create 102 type port-mstprstp 1
vlan mlt 102 1
vlan members add 102 4/2 portmember
interface vlan 102
vrf green
ip address 10.1.102.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 13990001
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf green
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS
```

```
isis apply redistribute direct vrf green
```

VRF red configuration

The following figure shows the red VRF in this Layer 3 VSN example.

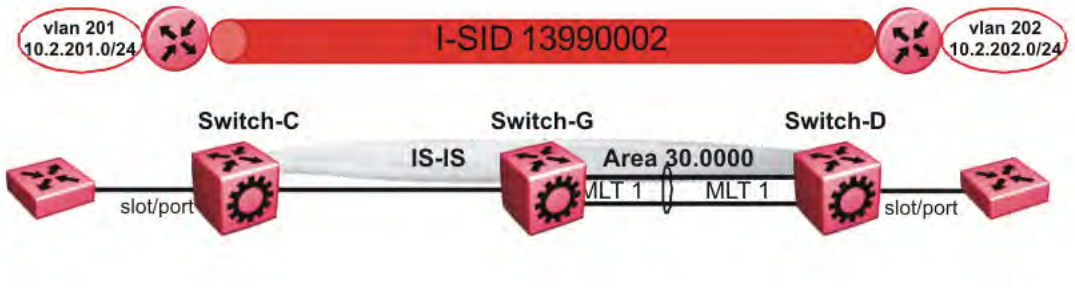

Figure 30: Layer 3 VSN — VRF red

The following sections show the steps required to configure the red VRF parameters in this example.

Note:

The following steps use slot/port 4/2, which is appropriate for a VSP 9000 configuration. The slot/port configuration for your product may be different.

VRF red – Switch-C

```
VRF CONFIGURATION

ip vrf red vrfid 2

VLAN CONFIGURATION

vlan create 201 type port-mstprstp 1
vlan mlt 201 1
vlan members 201 4/2 portmember
interface Vlan 201
vrf red
ip address 10.2.201.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf red
ipvpn
i-sid 13990002
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf red
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS
```

```
isis apply redistribute direct vrf red
```

VRF red – Switch-D

```
VRF CONFIGURATION

ip vrf red vrfid 2

VLAN CONFIGURATION

vlan create 202 type port-mstprstp 1
vlan mlt 101 1
vlan members 202 4/2 portmember
interface Vlan 202
vrf red
ip address 10.3.202.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf red
ipvpn
i-sid 13990002
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf red
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct vrf red
```

Verifying Layer 3 VSN operation

The following sections show the steps required to verify the Layer 3 VSN configuration in this example.

Note:

The following example uses slot/port numbers that are applicable to the VSP 9000, such as 4/20 and 4/30. The slot/port configuration for your product may be different.

Switch-C

```
Switch-C:1# show isis spbm ip-unicast-fib
================================================================================
                    SPBM IP-UNICAST FIB ENTRY INFO
================================================================================
     VRF   DEST                                 OUTGOING  SPBM   PREFIX IP ROUTE
VRF  ISID  ISID  Destination  NH BEB   VLAN INTERFACE  COST   COST   PREFERENCE
--------------------------------------------------------------------------------
GRT   -     -     10.0.0.2/32  Switch-D 4000  4/30      20      1       7
GRT   -     -     10.0.14.0/24 Switch-D 4000  4/30      20      1       7
--------------------------------------------------------------------------------
```

```
 Total number of SPBM IP-UNICAST FIB entries 2
-------------------------------------------------------------------------------
```

```
Switch-C:1# show isis spbm ip-unicast-fib id 13990001
===============================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO
===============================================================================================
    VRF   DEST                                    OUTGOING  SPBM  PREFIX  IP ROUTE
VRF ISID  ISID     Destination   NH BEB   VLAN    INTERFACE COST  COST    PREFERENCE
-----------------------------------------------------------------------------------------------
green -   13990001 10.1.101.0/24 Switch-D 4000    4/20      20      1       7
-----------------------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 1
-----------------------------------------------------------------------------------------------

Switch-C:1# show isis spbm ip-unicast-fib id 13990002
===============================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO
===============================================================================================
    VRF  DEST                                    OUTGOING SPBM    PREFIX IP ROUTE
VRF ISID ISID     Destination   NH BEB   VLAN     INTERFACE COST   COST   PREFERENCE
-----------------------------------------------------------------------------------------------
red -    13990002 10.2.202.0/24 Switch-D 4000     4/30       20     1    7

-----------------------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 1
-----------------------------------------------------------------------------------------------
```

```
Switch-C:1# show isis spbm ip-unicast-fib id all
===============================================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO

===============================================================================================
      VRF    DEST                                      OUTGOING  SPBM   PREFIX IP ROUTE
VRF   ISID   ISID      Destination   NH BEB   VLAN    INTERFACE COST   COST   PREFERENCE
-----------------------------------------------------------------------------------------------
GRT    -     -        10.0.0.2/32   Switch-D 4000    4/30      20     1       7
GRT    -     -        10.0.14.0/24  Switch-D 4000    4/30      20     1       7
green  -     13990001  10.1.102.0/24 Switch-D 4000     4/30      20     1       7
red    -     13990002  10.2.202.0/24 Switch-D 4000     4/30      20     1       7

-----------------------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 4
-----------------------------------------------------------------------------------------------
```

Switch-D

```
Switch-D:1# show isis spbm ip-unicast-fib
===============================================================================
     VRF   DEST                                  OUTGOING  SPBM  PREFIX IP ROUTE
VRF  ISID  ISID  Destination   NH BEB   VLAN    INTERFACE COST  COST   PREFERENCE
-------------------------------------------------------------------------------
GRT   -    -     10.0.0.1/32   Switch-C 4000    4/20      20    1      7
GRT   -    -     10.0.13.0/24  Switch-C 4000    4/20      20    1      7
-------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 2
-------------------------------------------------------------------------------
```

```
Switch-D:1# show isis spbm ip-unicast-fib id 13990001
===============================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO
===============================================================================================
     VRF   DEST                                    OUTGOING  SPBM  PREFIX IP ROUTE
VRF  ISID  ISID     Destination   NH BEB    VLAN   INTERFACE COST  COST   PREFERENCE
-----------------------------------------------------------------------------------------------
green -    13990001 10.1.101.0/24 Switch-C  4000    4/20      20    1     7
```

```
-------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 1
-------------------------------------------------------------------------------

Switch-D:1# show isis spbm ip-unicast-fib id 13990002
===============================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO
===============================================================================
     VRF   DEST                                    OUTGOING  SPBM  PREFIX IP ROUTE
VRF  ISID  ISID      Destination   NH BEB    VLAN  INTERFACE COST  COST   PREFERENCE
-------------------------------------------------------------------------------
red   -    13990002  10.2.201.0/24 Switch-C  4000   4/20     20        1      7
-------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 1
-------------------------------------------------------------------------------
```

```
Switch-D:1# show isis spbm ip-unicast-fib id all
===============================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO
===============================================================================
     VRF   DEST                                    OUTGOING  SPBM  PREFIX IP ROUTE
VRF  ISID  ISID      Destination   NH BEB    VLAN  INTERFACE COST  COST   PREFERENCE
-------------------------------------------------------------------------------
GRT   -     -       10.0.0.1/32   Switch-C 4000    4/20     20     1      7
GRT   -     -        10.0.13.0/24  Switch-C 4000     4/20     20        1      7
green -    13990001  10.1.101.0/24 Switch-C 4000     4/20     20        1      7
red   -    13990002  10.2.201.0/24 Switch-C 4000     4/20     20        1      7
-------------------------------------------------------------------------------
 Total number of SPBM IP-UNICAST FIB entries 4
-------------------------------------------------------------------------------
```

VRF green—Switch-C

```
Switch-C:1# show ip route vrf green
===============================================================================
                              IP Route - VRF green
===============================================================================
                                                NH                INTER
DST              MASK            NEXT          VRF/ISID    COST FACE PROT AGE TYPE PRF
-------------------------------------------------------------------------------
10.1.101.0     255.255.255.0   10.1.101.1     -             1    101   LOC  0   DB   0
10.1.102.0     255.255.255.0   Switch-D       vrf green     20   4000  ISIS 0   IBSV 7

2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.
-------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Rout
e,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

VRF green—Switch-D

```
Switch-D:1# show ip route vrf green
===============================================================================
                              IP Route - VRF green
===============================================================================
                                               NH                INTER
DST              MASK            NEXT          VRF/ISID      COST FACE PROT AGE TYPE PRF
-------------------------------------------------------------------------------
10.1.101.0     255.255.255.0   Switch-C       vrf green     20   4000 ISIS 0   IBSV 7
10.1.102.0     255.255.255.0   10.1.102.1     -              1    102  LOC  0   DB   0

2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.
-------------------------------------------------------------------------------
```

```
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Rout
e,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

VRF red—Switch-C

```
Switch-C:1# show ip route vrf red
================================================================================
                                IP Route - VRF red
================================================================================
                                          NH                INTER
DST              MASK            NEXT     VRF/ISID     COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.2.201.0     255.255.255.0   10.2.201.1   -            1   201  LOC  0   DB   0
10.2.202.0     255.255.255.0   Switch-D     vrf red      20  4000 ISIS 0   IBSV 7

2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.
--------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Rout
e,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

VRF red—Switch-D

```
Switch-D:1# show ip route vrf red
================================================================================
                                IP Route - VRF red
================================================================================
                                          NH                INTER
DST              MASK            NEXT     VRF/ISID     COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.2.201.0     255.255.255.0   Switch-C     vrf red      20   4000 ISIS 0   IBSV 7
10.2.202.0     255.255.255.0   10.2.202.1   -            1    202 LOC  0   DB   0
2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.
--------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Rout
e,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

Layer 3 VSN with IP multicast over SPBM configuration example

The example below shows the configuration to enable IP multicast over SPBM support on VLANs 500 and 501 that are part of VRF Green:

```
ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable


VRF CONFIGURATION

ip vrf green vrfid 2
```

```
VLAN CONFIGURATION - PHASE 1

vlan 110 i-sid 100
interface vlan 500
vrf green
ip address 192.0.2.1 255.255.255.0 1
ip spb-multicast enable
exit

vlan 111 i-sid 100
interface vlan 501
vrf green
ip address 192.0.2.2 255.255.0 0
ip spb-multicast enable
exit


ISIS SPBM IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 100
mvpn enable
exit
```

When using IGMPv3, the configuration is:

```
ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable


VRF CONFIGURATION

ip vrf green vrfid 2


VLAN CONFIGURATION - PHASE 1

vlan 110 i-sid 100
interface vlan 500
vrf green
ip address 192.0.2.1 255.255.255.0 1
ip spb-multicast enable
ip igmp version 3
exit

vlan 111 i-sid 100
interface vlan 501
vrf green
ip address 192.0.2.2 255.255.0 0
ip spb-multicast enable
ip igmp version 3
exit


ISIS SPBM IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 100
mvpn enable
exit
```

Inter-VSN routing configuration

This section provides concepts and procedures to configure Inter-Virtual Services Network (VSN) routing.

Inter-VSN routing configuration fundamentals

This section provides fundamental concepts on Inter-VSN Routing.

Inter-VSN routing

Inter-VSN routing with SPBM allows routing between Layer 2 VLANs with different I-SIDs.

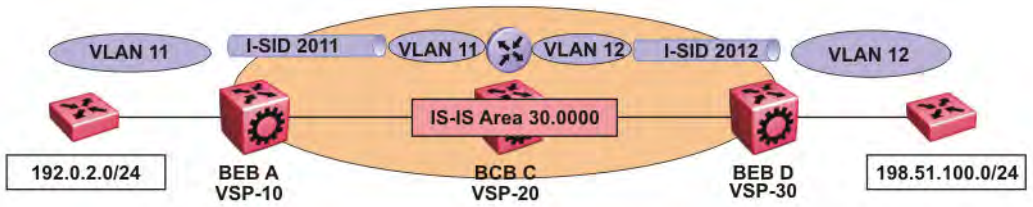

Figure 31: Inter-VSN routing

Inter-VSN routing provides a routing hub for Layer 2 Virtual Services Network edge devices. Layer 3 devices, routers, or hosts connect to the SPBM cloud using the SPBM Layer 2 VSN service. To go between a routed network, a Layer 2 VSN termination point provides the routing services to hop onto another Layer 2 VSN, using the I-SID.

Note:

The Layer 2 VLANs must be in the same VRF. You cannot route traffic between two different VRFs with Inter-VSN routing.

In this example, the C-VLANs are associated with I-SIDs on the BEBs using SPBM Layer 2 VSN. With Inter-VSN routing enabled, BCB C can transmit traffic between VLAN 11 (I-SID 2011) and VLAN 12 (I-SID 2012). The BEB switches can forward traffic between VLANs 11 and 12 on the VRF instance configured on the BCB.

IP interfaces are where the routing instance exists. In this case, on VSP-20.

Note:

Virtual Services Platform 9000 does not support IP multicast over SPBM routing on inter-VSN routing interfaces.

Inter-VSN routing configuration using ACLI

This section provides a procedure to configure Inter-VSN routing using ACLI.

Configuring SPBM Inter-VSN Routing

Inter-VSN routing allows routing between IP networks on Layer 2 VLANs with different I-SIDs.

Note:

The Layer 2 VLANs must be in the same VRF. You cannot route traffic between two different VRFs with Inter-VSN routing.

Inter-VSN routing is only typically used when you have to extend a VLAN as Layer 2 Virtual Services Networks (VSNs) for applications such as vMotion. Normally, Avaya recommends you use either IP shortcuts or Layer 3 VSNs to route traffic.

You must configure both the Backbone Edge Bridges (BEBs) and the Backbone Core Bridge (BCB). For a full configuration example, see Inter-VSN routing with SPBM configuration example on page 306.

Note:

To enable inter-VSN routing, you must configure an IP interface where the routing instance exists.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.

Procedure

1. Use the following procedure on the Backbone Edge Bridges (BEBs) containing the VSNs you want to route traffic between.
2. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

3. Create a customer VLAN (C-VLAN) by port:

   ```
   vlan create <2-4084> type port-mstprstp <0-63>
   ```

4. Add ports in the C-VLAN:

   ```
   vlan members add <1-4084> {slot/port [-slot/port][,...]}
   ```

5. Map a C-VLAN to an I-SID:

   ```
   vlan i-sid <1-4084> <0-16777215>
   ```

Important:

When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.

6. On the Backbone Core Bridge (BCB), use the following procedure to create a VRF, and add a VLAN for each VSN.
7. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

8. Create a VRF:

   ```
   ip vrf WORD<1-16> vrfid <1-511>
   ```

9. Create a VLAN to associate with each VSN:

   ```
   vlan create <2-4084> type port-mstprstp <0-63>
   ```

10. Enter VLAN Interface Configuration mode:

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

11. Add a VLAN to the VRF you created:

    ```
    vrf WORD<1-16>
    ```

12. Associate an I-SID with the VLAN:

    ```
    vlan i-sid <1-4084> <0-16777215>
    ```

    Important:

    When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.

    Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

13. Configure a circuitless IP interface (CLIP) address for the VLAN:

    ```
    ip address {A.B.C.D/X}
    ```

14. Repeat the preceding steps for every VLAN you want to route traffic between.

Variable definitions

Use the data in the following table to use the **vlan create** command.

Variable	Value
<1–4084>	Specifies the VLAN ID.
type port-mstprstp *<0–63>* [color *<0–32>*]	Creates a VLAN by port: • *<0–63>* is the STP instance ID. • *color <0–32>* is the color of the VLAN.

Use the data in the following table to use the **vlan members add** command.

Variable	Value
<1–4084>	Specifies the VLAN ID.
{slot/port [slot/port][,...]}	Specifies the port number using slot/port notation.

Use the data in the following table to use the **vlan i-sid** command.

Variable	Value
vlan i-sid *<1–4084> <0–16777215>*	Specifies the C-VLAN to associate with the I-SID. Use the no or default options to remove the I-SID from the specified VLAN. **Note:** Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

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

Variable	Value
WORD *<1–16>*	Create the VRF and specify the name of the VRF instance.
vrfid *<1–511>*	Specifies the VRF instance by ID number.

Use the data in the following table to use the **vrf** command.

Variable	Value
WORD *<1–16>*	Specifies the VRF name. Associates a port to a VRF.

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

Variable	Value
{A.B.C.D/X}	Configure a circuitless IP interface (CLIP) to provide a virtual interface not associated with a physical port.

Inter-VSN routing configuration using EDM

This section provides procedures to configure Inter-VSN routing using Enterprise Device Manager (EDM).

Configuring BEBs for Inter-VSN routing

Inter-VSN routing allows routing between IP networks on Layer 2 VLANs with different I-SIDs.

Note:

The Layer 2 VLANs must be in the same VRF. You cannot route traffic between two different VRFs with Inter-VSN routing.

Inter-VSN routing is only typically used when you have to extend a VLAN as Layer 2 Virtual Services Networks (VSNs) for applications such as vMotion. Normally, Avaya recommends you use either IP Shortcuts or Layer 3 VSNs to route traffic. You must configure both the Backbone Edge Bridges (BEBs) and the Backbone Core Bridge (BCB).

Note:

To enable inter-VSN routing, you must configure the IP interface where the routing instance exists.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.

About this task

Follow the procedures below on the Backbone Edge Bridges (BEBs) containing the VSNs you want to route traffic between.

Procedure

1. Create a customer VLAN (C-VLAN) by port and add ports in the C-VLAN. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. In the **Basic** tab, click **Insert**.
4. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
5. In the **Name** box, type the VLAN name, or use the name provided.
6. In the **Color Identifier** box, click the down arrow and choose a color from the list, or use the color provided.
7. In the **Type** box, select **byPort**.

8. In the **PortMembers** box, click the **(...)** button.
9. Click on the ports to add as member ports.

   The ports that are selected are recessed, while the nonselected ports are not recessed. Port numbers that appear dimmed cannot be selected as VLAN port members.
10. Click **OK**.
11. Click **Insert**.
12. Collapse the **VLANs** tab.

    The VLAN is added to the Basic tab.
13. Map a C-VLAN to an I-SID. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
14. Click **VLANs**.
15. Click the **Advanced** tab.
16. To map a C-VLAN to an I-SID, in the **I-sid** field, specify the I-SID to associate with the specified VLAN.

    Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.
17. Click **Apply**.

    Important:

    When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.
18. Configure the Backbone Core Bridge (BCB) for Inter-VSN Routing. For more information, see Configuring BCBs for Inter-VSN routing on page 299.

Basic field descriptions

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

Name	Description
Id	Specifies the VLAN ID for the VLAN.
Name	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.
Color Identifier	Specifies a proprietary color scheme to associate a color with the VLAN. Color does not affect how frames are forwarded.

Name	Description
Type	Specifies the type of VLAN: • byPort • byIpSubnet • byProtocolId • bySrcMac • spbm-bvlan
MstpInstance	Identifies the MSTP instance.
VrfId	Indicates the Virtual Router to which the VLAN belongs.
VrfName	Indicates the name of the Virtual Router to which the VLAN belongs.
PortMembers	Specifies the slot/port of each VLAN member.
ActiveMembers	Specifies the slot/port of each VLAN member.
StaticMembers	Specifies the slot/port of each static member of a policy-based VLAN.
NotAllowToJoin	Specifies the slot/ports that are never allowed to become a member of the policy-based VLAN.
OspfPassiveMembers	Specifies the slot/ports of each Open Shortest Path First (OSPF) passive member.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers of remote function call (RFC). • ip (IP version 4) • ipx802dot3 (Novell Internetwork Packet Exchange (IPX) on Ethernet 802.3 frames) • ipx802dot2 (Novell IPX on IEEE 802.2 frames) • ipxSnap (Novell IPX on Ethernet Standard Network Access Protocol (SNAP) frames) • ipxEthernet2 (Novell IPX on Ethernet Type 2 frames) • appleTalk [AppleTalk on Ethernet Type 2 and Ethernet Symbolic Network Analysis Program (SNAP) frames] • decLat (Digital Equipment Corporation Local Area Transport (DEC LAT) protocol) • decOther (Other DEC protocols) • sna802dot2 (IBM SNA on IEEE 802.2 frames) • snaEthernet2 (IBM SNA on Ethernet Type 2 frames)

Name	Description
	• netBIOS (NetBIOS protocol) • xns (Xerox XNS) • vines (Banyan VINES) • ipv6 (IP version 6) • usrDefined (user-defined protocol) • rarp (Reverse Address Resolution Protocol) • PPPoE (Point-to-Point Protocol over Ethernet) If the VLAN type is port-based, none is displayed in the Basic tab ProtocolId field.
SubnetAddr	Specifies the source IP subnet address (IP subnet-based VLANs only).
SubnetMask	Specifies the source IP subnet mask (IP subnet-based VLANs only).

Advanced field descriptions

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

Name	Description
Id	Specifies the VLAN ID.
Name	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.
Type	Specifies the type of VLAN: • byPort • byIpSubnet • byProtocolId • bySrcMac • spbm-bvlan
I-sid	Specifies the I-SID number assigned to a customer VLAN (C-VLAN). The range is 0 — 16777215. The default value is 0, which indicates that no I-SID is assigned. **Note:** Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID,

Name	Description
	therefore I-SID 0x00ffffff cannot be used for any other service.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers RFC: • ip (IP version 4) • ipx802dot3 (Novell IPX on Ethernet 802.3 frames) • ipx802dot2 (Novell IPX on IEEE 802.2 frames) • ipxSnap (Novell IPX on Ethernet SNAP frames) • ipxEthernet2 (Novell IPX on Ethernet Type 2 frames) • appleTalk (AppleTalk on Ethernet Type 2 and Ethernet SNAP frames) • decLat (DEC LAT protocol) • decOther (Other DEC protocols) • sna802dot2 (IBM SNA on IEEE 802.2 frames) • snaEthernet2 (IBM SNA on Ethernet Type 2 frames) • netBIOS (NetBIOS protocol) • xns (Xerox XNS) • vines (Banyan VINES) • ipv6 (IP version 6) • usrDefined (user-defined protocol) • RARP (Reverse Address Resolution protocol) • PPPoE (Point-to-point protocol over Ethernet) If the VLAN type is not protocol-based, None is displayed in the Basic tab ProtocolId field.
Encap	Specifies the encapsulation method. Values are: • Ethernet II • SNAP — SubNetwork Access Protocol (SNAP) • LLC — IEEE 802.2 Logic Link Control (LLC) This is the encapsulation type for user-defined protocol-based VLANs. The **Encap** option is not meaningful for other types of VLAN. By default, there is no encapsulation method configured for the VLAN.
AgingTime	Specifies the timeout period for dynamic VLAN membership; a potential VLAN port is made ACTIVE

Name	Description
	after it receives a packet that matches the VLAN; if no such packet is received for AgingTime seconds, the port is no longer active. The default is 600.
MacAddress	Specifies the MAC address assigned to the virtual router interface for this VLAN. This field is relevant only after the VLAN is configured for routing. This MAC address is used as the Source MAC in routed frames, ARP replies, or Routing Information Protocol (RIP) and OSPF frames.
Vlan Operation Action	Performs an operation on the VLAN. The values are: • none • flushMacFdb: Configures action to flushMacFdb. This action removes the learned MAC addresses from the forwarding database for the selected VLAN. • flushArp: Configures action to flushArp. This action removes the ARP entries from the address table for the selected VLAN. • flushIp: Configures action to flushIp. This action removes the learned IP addresses from the forwarding table for the selected VLAN. • flushDynMemb: Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based VLAN and removes MAC addresses learned on those ports. • all: Configures action to all. This action performs all the supported actions; it does does not perform the Snoop-related actions. • flushSnoopMemb: This action is not supported. • triggerRipUpdate: Configures action to triggerRipUpdate. After you execute this command the Virtual Services Platform 9000 immediately sends a RIP request to solicit the updated RIP routes. • flushSnoopMRtr: This action is not supported. The default is none.
Result	Specifies the result code after you perform an action. The default is none.
UserDefinedPid	Specifies the 16-bit user-defined network protocol identifier of a protocol-based VLAN with User Defined protocol.

Name	Description
NlbMode	Specifies if the NLB administrative privileges are enabled or disabled. The default value is disable. SPBM supports Network Load Balancing (NLB) unicast. SPBM does not support NLB multicast or NLB multicast with IGMP.

Configuring BCBs for Inter-VSN routing

Inter-VSN routing allows routing between IP networks on Layer 2 VLANs with different I-SIDs.

Note:

The Layer 2 VLANs must be in the same VRF. You cannot route traffic between two different VRFs with Inter-VSN routing.

Inter-VSN routing is only typically used when you have to extend a VLAN as Layer 2 Virtual Services Networks (VSNs) for applications such as vMotion. Normally, Avaya recommends you use either IP shortcuts or Layer 3 VSNs to route traffic. You must configure both the Backbone Edge Bridges (BEBs) and the Backbone Core Bridge (BCBs).

Note:

To enable inter-VSN routing, you must configure the IP interface where the routing instance exists.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must configure the Backbone Edge Bridges (BEBs) containing the VSNs you want to route traffic between. For more information, see Configuring BEBs for Inter-VSN routing on page 293.

About this task

Follow the procedures below to configure the Backbone Core Bridge (BCB) for inter-VSN routing.

Procedure

1. On the Backbone Core Bridge (BCB), create a VRF. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **VRF**.
3. Click **Insert**.
4. Specify the VRF ID.
5. Name the VRF instance.
6. Configure the other parameters as required.
7. Click **Insert**.
8. Create a VLAN to associate with each VSN. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.

9. Click **VLANs**.
10. In the **Basic** tab, click **Insert**.
11. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
12. In the **Name** box, type the VLAN name, or use the name provided.
13. In the **Color Identifier** box, click the down arrow and choose a color from the list, or use the color provided.
14. In the **Type** box, select **byPort**.
15. In the **PortMembers** box, click the **(...)** button.
16. Click on the ports to add as member ports.

    The ports that are selected are recessed, while the nonselected ports are not recessed. Port numbers that appear dimmed cannot be selected as VLAN port members.
17. Click **OK**.
18. Click **Insert**.
19. Collapse the **VLANs** tab.

    The VLAN is added to the **Basic** tab.
20. Associate the VLAN with an I-SID. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
21. Click **VLANs**.
22. In the VLANs tab, click the **Advanced** tab.
23. In the **I-sid** box, specify the I-SID to associate with the VLAN.
24. Click **Apply**.
25. Configure a circuitless IP interface (CLIP). In the navigation tree, expand the following folders: **Configuration** > **IP**.
26. Click **IP**.
27. Click the **Circuitless IP** tab.
28. Click **Insert**.
29. In the **Interface** field, assign a CLIP interface number.
30. Enter the IP address.
31. Enter the network mask.
32. Click **Insert**.

VRF field descriptions

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

Name	Dexcription
Id	Specifies the ID number of the VRF instance. VRF ID 0 is reserved for the GlobalRouter.
Name	Names the VRF instance.
ContextName	Identifies the VRF. The SNMPv2 Community String or SNMPv3 contextName denotes the VRF context and is used to logically separate the MIB module management.
TrapEnable	Enables the VRF to send VRF Lite-related traps (VrfUp and VrfDown). The default is true.
MaxRoutes	Configures the maximum number of routes allowed for the VRF. The default value is 10000, except for the GlobalRouter, which is 250000.
RpTrigger	Specifies the Routing Protocol (RP) triggers for the VRF. The triggers are used to initiate or shutdown routing protocols on a VRF. The protocols include RIP, OSPF and BGP. Multiple RPs can be acted upon simultaneously. Also, you can use this option to bring individual RPs up in steps.
MaxRoutesTrapEnable	Enables the generation of the VRF Max Routes Exceeded traps. The default is true.

Basic field descriptions

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

Name	Description
Id	Specifies the VLAN ID for the VLAN.
Name	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.
Color Identifier	Specifies a proprietary color scheme to associate a color with the VLAN. Color does not affect how frames are forwarded.
Type	Specifies the type of VLAN: • byPort • byIpSubnet • byProtocolId • bySrcMac • spbm-bvlan
MstpInstance	Identifies the MSTP instance.
VrfId	Indicates the Virtual Router to which the VLAN belongs.
VrfName	Indicates the name of the Virtual Router to which the VLAN belongs.

Name	Description
PortMembers	Specifies the slot/port of each VLAN member.
ActiveMembers	Specifies the slot/port of each VLAN member.
StaticMembers	Specifies the slot/port of each static member of a policy-based VLAN.
NotAllowToJoin	Specifies the slot/ports that are never allowed to become a member of the policy-based VLAN.
OspfPassiveMembers	Specifies the slot/ports of each Open Shortest Path First (OSPF) passive member.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers of remote function call (RFC). • ip (IP version 4) • ipx802dot3 (Novell Internetwork Packet Exchange (IPX) on Ethernet 802.3 frames) • ipx802dot2 (Novell IPX on IEEE 802.2 frames) • ipxSnap (Novell IPX on Ethernet Standard Network Access Protocol (SNAP) frames) • ipxEthernet2 (Novell IPX on Ethernet Type 2 frames) • appleTalk [AppleTalk on Ethernet Type 2 and Ethernet Symbolic Network Analysis Program (SNAP) frames] • decLat (Digital Equipment Corporation Local Area Transport (DEC LAT) protocol) • decOther (Other DEC protocols) • sna802dot2 (IBM SNA on IEEE 802.2 frames) • snaEthernet2 (IBM SNA on Ethernet Type 2 frames) • netBIOS (NetBIOS protocol) • xns (Xerox XNS) • vines (Banyan VINES) • ipv6 (IP version 6) • usrDefined (user-defined protocol) • rarp (Reverse Address Resolution Protocol) • PPPoE (Point-to-Point Protocol over Ethernet) If the VLAN type is port-based, none is displayed in the Basic tab ProtocolId field.

Name	Description
SubnetAddr	Specifies the source IP subnet address (IP subnet-based VLANs only).
SubnetMask	Specifies the source IP subnet mask (IP subnet-based VLANs only).

Advanced field descriptions

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

Name	Description
Id	Specifies the VLAN ID.
Name	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.
Type	Specifies the type of VLAN: • byPort • byIpSubnet • byProtocolId • bySrcMac • spbm-bvlan
I-sid	Specifies the I-SID number assigned to a customer VLAN (C-VLAN). The range is 0 — 16777215. The default value is 0, which indicates that no I-SID is assigned. **Note:** Virtual Services Platform 9000 reserves I-SID 0x00ffffff. Virtual Services Platform 9000 uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers RFC: • ip (IP version 4) • ipx802dot3 (Novell IPX on Ethernet 802.3 frames) • ipx802dot2 (Novell IPX on IEEE 802.2 frames) • ipxSnap (Novell IPX on Ethernet SNAP frames) • ipxEthernet2 (Novell IPX on Ethernet Type 2 frames)

Name	Description
	• appleTalk (AppleTalk on Ethernet Type 2 and Ethernet SNAP frames) • decLat (DEC LAT protocol) • decOther (Other DEC protocols) • sna802dot2 (IBM SNA on IEEE 802.2 frames) • snaEthernet2 (IBM SNA on Ethernet Type 2 frames) • netBIOS (NetBIOS protocol) • xns (Xerox XNS) • vines (Banyan VINES) • ipv6 (IP version 6) • usrDefined (user-defined protocol) • RARP (Reverse Address Resolution protocol) • PPPoE (Point-to-point protocol over Ethernet) If the VLAN type is not protocol-based, None is displayed in the Basic tab ProtocolId field.
Encap	Specifies the encapsulation method. Values are: • Ethernet II • SNAP — SubNetwork Access Protocol (SNAP) • LLC — IEEE 802.2 Logic Link Control (LLC) This is the encapsulation type for user-defined protocol-based VLANs. The **Encap** option is not meaningful for other types of VLAN. By default, there is no encapsulation method configured for the VLAN.
AgingTime	Specifies the timeout period for dynamic VLAN membership; a potential VLAN port is made ACTIVE after it receives a packet that matches the VLAN; if no such packet is received for AgingTime seconds, the port is no longer active. The default is 600.
MacAddress	Specifies the MAC address assigned to the virtual router interface for this VLAN. This field is relevant only after the VLAN is configured for routing. This MAC address is used as the Source MAC in routed frames, ARP replies, or Routing Information Protocol (RIP) and OSPF frames.
Vlan Operation Action	Performs an operation on the VLAN. The values are: • none • flushMacFdb: Configures action to flushMacFdb. This action removes the learned MAC addresses

Name	Description
	from the forwarding database for the selected VLAN. • flushArp: Configures action to flushArp. This action removes the ARP entries from the address table for the selected VLAN. • flushIp: Configures action to flushIp. This action removes the learned IP addresses from the forwarding table for the selected VLAN. • flushDynMemb: Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based VLAN and removes MAC addresses learned on those ports. • all: Configures action to all. This action performs all the supported actions; it does does not perform the Snoop-related actions. • flushSnoopMemb: This action is not supported. • triggerRipUpdate: Configures action to triggerRipUpdate. After you execute this command the Virtual Services Platform 9000 immediately sends a RIP request to solicit the updated RIP routes. • flushSnoopMRtr: This action is not supported. The default is none.
Result	Specifies the result code after you perform an action. The default is none.
UserDefinedPid	Specifies the 16-bit user-defined network protocol identifier of a protocol-based VLAN with User Defined protocol.
NlbMode	Specifies if the NLB administrative privileges are enabled or disabled. The default value is disable. SPBM supports Network Load Balancing (NLB) unicast. SPBM does not support NLB multicast or NLB multicast with IGMP.

Circuitless IP field descriptions

Use the data in the following table to use the **Circuitless IP** tab.

Name	Description
Interface	Specifies the number assigned to the interface, from 1 to 256.
Ip Address	Specifies the IP address of the CLIP.
Net Mask	Specifies the network mask.

Inter-VSN routing configuration example

This section provides a configuration example for Inter-VSN routing.

Inter-VSN routing with SPBM configuration example

The following figure shows a sample Inter-VSN deployment.

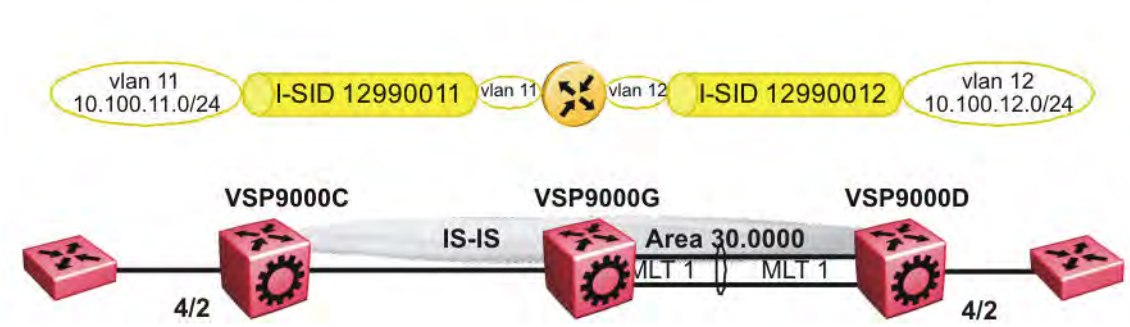

Figure 32: Inter-VSN routing configuration

The following sections show the steps required to configure the Inter-VSN parameters in this example. You must first configure basic SPBM and IS-IS infrastructure. For more information, see: SPBM configuration examples on page 114.

Note that the IP interfaces are configured where the routing instance exists, namely, on VSP9000G.

Note:

The Layer 2 VSNs must be in the same VRF. You cannot route traffic between two different VRFs with Layer 2 VSNs.

VSP9000C

```
VLAN CONFIGURATION

vlan create 11 type port-mstprstp 1
vlan members 11 4/2 portmember
vlan i-sid 11 12990011
```

VSP9000G

```
VRF CONFIGURATION

ip vrf blue vrfid 100

VLAN CONFIGURATION

vlan create 11 type port-mstprstp 1
vlan i-sid 11 12990011
interface Vlan 11
vrf blue

CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
interface loopback 3
```

```
ip address 10.100.11.1 255.255.255.0
exit

VLAN CONFIGURATION

vlan create 12 type port-mstprstp 1
vlan i-sid 12 12990012
interface Vlan 12
vrf blue
exit

CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter

interface loopback 4
ip address 10.100.12.1 255.255.255.0
exit
```

VSP9000D

```
VLAN CONFIGURATION

vlan create 12 type port-mstprstp 1
vlan members 12 4/2 portmember
vlan i-sid 12 12990012
```

Verifying Inter-VSN Routing operation

The following sections show how to verify Inter-VSN Routing operation in this example.

VSP9000G

```
VSP9000G:1# show ip route vrf blue
================================================================================
                              IP Route - VRF blue
================================================================================
                                                NH              INTER
DST              MASK            NEXT           VRF/ISID   COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.100.11.0    255.255.255.0   10.100.11.1    -          1    11     LOC  0   DB   0
10.100.12.0    255.255.255.0   10.100.12.1    -          1    12     LOC  0   DB   0

VSP9000G:1# show ip arp vrf blue
================================================================================
                              IP Arp - VRF blue
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT   TYPE    TTL(10 Sec)    TUNNEL
--------------------------------------------------------------------------------
10.100.11.1     00:0e:62:25:a2:00  11       -        LOCAL   2160
10.100.11.255   ff:ff:ff:ff:ff:ff  11       -        LOCAL   2160
10.100.12.1     00:0e:62:25:a2:01  12       -        LOCAL   2160
10.100.12.255   ff:ff:ff:ff:ff:ff  12       -        LOCAL   2160

================================================================================
                            IP Arp Extn - VRF blue
================================================================================
MULTICAST-MAC-FLOODING    AGING(Minutes)       ARP-THRESHOLD
--------------------------------------------------------------------------------
disable                   360                  500

4 out of 50 ARP entries displayed
```

VSP9000G

```
VSP9000G:1# show vlan mac-address-entry  11
================================================================================
```

```
                                    Vlan Fdb
================================================================================
VLAN            MAC                                    SMLT
ID   STATUS     ADDRESS            INTERFACE           REMOTE     TUNNEL
--------------------------------------------------------------------------------
11   learned    00:00:00:00:01:02  Port-4/2              false   VSP9000C
11   self       00:0e:62:25:a2:00  Port-cpp            false      -

2 out of 4 entries in all fdb(s) displayed.

VSP9000G:1# show vlan mac-address-entry  12
================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                    SMLT
ID   STATUS     ADDRESS            INTERFACE           REMOTE     TUNNEL
--------------------------------------------------------------------------------
12   learned    00:00:00:00:02:02  Port-4/2            false    VSP9000D
12   self       00:0e:62:25:a2:01  Port-cpp            false      -

2 out of 4 entries in all fdb(s) displayed.
```

VSP9000C

```
VSP9000C:1# show vlan mac-address-entry 11
================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                    SMLT
ID   STATUS     ADDRESS            INTERFACE           REMOTE         TUNNEL
--------------------------------------------------------------------------------
11   learned    00:00:00:00:01:02  Port-4/2            false        VSP9000D
11   learned    00:0e:62:25:a2:00  Port-4/2            false        VSP9000D

2 out of 2 entries in all fdb(s) displayed.
```

VSP9000D

```
VSP9000D:1# show vlan mac-address-entry 12
================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                    SMLT
ID   STATUS     ADDRESS            INTERFACE           REMOTE         TUNNEL
--------------------------------------------------------------------------------
12   learned    00:00:00:00:02:02  Port-4/2            false          VSP9000C
12   learned    00:0e:62:25:a2:01  Port-4/2            false          VSP9000C

2 out of 2 entries in all fdb(s) displayed.
```

Chapter 5: Operations and Management

CFM fundamentals

In a network, you need the ability to isolate a connectivity fault to correct it. When the network consists of an SPBM cloud as well as a customer domain, Connectivity Fault Management (CFM) helps determine where the problem exists to debug connectivity issues and isolate faults. By allowing CFM to break a network into sections using MEPs and MIPs, you can determine where the problem lies.

Typically the backbone nodes only learn Backbone MAC (B-MAC) addresses, while only the appropriate Backbone Edge Bridges (BEBs), which terminate the virtual services networks (VSN), learn the Customer MAC (C-MAC) addresses. As such, the nodes within the SPBM backbone have no knowledge of the C-MAC or IP addresses used within the Virtual Services Networks (VSNs) and only need to provide reachability to the B-MAC addresses within the backbone.

CFM divides or separates a network into administrative domains called Maintenance Domains (MD). CFM further subdivides each MD into logical groupings called Maintenance Associations (MA). A single MD can contain several MAs.

Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Two types of MP exist:

- Maintenance End Point (MEP)
- Maintenance Intermediate Point (MIP)

By allowing CFM to break a network into sections using MEPs and MIPs, the user can determine where in the network the problem exists.

You can explicitly configure MDs, MAs, MEPs and MIPs and associate them with multiple VLANs or you can use autogenerated CFM commands that create a MEP and MIP at a specified level for every SPBM B-VLAN or CMAC C-VLAN. If you choose to autogenerate CFM commands, the VSP 9000 creates the MD, MA and MEP ID used for each MEP.

Note:

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you

must first remove the existing MEP and MIP on the SPBM B-VLAN. VSP 9000 only supports one MEP or MIP on the SPBM B-VLAN.

CFM provides loopback messages (LBM), which act like ping. CFM also provides linktrace messages (LTM), which act like traceroute. As a result, you can debug Layer 2 with CFM. The Virtual Services Platform 9000 wraps the CFM LBM and LTM into easier commands, namely **`l2 ping`** and **`l2 traceroute`** respectively. The l2 commands only require a VLAN and destination target MAC address to use.

The MEPs and MIPs that you configure for SPBM B-VLANs do not respond to CFM messages sent from C-MAC VLANs because the VLAN used is not the same and packet encapsulation is different. You must use autogenerated CFM MEP and MIP level for every C-VLAN on the chassis. You can use either explicitly configured or autogenerated CFM MEP and MIP for SPBM B-VLANs.

CFM is based on the IEEE 802.1ag standard. To support troubleshooting of the SPBM cloud, Virtual Services Platform 9000 supports a subset of CFM functionality. CFM is the standard for Layer 2 ping, Layer 2 traceroute, and the end-to-end connectivity check of the Ethernet network.

On Virtual Services Platform 9000, CFM is implemented to support Loopback Messages (LBM), and Linktrace Messages (LTM). Messages are sent between Maintenance Points (MP) in the system. Continuity Check Messages (CCM) are not required or supported in the current release.

Autogenerated CFM and explicitly configured CFM

VSP 9000 simplifies CFM configuration with autogenerated CFM. With autogenerated CFM, you use the commands **`cfm spbm enable`** and **`cfm cmac enable`** and VSP 9000 creates default MD, MA, MEPs, and MIPs for SPBM B-VLANs and C-VLANs respectively.

- For SPBM B-VLANs, VSP 9000 provides two methods to configure CFM: autogenerated and explicitly configured. You cannot use both.
- For C-VLANs, you can only use autogenerated CFM.

Autogenerated CFM

You can use autogenerated CFM at a global level to create a MEP and a MIP at a specified level for every SPBM B-VLAN and C-VLAN on the chassis. If you use autogenerated CFM commands, you do not have to configure explicit MDs, MAs, MEPs, or MIPs, and associate them with multiple VLANs.

If you do not want to use autogenerated CFM commands, you can choose to configure explicit MDs, MAs, MEPs, and MIPs for SPBM B-VLANs. However, you cannot use both an autogenerated CFM configuration and an explicit CFM configuration together.

Note:

Previous explicit CFM configurations of MDs, MAs, and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLANs. VSP 9000 only supports one type of MEP or MIP for each SPBM B-VLAN.

For autogenerated CFM configuration information for ACLI see:

- Configuring autogenerated CFM on SPBM B-VLANs on page 321.
- Configuring autogenerated CFM on C-VLANs on page 323.

For autogenerated CFM configuration information for EDM see:

- Configuring autogenerated CFM on SPBM B-VLANs on page 349.
- Configuring autogenerated CFM on C-VLANs on page 351.

Explicitly configured CFM

If you choose to explicitly configure CFM, you must configure an MD, MA, MEPs, and MIPs.

For explicit configuration information for ACLI see:

- Configuring CFM MD on page 326.
- Configuring CFM MA on page 327.
- Assigning a MEP MIP level to an SPBM B-VLAN on page 328.
- Configuring CFM MEP on page 330.

For explicit configuration information for EDM see:

- Configuring CFM MD on page 352.
- Configuring CFM MA on page 353.
- Configuring CFM MEP on page 354.
- Configuring CFM nodal MEP on page 355.

Using CFM

For SPBM B-VLANs, the autogenerated MEPs and MIPs respond to **`l2 ping`**, **`l2 traceroute`**, and **`l2 tracetree`** in the same manner as the MEPs and MIPs created explicitly. For C-VLANs, the autogenerated MEPs and MIPs respond to **`l2 ping`** and **`l2 traceroute`**, but not to **`l2 tracetree`** because no multicast trees exist on C-VLANs. The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **`l2 ping`** and **`l2 traceroute`** requests.

Customer VLAN vs. SPBM B-VLAN configurations

CFM breaks the network into sections, called MEPs, so you can determine exactly where the problem exists.

The MEPs and MIPs configured for SPBM B-VLANs do not respond to CFM messages sent by C-VLANs.

To forward customer traffic across the core network backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a B-MAC destination address (BMAC-DA) and a B-MAC source address (BMAC-SA). In SPBM, each node

populates its forwarding database (FDB) with the B-MAC information derived from the IS-IS shortest path tree calculations.

Typically the SPBM Backbone Core Bridges (BCBs) in the SPBM cloud only learn the B-MAC addresses. The Backbone Edge Bridges (BEBs) know the Customer MACs on the appropriate BEBs that terminate the virtual services networks (VSNs). As such, the nodes within the SPBM cloud have no knowledge of the C-MAC addresses in the VSNs.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For C-VLANs, you have to trigger an **`l2 ping`** to learn the C-MAC address.
- For B-VLANs, you do not have to trigger an **`l2 ping`** to learn the C-MAC address because IS-IS populates the MAC addresses in the FDB table.

In both cases, linktrace traces the path up to the closest device to that MAC address that supports CFM in the SPBM cloud.

C-VLANs source addresses

CFM uses either the VLAN MAC or the CFM C-MAC for the BMAC-SA for the C-VLANs. The CFM C-MAC is the value of the management base MAC, which ends in 0x3c0. The system creates the VLAN MAC after a user adds an IP address to a VLAN.

If a VLAN has a MAC address, the system uses the VLAN MAC as the BMAC-SA by default. If a VLAN does not have a MAC address, the system uses the CFM C-MAC for the BMAC-SA. You may also configure the system to use the CFM C-MAC, even if a VLAN MAC exists.

Maintenance Domain (MD)

A Maintenance Domain (MD) is the part of a network that is controlled by a single administrator. For example, a customer can engage the services of a service provider, who, in turn, can engage the services of several operators. In this scenario, there can be one MD associated with the customer, one MD associated with the service provider, and one MD associated with each of the operators.

You assign one of the following eight levels to the MD:

- 0–2 (operator levels)
- 3–4 (provider levels)
- 5–7 (customer levels)

The levels separate MDs from each other and provide different areas of functionality to different devices using the network. An MD is characterized by a level and an MD name (optional).

A single MD can contain several Maintenance Associations (MA).

Maintenance Association (MA)

An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance association End Points (MEPs), each configured with the same

Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

The following figure shows MD level assignment in accordance with the 802.1ag standard. As shown in the figure, MIPs can be associated with MEPs. However, MIPs can also function independently of MEPs.

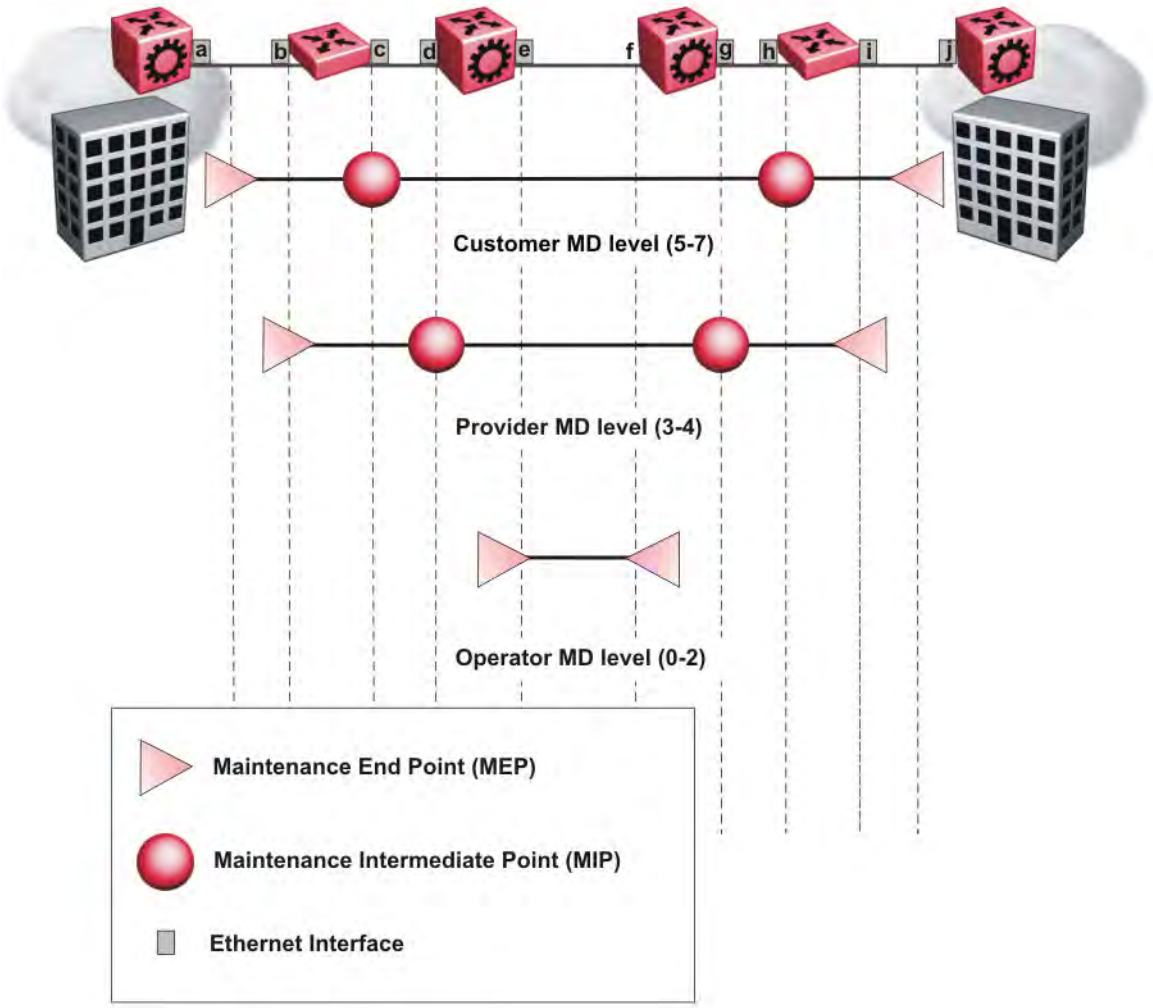

Figure 33: MD level assignment

Maintenance endpoints (MEP)

A Maintenance Endpoint (MEP) defines the end of a link and a Maintenance Intermediate Point is a point in the middle of the network.

A Maintenance Endpoint (MEP) represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA. MEP functionality can be divided into the following functions:

- Fault Detection
- Fault Verification
- Fault Isolation
- Fault Notification

Fault detection and notification are achieved through the use of Continuity Check Messages (CCM). CCM messages are not supported in the current release.

Maintenance domain intermediate points (MIP)

MIPs do not initialize any CFM messages. MIPs passively receive CFM messages, process the messages received and respond back to the originating MEP. By responding to received CFM messages, MIPs can support discovery of hop-by-hop path among MEPs, allow connection failures to be isolated to smaller segments of the network to help discover location of faults along the paths. MIPs can be created independent of MEPs. MIP functionality can be summarized as:

- Respond to Loopback (ping) messages at the same level as itself and addressed to it.
- Respond to Linktrace (traceroute) messages.
- Forward Linktrace messages after decrementing the TTL.

Fault verification

Fault verification is achieved through the use of Loopback Messages (LBM). An LBM is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a MEP or Maintenance Intermediate Point (MIP) but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR.

LBM message

The LBM packet is often compared to a ping. A MEP transmits the LBM packet. This packet can be addressed to another MEP or to the MAC address of the MP; in the case of SPBM, this is the SPBM system ID or its virtual SMLT MAC. Only the MP for which the packet is addressed responds with an LBR message.

- Provides “ICMP ping like” functionality natively at Layer-2.
- DA is the MAC address of the target.

- Includes a transaction identifier that allows the corresponding LBR to be identified when more than one LBM request is waiting for a response.
- Bridges forward the frame using the normal FDB rules.
- Only the target (MIP or MEP) responds.
- Initiator can choose the size and contents data portion of the LBM frame.
- Can be used to check the ability of the network to forward different sized frames.

l2 ping

The **l2 ping** command is a proprietary command that allows a user to trigger an LBM message.

- For B-VLANs, specify either the destination MAC address or node name.
- For C-VLANs, specify the destination MAC address.

Note:

To use Layer 2 ping for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the **cfm cmac enable** command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the **l2 ping** command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **l2 ping** command to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **l2 ping** and **l2 traceroute** requests.

The **l2 ping** command provides a simpler command syntax than the standard LBM commands, which require you to specify the MD, MA, and MEP ID information. The **l2 ping** command provides a ping equivalent at Layer 2 for use with nodes on the SPBM B-VLAN or C-VLANs.

The options supported for the **l2 ping** command vary based on the VLAN type. Only SPBM B-VLANs support the SMLT virtual option for the source mode. Only C-VLANs support the no VLAN MAC option on the source mode.

l2 ping with IP address

The **l2 ping** command also allows you to specify an IP address as the destination address. In this case, the IP address can be either C-VLAN or B-VLAN in the SPBM cloud.

The **l2 ping** command converts Layer 3 IP information to an appropriate Layer 2 VLAN and MAC combination. The system can also target IP addresses that are not SPBM derived routes.

When the **l2 ping** command is executed with an IP address as the destination, the operation finds all the valid MAC combinations that provide valid paths to the destination. If ECMP is enabled, there

can be multiple paths to the destination. In this case, **`l2 ping`** runs internally for each of the VLAN paths returned, and displays a summary of the results. If ECMP is disabled, the results display only one path.

Fault isolation

Fault isolation is achieved through the use of Linktrace Messages (LTM). LTM is intercepted by all the MPs on the way to the destination MP. Virtual Services Platform 9000 supports two types of LTM.

The first type, the unicast LTM, can be addressed to either MEP or MIP MAC address. Each MP on the way decrements the TTL field in the LTM frame, sends Linktrace Reply (LTR), and forwards the original LTM to the destination. The LTM is forwarded until it reaches its destination or the TTL value is decremented to zero. LTR is a unicast message addressed to the originating MEP.

The second type, the proprietary LTM, is used to map the MAC addresses of the SPBM network; in this case the target MAC is not an MP, but rather a service instance identifier (I-SID).

Link trace message

Connectivity Fault Management offers link trace messaging for fast fault detection. Link trace messages allow operators, service providers and customers to verify the connectivity that they provide or use and to debug systems.

Link trace message — unicast

The link trace message (LTM) is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP, which is the SPBM system id or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

- Trace the path to any certain MAC address.
- DA is unicast
- LTM contains:
  - Time to live (TTL)
  - Transaction Identifier
  - Originator MAC address
  - Target MAC address
- CFM unaware entities forward the frame as is like any other data frame.
- MIP or MEP that is not on the path to the target discards the LTM and does not reply.
- MIP that is on the path to the target:
  - Forwards the LTM after decrementing the TTL and replacing the SA with its own address.
  - Sends a reply (LTR) to the originator.
  - Identifies itself in the forwarded LTM and LTR by modifying TLV information.

- If the MIP or MEP is a target:
  - Sends an LTR to the originator.
  - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- A MEP that is not the target but is on the path to the target:
  - Generates a reply as described in the preceding information.
  - It also sets one of the flags fields in the reply to indicate that it is the terminal MEP.

Link trace message — multicast

The multicast link trace message (LTM) can be used to trace the multicast tree from any node on any I- SID using the nickname MAC address and the I-SID multicast address.

Specifying a multicast target address for an LTM allows for the tracing of the multicast tree corresponding to that destination address (DA). With a multicast target every node that is in the active topology for that multicast address responds with a Linktrace reply and also forwards the LTM frame along the multicast path. Missing Linktrace replies (LTRs) from the nodes in the path indicate the point of first failure.

This functionality allows you to better troubleshoot I-SID multicast paths in a SPBM network.

l2 traceroute

The **`l2 traceroute`** command is a proprietary command that allows a user to trigger an LTM message.

- For B-VLANs, specify either the destination MAC address or node name.
- For C-VLANs, specify the destination MAC address.

Note:

To use Layer 2 traceroute for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the **`cfm cmac enable`** command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the **`l2 traceroute`** command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **`l2 traceroute`** command to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **`l2 ping`** and **`l2 traceroute`** requests.

This command provides a simpler command syntax than the standard LTM commands, which require you to specify the MD, MA, and MEP ID information. The **l2 traceroute** command provides a trace equivalent at Layer 2 for use with nodes on the SPBM B-VLAN or C-VLANs.

The options supported for the **l2 traceroute** command vary based on the VLAN type. Only SPBM B-VLANs support the SMLT virtual option for the source mode. Only C-VLANs support the no VLAN MAC option on the source mode.

l2 traceroute with IP address

The **l2 traceroute** command allows you to specify an IP address as the destination address. In this case, the IP address can be either C-VLAN or B-VLAN in the SPBM cloud.

The **l2 traceroute** command converts Layer 3 IP information to an appropriate Layer 2 VLAN and MAC combination. The system can also target IP addresses that are not SPBM derived routes.

If ECMP is enabled, **l2 traceroute** runs internally for each of the VLAN paths returned, and displays a summary of the results. If ECMP is disabled, the results display only one path.

Destination addresses for C-VLAN l2 traceroute and linktrace messages

For C-VLANs, CFM uses the following destination MAC addresses for the corresponding maintenance domain (MD) levels for l2 traceroute and linktrace messages.

VSP 9000 supports both l2 traceroute and linktrace for C-VLANs, but Avaya prefers you use l2 traceroute.

Table 4: MD levels and corresponding destination addresses for CFM for C-VLANs

CFM MD level	Destination MAC address
0	01:80:c2:00:00:38
1	01:80:c2:00:00:39
2	01:80:c2:00:00:3a
3	01:80:c2:00:00:3b
4	01:80:c2:00:00:3c
5	01:80:c2:00:00:3d
6	01:80:c2:00:00:3e
7	01:80:c2:00:00:3f

l2 tracetree

The **l2 tracetree** command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

VSP 9000 only supports this command on SPBM B-VLANs. VSP 9000 does not support **l2 tracetree** command for C-VLANs.

Layer 2 tracemroute

The **l2tracemroute** command is a proprietary command that allows the user to trace the multicast tree for a certain multicast flow. The user specifies source, group, and service context (either VLAN or VRF) for the multicast flow to trace.

CFM sends a multicast LTM using an internal calculation to map the source, group, and context to the corresponding target address. The LTR comes from all leaves of the multicast tree for that flow, as well as transit nodes. The target MAC used in the LTM is a combination of the data I-SID and the nickname and the packet is sent on the appropriate SPBM B-VLAN. The user can see the generated multicast tree for that flow, which includes the data I-SID and nickname.

Nodal MPs

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. The Nodal MEP provides traceability and troubleshooting at the system level for a certain B-VLAN. Each node (chassis) has a certain MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP. The Nodal B-VLAN MPs supports eight levels of CFM and you configure the Nodal B-VLAN MPs on a per B-VLAN basis. Virtual SMLT MAC addresses are also able to respond for LTM and LBM.

Nodal B-VLAN MEPs

The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs. To support this behavior a MAC Entry is added to the FDB and a new CFM data-path table containing the B-VLAN and MP level are added to direct CFM frames to the CP as required.

Nodal B-VLAN MIPs

The Nodal MIP is associated with a B-VLAN. VLAN and level are sufficient to specify the Nodal MIP entity. The Nodal MIP MAC address is the SPBM system ID for the node on which it resides. If the fastpath sends a message to the CP, the MIP responds if it is not the target and the MEP responds if it is the target.

Nodal B-VLAN MEPs and MIPs with SMLT

When Nodal MEPs or MIPs are on SPBM B-VLANs the LTM code uses a unicast MAC DA. The LTM DA is the same as the target MAC address, which is the SPBM MAC address or the SMLT MAC address of the target node.

Virtual Services Platform 9000 supports SMLT interaction with SPBM. This is accomplished by using two B-VIDs into the core from each pair of SMLT terminating nodes. Both nodes advertise the Nodal B-MAC into the core on both B-VIDS. In addition each node advertises the SMLT virtual B-MAC on one of the two B-VLANs.

The Nodal MEP and MIP are expanded to respond to both the Nodal MAC address as well as the Virtual SMLT MAC address if both MACs are being advertised on its B-VLAN. In addition a source mode is added to the LTM and LBM command to use either the Nodal MAC or the SMLT virtual MAC address as the source MAC in the packet.

Configuration considerations

When you configure CFM, be aware of the following configuration considerations:

- A single switch has a limit of two nodal MEPs and two nodal MIPs
- All nodal MEPs and MIPs are restricted to SPBM B-VIDs.
- The Maintenance level for MEPs and MIPs on a certain B-VID (in a network) must be configured to the same level for them to respond to a certain CFM command.

Limitations

When you configure CFM, be aware of the following configuration limitations:

- CFM does not support CCM messages.
- Only an autogenerated MEP and MIP can exist for C-VLANs.
- Only one MEP can exist on each C-VLAN and SPBM B-VLAN.
- Only one MIP can exist on each C-VLAN and SPBM B-VLAN.
- SMLT Virtual MAC for C-VLAN does not exist, so VSP 9000 does not support this option for l2 ping and l2 traceroute.
- VSP 9000 does not support l2tracetree on C-VLANs because no multicast tree exists on C-VLANs.
- The autogenerated MEPs do not have a uniqueness across the entire network until you configure the global MEP ID on each box to a different value. You must configure a unique MEP ID at a global level for CFM.
- MEPs and MIPs configured for SPBM VLANs do not respond to CFM messages sent from C-MAC VLANs because the VLAN and packet encapsulation are different.
- Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN.
- You can only configure global CFM at one MD level for each chassis for each VLAN type.

CFM configuration using ACLI

This section provides procedures to configure and use Connectivity Fault Management (CFM) using Avaya Command Line Interface (ACLI). The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults, which is performed at Layer 2, not Layer 3. To support troubleshooting of the SPBM cloud, Virtual Services Platform 9000 supports a subset of CFM functionality.

Note:

When you enable CFM in an SPBM network, Avaya recommends that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

Autogenerated CFM

CFM provides two methods for configuration: autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure a MD, MA, and MEP ID to create a MEP.

Note:

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN.

VSP 9000 only supports one MEP and one MIP, either autogenerated or explicitly configured, on the SPBM B-VLAN. Similarly, VSP 9000 only supports one MEP and one MIP on the C-VLAN. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN.

For autogenerated CFM configuration information for ACLI see:

- Configuring autogenerated CFM on SPBM B-VLANs on page 321.
- Configuring autogenerated CFM on C-VLANs on page 323.

Configuring autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This eliminates the need to explicitly configure an MD, MA, and MEP ID, and to associate the MEP and MIP level to the SPBM B-VLAN.

When you enable this feature, the device creates a global MD (named spbm) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1.

The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.

Important:

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs:

   ```
   cfm spbm level <0-7>
   ```

   You can change this level from the default of 4 either before or after the feature is enabled.

   Only configure global CFM at one MD level for each chassis for each VLAN type.

3. Assign a global CFM MEP ID for all CFM SPBM MEPs:

   ```
   cfm spbm mepid <1-8191>
   ```

4. Enable the autogenerated CFM for SPBM B-VLANs globally:

   ```
   cfm spbm enable
   ```

5. **(Optional)** Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs to the default:

   ```
   default cfm spbm level
   ```

6. **(Optional)** Assign a global CFM MEP ID for all CFM SPBM MEPs to the default:

   ```
   default cfm spbm mepid
   ```

7. **(Optional)** Disable the global CFM MEPs and MIPs:

   ```
   no cfm spbm enable
   ```

8. Display the global CFM MEP configuration:

   ```
   show cfm spbm
   ```

Example

Configure autogenerated CFM MEPs and MIPs:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#cfm spbm level 6
VSP-9012:1(config)#cfm spbm mepid 4
VSP-9012:1(config)#cfm spbm enable
VSP-9012:1(config)#show cfm spbm

LEVEL ADMIN    MEPID    MAC
```

```
=======================================
6         enable         4         00:15:e8:b8:a3:df
```

Variable definitions

Use the data in the following table to use the **cfm spbm** command.

Variable	Value
level<*0-7*>	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
mepid<*1-8191*>	Specifies the global MEP ID within the range of 1–8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
enable	Enables autogenerated CFM on all SPBM B-VLANs.

Job aid

The following table describes the fields for the **show cfm spbm** command.

Parameter	Description
LEVEL	Specifies the global SPBM CFM maintenance level for the chassis. The default is 4.
ADMIN	Specifies if CFM MEPs and MIPs are globally enabled.
MEP ID	Specifies the global MEP ID. The default is 1.
MAC	Specifies the MAC address.

Configuring autogenerated CFM on C-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every C-VLAN on the chassis.

Important:

CFM supports one MEP or MIP for each C-VLAN. Only autogenerated CFM provides support for configuring MEP and MIPs on C-VLANs. You cannot explicitly configure C-VLANs.

About this task

When you enable this feature, you create a global MD (named cmac) for all the customer MAC (C-MAC) MEPs. This global MD has a default maintenance level of 4, which you can change with the level attribute. The autogenerated CFM commands also create an MA for each C-VLAN, a MEP for each C-VLAN, associate the MEP with the corresponding C-VLAN, and a MIP with the C-VLAN.

All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The device automatically associates the MEPs with the C-VLANs configured.

The MIP level maps to the global level. The device automatically associates the MIP level with the C-VLANs configured when you enable the feature.

The feature is disabled by default.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the maintenance level for every CFM C-MAC MEP and MP level on all the C-VLANs:

   ```
   cfm cmac level <0-7>
   ```

   Only configure global CFM at one MD level for each chassis for each VLAN type.

3. Assign a global CFM MEP ID for all CFM C-MAC MEPs:

   ```
   cfm cmac mepid <1-8191>
   ```

4. Enable the autogenerated CFM for C-VLANs:

   ```
   cfm cmac enable
   ```

5. **(Optional)** Configure the maintenance level for every CFM C-MAC MEPs and MP level on all the C-VLANs to the default:

   ```
   default cfm cmac level
   ```

6. **(Optional)** Assign a global CFM MEP ID for all CFM C-MAC MEPs to the default:

   ```
   default cfm cmac mepid
   ```

7. **(Optional)** Disable the global CFM MEPs and MIPs:

   ```
   no cfm cmac enable
   ```

8. Display the global CFM MEP configuration:

   ```
   show cfm cmac
   ```

Example

Configure autogenerated CFM MEPs and MIP level:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#cfm cmac level 0
VSP-9012:1(config)#cfm cmac mepid 4
VSP-9012:1(config)#cfm cmac enable
VSP-9012:1(config)#show cfm cmac

LEVEL ADMIN    MEPID    MAC
=====================================
0        enable        4        00:15:e8:b8:a3:de
```

Variable definitions

Use the data in the following table for the **cfm cmac** command.

Variable	Value
level<*0-7*>	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
mepid<*1-8191*>	Specifies the global MEP ID within the range of 1–8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1. **Note:** The MA takes its name from this value for autogenerated CFM. For example, if you specify 500 as the MEP ID, the MA will also be 500.
enable	Enables autogenerated CFM for all C-MAC VLANs.

Job aid

The following table describes the fields for the **show cfm cmac** command.

Parameter	Description
LEVEL	Specifies the global C-VLAN CFM maintenance level for the chassis. The default is 4.
ADMIN	Specifies if CFM C-VLAN MEPs and MIPs are globally enabled.
MEP ID	Specifies the global MEP ID. The default is 1.
MAC	Specifies the MAC address.

Configuring explicit CFM

For SPBM B-VLANs, CFM provides two methods for configuration: autogenerated and explicit. You cannot use both. Use the procedures in this section to configure CFM explicitly. For C-VLANs, you can only use the autogenerated method.

If you want to create explicit CFM MEPs that require you to configure an MD, MA, and MEP ID, see the procedures in the following sections:

- Configuring CFM MD on page 326.
- Configuring CFM MA on page 327.
- Assigning a MEP and MIP level to an SPBM B-VLAN on page 328.
- Configuring CFM MEP on page 330.

Note:

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

Configuring CFM MD

Use this procedure to configure the Connectivity Fault Management (CFM) Maintenance Domain (MD) explicitly. An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

Note:

If you use autogenerated CFM, you do not configure CFM MD because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create the CFM MD:

   ```
   cfm maintenance-domain WORD<1-22> [index <1-2147483647>]
   [maintenance-level <0-7>] [level <0-7>]
   ```

3. Display the CFM MD configuration:

   ```
   show cfm maintenance-domain
   ```

4. Delete the CFM MD:

   ```
   no cfm maintenance-domain WORD<1-22>
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# cfm maintenance-domain md1 index 99 maintenance-level
3
VSP-9012:1(config)# show cfm maintenance-domain
================================================================================
                         Maintenance Domain
================================================================================
Domain Name             Domain Index    Level Domain Type
--------------------------------------------------------------------------------
md1                     99              3     NONE

 Total number of Maintenance Domain entries: 1.
VSP-9012:1(config)# no cfm maintenance-domain md1
VSP-9012:1(config)#show cfm maintenance-domain
```

```
================================================================================
                          Maintenance Domain
================================================================================
Domain Name             Domain Index    Level Domain Type
--------------------------------------------------------------------------------

 Total number of Maintenance Domain entries: 0.
```

Variable definitions

Use the data in the following table to use the **cfm maintenance-domain** command.

Variable	Value
WORD<1–22>	Specifies the maintenance domain name.
index *<1–2147483647>*	Specifies a maintenance domain entry index.
maintenance-level *<0–7>*	Specifies the MD maintenance level when creating the MD. The default is 4.
level *<0–7>*	Modifies the MD maintenance level for an existing MD. The default is 4.

Configuring CFM MA

Use this procedure to configure the CFM Maintenance Association (MA) explicitly. An MA represents a logical grouping of monitored entities within its domain. It can therefore represent a set of Maintenance Association End Points (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

Note:

If you use autogenerated CFM, you do not configure CFM MA because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Create the CFM MA:

   ```
   cfm maintenance-association WORD<1-22> WORD<1-22> [index <1-
   2147483647>]
   ```

3. Display the CFM MA configuration:

   ```
   show cfm maintenance-association
   ```

4. Use the following command, if you want to delete the CFM MA:

   ```
   no cfm maintenance-association WORD<1-22> WORD<1-22>
   ```

Example

```
VSP-9012:1> enable
```

```
VSP-9012:1# configure terminal
VSP-9012:1(config)# cfm maintenance-association md1 ma1 index 98
VSP-9012:1(config)# show cfm maintenance-association
================================================================================
                        Maintenance Association Status
================================================================================
Domain Name             Assn Name              Domain Idx  Assn Idx
--------------------------------------------------------------------------------
md1                     ma1                    1           98

 Total number of Maintenance Association entries: 1.

================================================================================
                        Maintenance Association config
================================================================================
Domain Name             Assn Name
--------------------------------------------------------------------------------
md1                     ma1

 Total number of MA entries: 1.
```

Variable definitions

Use the data in the following table to use the **`cfm maintenance-association`** command.

Variable	Value
WORD<1–22> WORD<1–22>	Creates the CFM MA. The first parameter, specifies the MD name. The second parameter, specifies the MA name.
index *<1–2147483647>*	Specifies a maintenance association entry index.

Assigning a MEP and MIP level to an SPBM B-VLAN

Use this procedure to assign a nodal MEP to an SPBM B-VLAN. The Nodal MEP provides traceability and troubleshooting at the system level for a specific B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the specific B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a specific MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Note:

If you use autogenerated CFM, you do not configure CFM MIP/MEP because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Add nodal MEPs to the B-VLAN:

```
vlan nodal-mep <1-4084> WORD<0-22> WORD<0-22> <1-8191>
```

3. Display the nodal MEP configuration:

```
show vlan nodal-mep <1-4084>
```

4. Remove the nodal MEPs from the B-VLAN:

```
no vlan nodal-mep <1-4084> WORD<0-22> WORD<0-22> <1-8191>
```

5. Add nodal MIP level to the B-VLAN:

```
vlan nodal-mip-level <1-4084> WORD<0-15>
```

6. Display the nodal MIP level configuration:

```
show vlan nodal-mip-level [<1-4084>]
```

7. Remove the nodal MIP level from the B-VLAN:

```
no vlan nodal-mip-level <1-4084> WORD<0-15>
```

Example

```
VSP-9012:1> enable
VSP-9012:1> configure terminal
VSP-9012:1> vlan nodal-mep 100 md1 ma1 2
VSP-9012:1> show vlan nodal-mep
```

```
================================================================================
                                Vlan Nodal Mep
================================================================================
VLAN_ID    DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
--------------------------------------------------------------------------------
1
100         md1.ma1.2
216
304
404
500
616
716
816
916
1000       spbm.1000.1
1001       spbm.1001.1
```

```
VSP-9012:> vlan nodal-mip 100 6
VSP-9012:> show vlan nodal-mip
```

```
================================================================================
                            Vlan Nodal Mip Level
================================================================================
VLAN_ID    NODAL_MIP_LEVEL_LIST
```

```
--------------------------------------------------------------------------------
1
100         6
216
304
404
500
616
716
816
916
1000        4
1001        4
```

Variable definitions

Use the data in the following table to use the **`vlan nodal-mep`** command.

Variable	Value
<1–4084>	Specifies the VLAN ID.
WORD<0–22>	The first parameter, specifies the Maintenance Domain name.
WORD<0–22>	The second parameter, specifies the Maintenance Association name.
<1–8191>	Specifies the nodal MEPs to add to the VLAN.

Use the data in the following table to use the **`vlan nodal-mip-level`** command.

Variable	Value
<1–4084>	Adds the nodal MIP level. Specifies the VLAN ID.
WORD<0–15>	Adds the nodal MIP level, which has up to eight levels, ranging from 0 to 7.

Configuring CFM MEP

Use this procedure to configure the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

Note:

If you use autogenerated CFM, you do not configure CFM MEPs because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create the CFM MEP:

```
cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191> [state
<enable>]
```

3. Enable an existing CFM MEP:

   ```
   cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191> enable
   ```

4. Disable an existing CFM MEP:

   ```
   no cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191> enable
   ```

5. Display the CFM MEP configuration:

   ```
   show cfm maintenance-endpoint
   ```

6. Delete an existing CFM MEP:

   ```
   no cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191>
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# cfm maintenance-endpoint md1 ma1 1 state enable
VSP-9012:1> show cfm maintenance-endpoint
================================================================================
                        Maintenance Endpoint Config
================================================================================
DOMAIN                  ASSOCIATION             MEP  ADMIN
NAME                    NAME                    ID
--------------------------------------------------------------------------------
md1                     ma1                     1    enable

 Total number of MEP entries: 1.

================================================================================
                        Maintenance Endpoint Service
================================================================================
DOMAIN_NAME             ASSN_NAME               MEP_ID TYPE    SERVICE_DESCRIPTION
--------------------------------------------------------------------------------
md1                     ma1                     1      unused

 Total number of MEP entries: 1.
```

Variable definitions

Use the data in the following table to use the **`cfm maintenance-endpoint`** command.

Variable	Value
WORD<1–22>	The first parameter, specifies the MD name.
WORD<1–22>	The second parameter, specifies the MA name.
<1–8191>	Specifies the MEP ID.
state {enable | disable}	Enables or disables the MEP when creating the MEP. The default is disabled.

Variable	Value
enable	Enables an existing MEP. Use this parameter with the no option to disable an existing MEP.

Triggering a loopback test (LBM)

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

Before you begin

- You must have a MEP that is associated with a B-VLAN or a C-VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger the loopback test:

```
loopback WORD<1-22> WORD<1-22> <1-8191>
<0x00:0x00:0x00:0x00:0x00:0x00> [burst-count <1-200>] [data-tlv-size
<0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode
{nodal|noVlanMac|smltVirtual}] [testfill-pattern <all-zero|all-zero-
crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>]
[time-out <1-10>]
```

Example

```
VSP-9012:1# loopback md1 4001 13 00:14:0D:A2:B3:DF burst-count 10 priority
3 time-out 5
```

```
Result of LBM from mep: md1.4002.13 to MAC address: 00:14:0D:A2:B3:DF :
    Sequence number of the first LBM is 10575
    The total number of LBMs sent out is 10
    The number of LBRs received is 10
    The number of LBRs lost is 0
    The percentage of LBMs lost is 0.00%
    The RTT Min is 764 microsecs, Max is 800 microsecs, Average is 783.00 microsecs
    The RTTDV min is 3 microsecs, Max is 23 microsecs, Average is 9.11 microsecs
    The Standard Deviation of RTT is 11.53 microsecs
```

Variable definitions

Use the data in the following table to use the **loopback** command.

Variable	Value
WORD<1–22>	The first parameter, specifies the MD name.
WORD<1–22>	The second parameter, specifies the MA name.
<1–8191>	Specifies the MEP ID.

Variable	Value
<0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the remote MAC address to reach the MEP/MIP.
burst-count *<1–200>*	Specifies the burst-count.
data-tlv-size *<0–400>*	Specifies the data TLV size.
frame-size *<64–1500>*	Specifies the frame-size. The default is 0.
priority *<0–7>*	Specifies the priority. The default is 7.
source-mode *{nodal|noVlanMac|smltVirtual}*]	Specifies the source mode: • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual—Use this value with B-VLANs only. The default is nodal.
testfill-pattern *{all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc}*	Specifies the testfill pattern: • all-zero — null signal without cyclic redundancy check • all-zero-crc — null signal with cyclic redundancy check with 32-bit polynomial • pseudo-random-bit-sequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudo-random-bit-sequence-crc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default is 1: all-zero.
time-out *<1–10>*	Specifies the time-out interval in seconds. The default is 3.

Triggering linktrace (LTM)

Use the following procedure to trigger a linktrace.

The Linktrace Message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virutal SMLT MAC. MPs on the path to the target address respond with an LTR.

Before you begin

- You must have a MEP that is associated with a B-VLAN or C-VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.

2. Trigger the linktrace:

```
linktrace WORD<1-22> WORD<1-22> <1-8191>
<0x00:0x00:0x00:0x00:0x00:0x00> [detail] [priority <0-7>] [source-
mode <nodal|noVlanMac|smltVirtual>] [ttl-value <1-255>]
```

Example

```
VSP-9012:1# linktrace md1 4001 13 00:bb:00:00:14:00 priority 7

Please wait for LTM to complete or press any key to abort

Received LTRs:

SeqNum: 10575 MD: md1 MA:4001    MepId:     13     Priority: 7
-------------------------------------------------------------
TTL SRC MAC                    FWDYES    TERMMEP    RELAY        ACTION
--- ------------------ ------  -------- ----------------------
63    00:bb:00:00:10:00         true        false        Fdb
62     00:bb:00:00:14:00      false        true        Hit
```

Variable definitions

Use the data in the following table to use the **linktrace** command.

Variable	Value
WORD<1–22>	The first parameter, specifies the MD name.
WORD<1–22>	The second parameter, specifies the MA name.
<1–8191>	Specifies the MEP ID.
<0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the target MAC address to reach the MEP.
detail	Displays linktrace result details.
priority *<0–7>*	Specifies the priority. The default is 7.
source-mode *<nodal|noVlanMac|smltVirtual>*	Specifies the source mode: • 1: nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • 2: smltVirtual—Use this value with B-VLANs only. The default is 1: nodal.
ttl-value *<1–255>*	Specifies the Time-to-Live value. The default is 64.

Triggering a Layer 2 ping

Use this procedure for C-VLANs or B-VLANs to trigger a Layer 2 ping, which acts like native **ping**. This feature enables CFM to debug Layer 2. Layer 2 ping can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

Note:

To use Layer 2 ping for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the **cfm cmac enable** command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the **l2 ping** command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **l2 ping** command to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **l2 ping** and **l2 traceroute** requests.

Before you begin

- You must configure and enable CFM.
- You must have a MEP that is associated with a VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger a Layer 2 ping:

```
l2 ping {vlan <1-4084>} {routernodename WORD<0-255> | mac
<0x00:0x00:0x00:0x00:0x00:0x00>} [burst-count <1-200>] [data-tlv-
size <0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode
<nodal|noVlanMac|smltVirtual>] [testfill-pattern <all-zero|all-zero-
crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>]
[time-out <1-10>]
```

```
l2 ping {ip-address WORD<0-255>} [burst-count <1-200>] [data-tlv-
size <0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode
<nodal|noVlanMac|smltVirtual>] [testfill-pattern <all-zero|all-zero-
crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>]
[time-out <1-10>] [vrf WORD<1-16>]
```

Example

Trigger a Layer 2 ping for MAC address 00.14.0d.bf.a3.d:

```
VSP-9012:1> l2 ping vlan 500 mac 00.14.0d.bf.a3.d
Please wait for l2ping to complete or press any key to abort
----00:14:0d:bf:a3:df  L2 PING Statistics----  0(64) bytes of data
1 packets transmitted, 0 packets received,   100.00% packet loss
```

Trigger a Layer 2 ping for router node name VSP-MONTIO:

```
VSP-9012:1> l2 ping vlan 500 routernodename VSP-MONTI0
Please wait for l2ping to complete or press any key to abort
```

```
----00:14:0d:a2:b3:df    L2 PING Statistics----  0(64) bytes of data
1 packets transmitted, 1 packets received,  0.00% packet loss
  round-trip (us)        min/max/ave/stdv =  26895/26895/26895.00/ 0.00
```

Trigger a Layer 2 ping for IP address 192.0.2.102:

```
VSP-9012:1>l2 ping ip-address 192.0.2.102

Please wait for l2ping to complete or press any key to abort

L2 PING  Statistics : IP 192.0.2.102, paths found 1, path attempted 1
================================================================================
                                          TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                             PKTS  PKTS  LOSS     MIN/MAX/AVE (us)
================================================================================

50   80:17:7d:75:aa:02  (80:17:7d:75:aa:02)  1     0     100.00%  0/0/0.00
```

Variable definitions

Use the data in the following table to configure the **`l2 ping`** command.

Variable	Value
{vlan *<1–4084>* routernodename *WORD<0–255>*} (vlan *<1–4084>* mac *<0x00:0x00:0x00:0x00:0x00:0x00>*} {ip-address *WORD<0–255>*}	Specifies the destination for the L2 ping: • <1–4084> — Specifies the VLAN ID. • WORD<0–255> — Specifies the Router node name. • <XX:XX:XX:XX:XX:XX> — Specifies the MAC address. • <A.B.C.D> — Specifies the IP address. **Note:** VSP 9000 does not support the routernodename option for C-VLANs.
burst-count *<1–200>*	Specifies the burst count.
data-tlv-size *<0–400>*	Specifies the data TLV size. The default is 0.
frame-size *<64–1500>*]	Specifies the frame size. The default is 0.
testfill-pattern *<all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>*	Specifies the testfill pattern: • all-zero — Null signal without cyclic redundancy check. • all-zero-crc — Null signal with cyclic redundancy check with 32-bit polynomial. • pseudo-random-bit-sequence — Pseudo-random-bit-sequence without cyclic redundancy check. • pseudo-random-bit-sequence-crc — Pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default is all-zero.
priority *<0–7>*]	Specifies the priority. The default is 7.
time-out *<1–10>*	Specifies the interval in seconds. The default is 3.

Variable	Value
source-mode <*nodal*|*noVlanMac*|*smltVirtual*>	Specifies the source mode: • nodal • noVlanMac — Use this value with C-VLANs only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual—Use this value with B-VLANs only. The default is nodal.
vrf *WORD<1–16>*	Specifies the VRF name.

Triggering a Layer 2 traceroute

Use this procedure for B-VLANs or C-VLANs to trigger a Layer 2 traceroute, which acts like native **`traceroute`**. This feature enables CFM to debug Layer 2 for SPBM B-VLANs or C-VLANs. Layer 2 traceroute can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

Note:

To use Layer 2 traceroute for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the **`cfm cmac enable`** command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the **`l2 traceroute`** command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **`l2 traceroute`** command to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **`l2 ping`** and **`l2 traceroute`** requests.

Before you begin

- You must configure and enable CFM.
- You must have a MEP that is associated with a VLAN.

About this task

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an **`l2 ping`** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.
- For C-VLANs, you have to trigger an **`l2 ping`** to learn the C-MAC address.

In both cases, **`linktrace`** traces the path up to the closest device to that MAC address that supports CFM.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger a Layer 2 traceroute:

```
l2 traceroute {vlan <1-4084> routernodename WORD<0-255> | vlan <1-4084> mac <0x00:0x00:0x00:0x00:0x00:0x00>} [priority <0-7>] [source-mode <nodal|noVlanMac|smltVirtual>] [ttl—value <1-255>]
```

```
l2 traceroute {ip-address WORD<0-255>} [priority <0-7>] [source-mode <nodal|noVlanMac|smltVirtual>] [ttl—value <1-255>] [vrf WORD<1-16]
```

Example

Trigger a Layer 2 traceroute for VLAN 500 with the router node name VSP-MONTIO:

```
VSP-9012:1#l2 traceroute vlan 500 routernodename VSP-MONTIO

Please wait for l2traceroute to complete or press any key to abort

l2traceroute to VSP-MONTI0  (00:14:0d:a2:b3:df),  vlan 500
0    VSP-PETER4         (00:15:9b:11:33:df)
1    VSP-MONTI0          (00:14:0d:a2:b3:df)
```

Trigger a Layer 2 traceroute for the IP address 192.0.2.1:

```
VSP-9012:1#l2 traceroute ip-address 192.0.2.1

Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 192.0.2.1, paths found 1
================================================================
VSP-SHAMIM  (00:1a:8f:08:53:df),  vlan 500
0    VSP-PETER4          (00:15:9b:11:33:df)
1    VSP-MONTI0          (00:14:0d:a2:b3:df)
```

The output for the l2 traceroute using C-VLAN 10 to target MAC 00:14:9b:11:30:00.

```
VSP-9012:1# l2 traceroute 10.00:14:9b:11:30:00

Please wait for l2traceroute to complete or press any key to abort
l2traceroute to   00:14:9b:11:30:00,  vlan 10
0 00:15:9b:11:30:00    (00:15:9b:11:30:00)
1 00:14:9b:11:30:00    (00:14:9b:11:30:00)
```

Variable definitions

Use the data in the following table to use the **`l2 traceroute`** command.

Variable	Value
{vlan *<1–4084>* routernodename *WORD<0–255>*} (vlan *<1–4084>* mac *<0x00:0x00:0x00:0x00:0x00:0x00>*} {ip-address *WORD<0–255>*}	Specifies the destination for the L2 traceroute: • <1–4084> — Specifies the VLAN ID. • WORD<0–255> — Specifies the router node name. • <XX:XX:XX:XX:XX:XX> — Specifies the MAC address. • WORD<0–255> — Specifies the IP address. **Note:** VSP 9000 does not support the routernodename option for C-VLANs.
ttl-value*<1–255>*	Specifies the TTL value. The default is 64.
priority *<0–7>*	Specifies the priority. The default is 7.
source-mode *<nodal|noVlanMac|smltVirtual>*	Specifies the source mode: • nodal • noVlanMac — Use this value with C-VLANs only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual— Use this value with B-VLANs only. The default is nodal.
vrf *WORD<1–16>*	Specifies the VRF name.

Triggering a Layer 2 tracetree

Use this procedure to trigger a Layer 2 tracetree. Layer 2 tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Note:

VSP 9000 only supports this command on SPBM B-VLANs only, not C-VLANs.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.

2. Trigger a Layer 2 tracetree:

```
l2 tracetree {<1-4084> <1-16777215> [routernodename WORD<0-255> |
<1-4084> <1-16777215>] [mac <0x00:0x00:0x00:0x00:0x00:0x00>}]
[priority <0-7>] [source-mode <nodal|smltVirtual>] [ttl—value <1-
255>]
```

Example

```
VSP-9012:1# l2 tracetree 500 1

Please wait for l2tracetree to complete or press any key to abort


l2tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10
hops 64
1   VSP-PETER4        00:15:9b:11:33:df -> VSP-MONTI0        00:14:0d:a2:b3:df
2   VSP-MONTI0        00:14:0d:a2:b3:df -> VSP-LEE2          00:15:e8:b8:a3:df
```

Variable definitions

Use the data in the following table to use the **`l2 tracetree`** command.

Variable	Value
{*<1–4084> <1–16777215> routernodename WORD<0–255>* | *<1–4084> <1–16777215> mac <0x00:0x00:0x00:0x00:0x00:0x00>*}	• <1–4084> — Specifies the VLAN ID. • <1–16777215> — Specifies the I-SID. • WORD<0–255> — Specifies the Router Node Name. • <0x00:0x00:0x00:0x00:0x00:0x00> — Specifies the MAC address.
ttl-value*<1–255>*	Specifies the TTL value. The default is 64.
priority *<0–7>*	Specifies the priority value. The default is 7.
source-mode *<nodal|smltVirtual>*	Specifies the source mode. • 1: nodal • 2: smltVirtual The default is nodal.

Triggering a Layer 2 tracemroute

Use this procedure to debug the IP multicast over SPBM stream path using **`l2 tracemroute`** on the VLAN (Layer 2) or the VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

Note:

The VLAN option is only valid for a VLAN that has an I-SID configured and IGMP snooping enabled.

Before you begin

- On the source and destination nodes, you must configure an autogenerated or an explicit CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger a Layer 2 tracemroute on the VLAN:

   ```
   l2 tracemroute source <A.B.C.D> group <A.B.C.D> vlan
   <1-4084>[priority <0-7>] [ttl—value <1-255>]
   ```

   Note:

   For the preceding command, if you do not specify a VLAN, **`l2 tracemroute`** uses the global default VRF.

   Wait for the l2 tracemroute to complete or press any key to abort.
3. Trigger a Layer 2 tracemroute on the VRF:

   ```
   l2 tracemroute source <A.B.C.D> group <A.B.C.D> vrf WORD<1-16>
   [priority <0-7>] [ttl—value <1-255>]
   ```

   Note:

   For the preceding command, if you do not specify a VRF, **`l2 tracemroute`** uses the global default VRF.

   Wait for the l2 tracemroute to complete or press any key to abort.

Example

The following is a sample output for a Layer 2 tracemroute on a VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#l2 tracemroute source 192.0.2.81 233.252.0.1 vlan 201

Please wait for l2 tracemroute to complete or press any key to abort.

Source 192.0.2.81

Group: 233.252.0.1

VLAN:201

BMAC: 03:00:03:f4:24:01

B-VLAN: 10

I-SID: 16000001

================================================================
1 VSP-PETER4 00:03:00:00:00:00 -> VSP-LEE1 00:14:0d:bf:a3:df
2 VSP-LEE1 00:14:0d:bf:a3:df -> VSP-LEE2 00:15:e8:b8:a3:df
```

The following is a sample output for a Layer 2 tracemroute on a VRF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#l2 tracemroute source 192.0.2.10 group 233.252.0.1 vrf red

Please wait for l2 tracemroute to complete or press any key to abort.

Source 192.0.2.10

Group: 233.252.0.1

VRF: redID 1

BMAC: 03:00:04:f4:24:01

B-VLAN: 20

I-SID: 16000001

================================================================
1 VSP-PETER4 00:03:00:00:00:00 -> VSP-LEE1 00:14:0d:bf:a3:df
2 VSP-LEE1 00:14:0d:bf:a3:df -> VSP-LEE2 00:15:e8:b8:a3:df
```

Variable definitions

Use the data in the following table to use the **`l2 tracemroute`** command.

Variable	Value
source *<A.B.C.D>*	Specifies the source IP address.
group *<A.B.C.D>*	Specifies the IP address of the multicast group.
vlan *<1–4084>*	Specifies the VLAN value.
vrf *WORD<1–16>*	Specifies the VRF name. If you do not specify a VRF name, then the results are shown for the flow in the Global Router (default) context.
priority *<0–7>*	Specifies the priority value.
ttl *<1–255>*	Specifies the time-to-live (TTL) for the trace packet, which is how many hops the trace packet takes before it is dropped.

Job aid

The following table describes the fields in the output for **`l2 tracemroute`** command for a VLAN.

Parameter	Description
Source	Specifies the source IP address of the flow where the multicast trace tree originates.
Group	Specifies the IP address of the multicast group.
VLAN	Specifies the VLAN.
BMAC	Specifies the backbone MAC address.
B-VLAN	Specifies the backbone VLAN.
I-SID	Specifies the service identifier.

The following table describes the fields in the output for **l2 tracemroute** command for a VRF.

Parameter	Description
Source	Specifies the source IP address of the flow where the multicast trace tree originates.
Group	Specifies the IP address of the multicast group.
VRF	Specifies the VRF.
BMAC	Specifies the backbone MAC address.
B-VLAN	Specifies the backbone VLAN.
I-SID	Specifies the service identifier.

Using trace CFM to diagnose problems

Use the following procedure to display trace information for CFM.

About this task

Use trace to observe the status of a software module at a certain time.

For example, if you notice a CPU utilization issue (generally a sustained spike above 90%) perform a trace of the control plane activity.

Use the **trace level 120 <0–4>** command to trace CFM module information, including ACLI, instrumentation, High Availability, show config, and platform dependent code. The CFM module ID is 120.

Use the **trace cfm level <0–4>** command to trace platform independent code and CFM protocol code.

Caution:

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Clear the trace:

   ```
   clear trace
   ```

3. Begin the trace operation:

   ```
   trace cfm level <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

4. Stop tracing:

   ```
   trace shutdown
   ```

5. View the trace results:

   ```
   show trace cfm
   ```

6. Begin the trace operation for the CFM module:

   ```
   trace level 120 <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

7. View trace results:

   ```
   trace screen enable
   ```

   Important:

   If you use trace level 3 (verbose) or trace level 4 (very verbose), Avaya recommends you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

8. Save the trace file to the Compact Flash card for retrieval.

   ```
   save trace [file WORD<1-99>]
   ```

   If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

9. Search trace results for a specific string value, for example, the word error:

   ```
   trace grep [WORD<0-128>]
   ```

   If you use this command and do not specify a string value, you clear the results of a previous search.

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)# clear trace
VSP-9012:1(config)# trace cfm level 3
VSP-9012:1(config)# trace shutdown
VSP-9012:1(config)# show trace cfm
================================================================================
                                CFM Tracing Info
================================================================================
 Status    : Enabled
 Level     : VERBOSE
VSP-9012:1(config)#trace level 120 3
VSP-9012:1(config)# save trace
VSP-9012:1(config)# trace grep error
VSP-9012:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

Use the data in the following table to use the **trace** command.

Table 5: Variable definitions

Variable	Value
cfm level *[<0-4>]*	Starts the trace by specifying the level. • *<0–4>* specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.
grep *[WORD<0-128>]*	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level *<0–192>[<0-4>]*	Starts the trace by specifying the module ID and level. • *<0–192>* specifies the module ID. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.
shutdown	Stops the trace operation.
screen {disable|enable}	Enables the display of trace output to the screen.

Use the data in the following table to use the **`save trace`** command.

Table 6: Variable definitions

Variable	Value
file *WORD<1–99>*	Specifies the file name in one of the following formats: • a.b.c.d:<file> • x:x:x:x:x:x:x:x <file> • /intflash/<file> • /extflash/<file> • /usb/<file> • /mnt/intflash/ <file> • /mnt/extflash/ <file> /mnt/intflash is the internal flash of the second CP module (the one to which you are not connected). /mnt/extflash is the external flash of the second CP module (the one to which you are not connected).

Using trace SPBM to diagnose problems

Use the following procedure to display trace information for SPBM IS-IS. In the case of IS-IS, this procedure also provides information related to the flags set.

About this task

Use the **trace level 119 <0–4>** command to trace IS-IS module information, including ACLI, instrumentation, High Availability, show config and platform dependent code. The IS-IS module ID is 119.

Use the **trace level 125 <0–4>** command to trace SPBM module information, including ACLI, instrumentation, High Availability, show config and platform dependent code. The SPBM module ID is 125.

Use the **trace spbm isis level** command to trace platform independent code, IS-IS protocol, IS-IS hello, IS-IS adjacency, LSP processing, and IS-IS computation.

Caution:

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Clear the trace:

   ```
   clear trace
   ```

3. Begin the trace operation:

   ```
   trace spbm isis level <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

4. Stop tracing:

   ```
   trace shutdown
   ```

5. Display the trace information for SPBM IS-IS:

   ```
   show trace spbm isis
   ```

6. Begin the trace operation for the SPBM module:

   ```
   trace level 125 <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

7. Begin the trace operation for the IS-IS module:

   ```
   trace level 119 <0-4>
   ```

   Wait approximately 30 seconds, and then stop trace.

8. View trace results:

   ```
   trace screen enable
   ```

Important:

If you use trace level 3 (verbose) or trace level 4 (very verbose), Avaya recommends you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

9. Save the trace file to the Compact Flash card for retrieval.

   ```
   save trace [file WORD<1-99>]
   ```

   If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

10. Search trace results for a specific string value, for example, the word error:

    ```
    trace grep [WORD<0-128>]
    ```

    If you use this command and do not specify a string value, you clear the results of a previous search.

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)# clear trace
VSP-9012:1(config)# trace spbm isis level 3
VSP-9012:1(config)# trace shutdown
VSP-9012:1(config)# show trace spbm isis
================================================================================
                              SPBM ISIS Tracing Info
================================================================================
 Status    : Enabled
 Level     : VERY_TERSE
 Flag Info :
VSP-9012:1(config)#trace level 125 3
VSP-9012:1(config)#trace level 119 3
VSP-9012:1(config)# save trace
VSP-9012:1(config)# trace grep error
VSP-9012:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

Use the data in the following table to use the **trace** command.

Table 7: Variable definitions

Variable	Value
grep *[WORD<0-128>]*	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level *<0–192>[<0-4>]*	Starts the trace by specifying the module ID and level. • *<0–192>* specifies the module ID. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.

Variable	Value
spbm isis level *[<0-4>]*	Starts the trace by specifying the level. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose. The default is 1, very terse.
shutdown	Stops the trace operation.
screen {disable|enable}	Enables the display of trace output to the screen.

Use the data in the following table to use the **`save trace`** command.

Table 8: Variable definitions

Variable	Value
file *WORD<1–99>*	Specifies the file name in one of the following formats: • a.b.c.d:<file> • x:x:x:x:x:x:x:x <file> • /intflash/<file> • /extflash/<file> • /usb/<file> • /mnt/intflash/ <file> • /mnt/extflash/ <file> /mnt/intflash is the internal flash of the second CP module (the one to which you are not connected). /mnt/extflash is the external flash of the second CP module (the one to which you are not connected).

CFM configuration using EDM

This section provides procedures to configure Connectivity Fault Management (CFM) using Enterprise Device Manager (EDM).

Note:

When you enable CFM in an SPBM network, Avaya recommends that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

Autogenerated CFM

CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure an MD, MA, and MEP ID to create a MEP.

Note:

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN. VSP 9000 only supports one MEP or MIP on the SPBM B-VLAN, either explicitly configured or autogenerated.

For autogenerated CFM configuration information for EDM see:

- Configuring autogenerated CFM on SPBM B-VLANs on page 349.
- Configuring autogenerated CFM on C-VLANs on page 351.

Configuring autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This eliminates the need to explicitly configure an MD, MA, and MEP ID and to associate the MEP and MIP level to the SPBM B-VLAN.

To configure autogenerated CFM on C-VLANs, see Configuring autogenerated CFM on C-VLANs on page 351.

About this task

When you enable this feature, the device creates a global MD (named spbm) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.

Important:

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **CFM**.

3. Click the **Global** tab.
4. Click **enable** next to **SpbmAdminState**.
5. Specify a maintenance level in the **SPBMLevel** field.
6. Specify an MEP ID in the **SpbmMepId** field.
7. Click **Apply**.

CFM Global field descriptions

Use the data in the following table to configure the global MEP and MIP parameters.

Name	Description
SpbmAdminState	Enables or disables autogenerated CFM for B-VLANs. The default is disable. Enabling **SpbmAdminState** creates one MIP level and one MEP on every B-VLAN at the specified SpbmLevel.
SpbmLevel	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each B-VLAN type.
SpbmMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
CmacAdminState	Enables or disables autogenerated CFM for C-VLANs. The default is disable. Enabling **CmacAdminState** creates one MIP level and one MEP on every C-VLAN at the specified CmacLevel
CmacLevel	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each C-VLAN type.
CmacMepId	Specifies the global CFM CMAC MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
Bmac	This is read-only only field. Specifies the B-MAC address of the node.
Cmac	This is read-only only field. Specifies the C-MAC address of the node.

Configuring autogenerated CFM on C-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every C-VLAN on the chassis.

To configure autogenerated CFM on SPBM B-VLANs, see Configuring autogenerated CFM on SPBM B-VLANs on page 349.

Important:

CFM supports only one MEP and one MIP per C-VLAN. The only method to configure CFM MEPs and MIPs on C-VLANs is to use the simplified commands above which autogenerate MEPs and MIP on C-VLANs. You cannot explicitly configure on some VLANs as this is possible with SPBM CFM.

About this task

When you enable this feature, you create a global MD (named cmac) for all the customer MAC (C-MAC) MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. The autogenerated CFM commands also create an MA for each C-VLAN, a MEP for each C-VLAN, and associate the MEP with the corresponding C-VLAN and a MIP with the C-VLAN.

All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The device automatically associates the MEPs with the C-VLANs configured. The MIP level maps to the global level. The device automatically associates the MIP level with the C-VLANs configured when you enable the feature.

The feature is disabled by default.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **CFM**.
3. Click the **Global** tab.
4. Click **enable** next to CmacAdminState.
5. Specify a C-MAC CFM maintenance level in the **CmacLevel** field.
6. Specify an MEP ID in the **CmacMepId** field.
7. Click **Apply**.

CFM Global field descriptions

Use the data in the following table to configure the global MEP and MIP parameters.

Name	Description
SpbmAdminState	Enables or disables autogenerated CFM for B-VLANs. The default is disable.
SpbmLevel	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4.

Name	Description
	Only configure global CFM at one MD level for each chassis for each VLAN type.
SpbmMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
CmacAdminState	Enables or disables autogenerated CFM for C-VLANs. The default is disable.
CmacLevel	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
CmacMepId	Specifies the global CFM MEP within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
Bmac	Specifies the B-MAC address of the node.
Cmac	Specifies the C-MAC address of the node.

Configuring explicit CFM

For SPBM B-VLANs, CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. Use the procedures in this section to configure MEPs explicitly.

If you want to create autogenerated CFM MEPs that eliminate the need to configure an MD, MA, and MEP ID, see the procedures in Autogenerated CFM on page 349. For C-VLANs, you can only use the autogenerated method.

Note:

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

For explicit configuration information for EDM see:

- Configuring CFM MD on page 352.
- Configuring CFM MA on page 353.
- Configuring CFM MEP on page 354.
- Configuring CFM nodal MEP on page 355.

Configuring CFM MD

Use this procedure to configure a Connectivity Fault Management (CFM) Maintenance Domain (MD). An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

Note:

If you use autogenerated CFM, you do not configure CFM MD because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **CFM**.
3. Click the **MD** tab.
4. Click **Insert**.
5. In the fields provided, specify an index value, name, and level for the MD.
6. Click **Insert**.

MD field descriptions

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

Name	Description
Index	Specifies a maintenance domain entry index.
Name	Specifies the MD name.
NumOfMa	Indicates the number of MAs that belong to this maintenance domain.
Level	Specifies the MD maintenance level. The default is 4.
NumOfMip	Indicates the number of MIPs that belong to this maintenance domain
Type	Indicates the type of domain.

Configuring CFM MA

Use this procedure to configure a CFM Maintenance Association (MA). An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance Endpoints (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

Note:

If you use autogenerated CFM, you do not configure CFM MA because the switch configures a default MD, MA, MEPs, and MIPs.

Before you begin

- You must configure a CFM MD.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.

2. Click **CFM**.
3. Click the **MD** tab.
4. Highlight an existing MD, and then click **MaintenanceAssociation**.
5. In the **MA** tab, click **Insert**.
6. In the fields provided, specify an index value and name for the MA.
7. Click **Insert**.

MA field descriptions

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

Name	Description
DomainIndex	Specifies the maintenance domain entry index.
AssociationIndex	Specifies a maintenance association entry index.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
NumOfMep	Indicates the number of MEPs that belong to this maintenance association.

Configuring CFM MEP

Use this procedure to configure the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

Note:

If you use autogenerated CFM, you do not configure CFM MEPs because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **CFM**.
3. Click the **MD** tab.
4. Highlight an existing MD, and then click **MaintenanceAssociation**.
5. In the **MA** tab, highlight an existing MA, and then click **MaintenanceEndpoint**.
6. Click **Insert**.
7. In the fields provided, specify the ID and the administrative state of the MEP.
8. Click **Insert**.

MEP field descriptions

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

Name	Description
DomainIndex	Specifies the MD index.
AssociationIndex	Specifies the MA index.
Id	Specifies the MEP ID.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
AdminState	Specifies the administrative state of the MEP. The default is disable.
MepType	Specifies the MEP type: • trunk • sg • endpt • vlan • port • endptClient • nodal • remotetrunk • remotesg • remoteendpt • remoteVlan • remotePort • remoteEndptClient **Note:** VSP products only support Nodal Mep Type.
ServiceDescription	Specifies the service to which this MEP is assigned.

Configuring CFM nodal MEP

Use this procedure to configure the CFM nodal Maintenance Endpoint (MEP). The Nodal MEP provides traceability and troubleshooting at the system level for a specific B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the specific B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and Maintenance Intermediate Point (MIP) functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a specific MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Note:

If you use autogenerated CFM, you do not configure CFM nodal MEPs because the switch configures a default MD, MA, MEPs, and MIPs.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Advanced** tab.
4. Select an SPBM VLAN.
5. Click **Nodal**.
6. In the **NodalMepList** field, specify the nodal MEPs to add to the VLAN.
7. Click **Apply**.

Nodal MEP/MIP field descriptions

Use the data in the following table to use the **Nodal MEP/MIP** tab.

Name	Description
NodalMepList	Specifies the nodal MEPs to add to the VLAN, in the format <mdName.maName.mepId>, for example md10.ma20.30.
NumOfNodalMep	Indicates the number of nodal MEPs assigned to this VLAN.
NodalMipLevelList	Specifies a MIP level list.
NumOfNodalMipLevel	Indicates the number of nodal MIP levels assigned to this VLAN that allows MIP functionality to be enabled for each level for each VLAN.

Configuring Layer 2 ping

Use this procedure to configure a Layer 2 ping for C-VLANs or B-VLANs. This feature enables CFM to debug Layer 2. Layer 2 ping can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

Note:

To use Layer 2 ping for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with **L2Ping**, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2Ping** option to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN or you can autogenerate CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerate CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. From the **L2Ping** tab, configure the Layer 2 ping properties.
4. To initiate a Layer 2 ping, highlight an entry and click the **Start** button.
5. To update a Layer 2 ping, click the **Refresh** button.
6. To stop the Layer 2 ping, click the **Stop** button.

L2Ping field descriptions

Use the data in the following table to use the **L2Ping** tab.

Name	Description
VlanId	Identifies the B-VLAN or the C-VLAN.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Ping transmission.
Messages	Specifies the number of L2Ping messages to be transmitted. The default is 1.
Status	Specifies the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the L2Ping messages.

Name	Description
	• abort: The service aborted or is about to abort the L2Ping messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: The L2Ping Messages will be (or have been) sent. • false: The L2Ping Messages will not be sent. The default is true.
Priority	Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
TimeoutInt	Specifies the interval to wait for an L2Ping time-out. The default value is 3 seconds.
TestPattern	Specifies the test pattern to use in the L2Ping PDU: • allZero: Null signal without cyclic redundancy check. • allZeroCrc: Null signal with cyclic redundancy check with 32-bit polynomial. • pseudoRandomBitSequence: Pseudo-random-bit-sequence without cyclic redundancy check. • pseudoRandomBitSequenceCrc: Pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default value is allZero.
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The default is 0.
FrameSize	Specifies the frame size. If the frame size is specified then the data size is internally calculated and the calculated data size is included in the data TLV. The default is 0.
SourceMode	Specifies the source modes of the transmit loopback service: • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC

Name	Description
	address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.
SeqNumber	The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
Result	Displays the Layer 2 Ping result.

Initiating a Layer 2 traceroute

Use this procedure for B-VLANs or C-VLANs to trigger a Layer 2 traceroute. This feature enables CFM to debug Layer 2. Layer 2 traceroute can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

Note:

To use Layer 2 traceroute for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with **L2Traceroute**, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2Traceroute** option to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN, or you can autogenerate CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerate CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.

About this task

If you configure **IsTraceTree** to false, then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true, then EDM performs TraceTree on the multicast tree.

For more information on how to configure tracetree, see Configuring Layer 2 tracetree on page 377.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an **L2Ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.
- For C-VLANs, you have to trigger an **L2Ping** to learn the C-MAC address.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 Traceroute/TraceTree** tab.
4. To start the traceroute, highlight an entry, and then click the **Start** button.
5. To update the traceroute, click the **Refresh** button.
6. To stop the traceroute, click the **Stop** button.

L2Traceroute field descriptions

Use the data in the following table to use the **L2Traceroute** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the B-VLAN or C-VLAN.
Priority	Specifies a 3–bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Specifies whether the host name is (true) or is not (false) used for the L2Trace transmission.
Isid	Specifies the Service Instance Identifier (I-SID).
IsTraceTree	Specifies whether the multicast tree or unicast path is traced. If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.
Status	Indicates the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the L2Trace messages. • abort: The service aborted or is about to abort the L2Trace messages.

Name	Description
	This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: The L2Trace messages will be (or have been) sent. • false: The L2Trace messages will not be sent. The default is true.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The value of the time-to-live (TTL) field in the L2Trace is defined by the originating MEP. The default value is 64.
SourceMode	Specifies the source mode: • nodal • noVlanMac — Use this value with C-VLANs only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.
SeqNumber	Specifies the transaction identifier/sequence number of the first linktrace message sent. The default is 0.
Flag	L2Trace result flag indicating L2Trace status or error code: • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled to perform L2Trace • noL2TraceResponse (5): No L2Trace response received

Name	Description
	• l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace • l2TraceLeafNode (9): On a leaf node in the I-SID tree • l2TraceNotInTree (10): Not in the I-SID tree • l2TraceSmltNotPrimary (11): Requested SMLT source from non-primary node

Viewing Layer 2 traceroute results

Use this procedure to view Layer 2 traceroute results. This feature enables CFM to debug Layer 2. It can also help you debug ARP problems by providing the ability to troubleshoot next hop ARP records.

About this task

You can display Layer 2 tracetree results to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID. For more information, see Viewing Layer 2 tracetree results on page 379.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2Traceroute/TraceTree** tab.
4. Click the **Refresh** button to update the results.
5. To view the traceroute results, highlight an entry, and then click **Result**.

L2 Traceroute Result field descriptions

Use the data in the following table to use the **L2 Traceroute Result** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the B-VLAN or C-VLAN.
SeqNumber	Specifies the transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which response of the L2Trace is going to be returned. The default is 0.

Name	Description
Hop	Specifies the number of hops away from L2Trace initiator.
ReceiveOrder	Specifies an index to distinguish among multiple L2Trace responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Specifies a time-to-Live (TTL) field value for a returned L2Trace response.
SrcMac	Specifies the MAC address of the MP that responds to the L2Trace request for this L2TraceReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	Specifies the host name of the node that forwarded the L2Trace to the responding node.

Configuring Layer 2 IP ping

Use this procedure for B-VLANs or C-VLANs to configure Layer 2 IP ping.

Layer 2 IP ping allows a user to specify an IP address as the destination address. In this case, the IP address can for a B-VLAN or C-VLAN.

Note:

To use Layer 2 IP Ping for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify an IP address with **L2 IP Ping**, the IP address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2 IP Ping** option to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.

- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN or you can autogenerate CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerate CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.
- If you want to run a Layer 2 IP ping for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Ping** tab.
4. To add a new entry, click **Insert**, specify the destination IP address and optional parameters, and click **Insert**.
5. To start the Layer 2 IP ping, highlight an entry, and then click **Start**.
6. To update the Layer 2 IP ping, click the **Refresh** button.
7. To stop the Layer 2 IP ping, click **Stop**.

L2 IP Ping field descriptions

Use the data in the following table to use the **L2 IP Ping** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address Only IPv4 is supported.
IpAddr	Specifies the destination IP address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Messages	Specifies the number of L2IpPing messages to be transmitted for each MAC/VLAN pair. Range is 1–200. The default is 1.
Status	Specifies the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the L2IpPing messages. • abort: The service is aborted or about to abort the L2IpPing messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.

Name	Description
	The default is ready.
ResultOk	Indicates the result of the operation: • true: L2IpPing Messages will be or have been sent. • false: L2IpPing Messages will not be sent. The default is true.
TimeoutInt	Specifies the interval to wait for an L2IpPing time-out with a range of 1–10 seconds with a default value of 3 seconds.
TestPattern	Specifies the test pattern to use in the L2IPPing PDU: • **allZero** — Null signal without cyclic redundancy check. • **allZeroCrc** — Null signal with cyclic redundancy check with 32-bit polynomial. • **pseudoRandomBitSequence** — Pseudo-random-bit-sequence without cyclic redundancy check. • **pseudoRandomBitSequenceCrc** — Pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. A cyclic redundancy check is a code that detects errors. The default value is allZero.
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The range is 0–400. The default is 0.
PathsFound	Specifies the number of paths found to execute the command. The default is 0.

Viewing Layer 2 IP ping results

Use this procedure to view Layer 2 IP ping results.

Note:

After you trigger Layer 2 IP ping, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Ping** tab.

4. To view the Layer 2 IP ping results, highlight an entry, and then click **Result**.

L2 IP Ping Result field descriptions

Use the data in the following table to use the **L2 IP Ping Result** tab.

Name	Description
IpAddrType	Specifies the address type of the destination IP address.
IpAddr	Specifies the destination IP address.
SendOrder	Specifies the order that sessions were sent. It is an index to distinguish among multiple L2Ping sessions. This value is assigned sequentially from 1. It correlates to the number of paths found.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN ID found from the Layer 3 lookup and used for transmission.
DestMacAddress	Indicates the target MAC address transmitted.
PortNum	Specifies either the value '0', or the port number of the port used for the l2 IP ping.
DestHostName	Specifies the host name of the responding node.
Size	Specifies the number of bytes of data sent.
PktsTx	Specifies the number of packets transmitted for this VLAN/MAC.
PktsRx	Specifies the number of packets received for this VLAN/MAC.
PercentLossWhole	Specifies the percentage of packet loss for this VLAN/MAC.
PercentLossFract	Specifies the percentage of packet loss for this VLAN/MAC.
MinRoundTrip	Specifies the minimum time for round-trip for this VLAN/MAC.
MaxRoundTrip	Specifies the maximum time for round-trip for this VLAN/MAC.
RttAvgWhole	Specifies the average time for round-trip for this VLAN/MAC.
RttAvgFract	Specifies the fractional portion of average time for round-trip.
Flag	Specifies the result flag indicating status or error code: • 1 - No error • 2 - Internal error • 3 - Invalid IP

Name	Description
	• 4 - L2Trace completed • 5 - Lookup failure for IP (no VLAN/MAC entries)

Configuring Layer 2 IP traceroute

Use this procedure for C-VLANs or B-VLANs to configure Layer 2 IP traceroute.

Note:

To use Layer 2 IP traceroute for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify an IP address with **L2 IP Traceroute**, the IP address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2 IP Traceroute** option to test reachability for all the B-MAC addresses in the SPBM network.

Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN or you can autogenerate CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerate CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.
- If you want to run a Layer 2 IP traceroute for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

About this task

If you configure **IsTraceTree** to false, then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true, then EDM performs TraceTree on the multicast tree.

For more information on how to configure tracetree, see Configuring Layer 2 tracetree on page 377.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an **L2Ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.
- For C-VLANs, you have to trigger an **L2Ping** to learn the C-MAC address.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**
3. Click the **L2 IP Traceroute** tab.
4. To add a new entry, click **Insert**, specify the destination IP address and, optionally, the TTL value, and then click **Insert**.
5. To start the Layer 2 IP traceroute, highlight an entry, and then click the **Start** button.
6. To update the L2 IP traceroute, click the **Refresh** button.
7. To stop the Layer 2 IP traceroute, click the **Stop** button.

L2 IP Traceroute field descriptions

Use the data in the following table to use the **L2 IP Traceroute** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address. Only IPv4 is supported.
IPAddr	Specifies the destination IP address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The default value is 64.
Status	Indicates the status of the transmit loopback service: • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit, the L2Trace messages. • abort: Specifies the service is aborted or about to abort the L2Trace messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more

Name	Description
	management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: the Trace Messages will be or have been sent. • false. the Trace Messages will not be sent The default is true.
PathsFound	Specifies the number of paths found to execute the L2Trace. The default is 0.

Viewing Layer 2 IP traceroute results

Use this procedure to view Layer 2 IP traceroute results.

Note:

After you trigger Layer 2 IP traceroute, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Traceroute** tab.
4. To view the Layer 2 IP traceroute results, highlight an entry, and then click **Result**.

L2 IP Traceroute Result field descriptions

Use the data in the following table to use the **L2 IP Traceoute Result** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address.
IpAddr	Specifies the destination IP address.
SendOrder	Denotes the order that sessions are sent. It is an index to distinguish among multiple L2Trace sessions. It correlates to the number of paths found. This value is assigned sequentially from 1.
Hop	Specifies the number of L2 hops away from L2Trace initiator.
ReceiveOrder	Specifies the order that sessions are sent. It is an index to distinguish among multiple L2Trace responses with the same Send Transaction Identifier field value. This value is assigned sequentially from

Name	Description
	1, in the order that the Linktrace Initiator received the responses.
Ttl	Specifies the time-to-live (TTL) field value for a returned L2Trace response.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN found from Layer 3 lookup and used for transmission.
DestMacAddress	Indicates the target MAC address transmitted.
PortNum	Specifies either the value '0', or the port number of the port used for the l2trace.
SeqNumber	Specifies the transaction identifier/sequence number used in linktrace message packet. The default is 0.
SrcMac	Specifies the MAC address of the MP that responded to L2Trace request for this L2traceReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	Specifies the host name of the node that forwarded the L2Trace to the responding node.
Flag	Indicates the L2Trace result flag status or error code: • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace

Triggering a loopback test

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN or C-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **CFM**.
3. Click the **LBM** tab.
4. Configure the loopback test properties as required.
5. Click **Apply**.
6. To trigger the loopback test, double-click in the **Status** field, select **transmit**.
7. Click **Apply**.
8. To update the loopback test, click the **Refresh** button.

LBM field descriptions

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

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the Maintenance Endpoint index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
DestMacAddress	Specifies the remote MAC address to reach the MEP/MIP.
Messages	Specifies the number of loopback messages to be transmitted. The default is 1.
VlanPriority	Specifies the priority. The default is 7.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	Indicates the result of the operation: • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. The default is true.

Name	Description
Status	Indicates the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the Loopback messages. • abort: The service is aborted or about to abort the Loopback messages. The default is ready.
Result	Displays the LBM result.
TimeoutInt	Specifies the timeout interval in seconds. The default value is 3 seconds.
InterFrameInt	Specifies the interval between LBM frames with a range of (0..1000) msecs and a default value of 500 msecs. The value of 0 msecs indicates to send the frames as fast as possible. The default is 500.
TestPattern	Specifies the testfill pattern: • **allZero** — null signal without cyclic redundancy check • **allZeroCrc** — null signal with cyclic redundancy check with 32-bit polynomial • **pseudoRandomBitSequence** — pseudo-random-bit-sequence without cyclic redundancy check • **pseudoRandomBitSequenceCrc** — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial A cyclic redundancy check is a code that detects errors. The default value is allZero.
DataSize	Specifies the data type-length-value (TLV) size. The default is 0.
FrameSize	Specifies the frame-size. The default is 0.
Sourcemode	Specifies the source mode: • **nodal** • **noVlanMac** — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • **smltVirtual** — Use the smltVirtual option with B-VLANs only. The default is nodal.

Triggering linktrace

Use the following procedure to trigger a linktrace. The link trace message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN or C-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **CFM**.
3. Click the **LTM** tab.
4. Configure the linktrace test properties as required.
5. Click **Apply**.
6. To trigger the linktrace test, double-click in the Status field, select **transmit**, and then click **Apply**.

   OR

   Highlight an entry, and then click **Start**.
7. To update the linktrace, click the **Refresh** button.
8. To stop the linktrace, click **Stop**.
9. To view the results of the linktrace, click **Result**.

LTM field descriptions

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

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the MEP index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.

Name	Description
VlanPriority	Specifies the VLAN priority, a 3–bit value to be used in the VLAN tag, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the remote MAC address to reach the MEP.
Ttl	Indicates the number of hops remaining to this LTM. This value is decremented by 1 by each bridge that handles the LTM. The decremented value is returned in the LTR. If the value is 0 on output, the LTM is not transmitted to the next hop. The value of the TTL field in the LTM is specified at the originating MEP. The default value is 64.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	Indicates the result of the operation: • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. The default is true.
Status	Indicates the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the LTM messages. • abort: The service is aborted, or about to abort, the LTM message. The default is ready.
Flag	Displays the LTM result flag indicating LTM status or error code. Each value represents a status or error case: • 1 - No error • 2 - LTM internal error • 3 - Unknown Remote Maintenance Endpoint • 4 - Invalid Remote Maintenance Endpoint MAC Address • 5 - Unset Remote Maintenance Endpoint MAC address • 6 - MEP must be enabled to perform LTM • 7 - No LTR response received • 8 - Linktrace to own MEP MAC is not sent

Name	Description
	• 9 - Endpoint must be enabled in order to perform LTM • 10 - Pbt-trunk must be enabled to perform LTM • 11 - LTM completed • 12 - LTM leaf node
SourceMode	Specifies the source mode: • **nodal** • **noVlanMac** — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • **smltVirtual** — Use the smltVirtual option with B-VLANs only. The default is nodal.

Viewing linktrace results

Use this procedure to view linktrace results.

Note:

After you trigger linktrace, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **CFM**.
3. Click the **LTM** tab.
4. Highlight an entry, and then click **Result**.

Link Trace Replies field descriptions

Use the data in the following table to use the **Link Trace Result** tab.

Name	Description
DomainIndex	Indicates the Maintenance Domain Index.
AssociationIndex	Indicates the Maintenance Association Index.
MepId	Indicates the Maintenance EndPoint ID.
SeqNumber	Indicates the transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which LTM response is going to be returned. The default is 0.

Name	Description
Hop	Indicates the number of hops away from the LTM initiator.
ReceiveOrder	Indicates the index value used to distinguish among multiple LTRs with the same LTR Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the LTRs.
Ttl	Indicates the TTL field value for a returned LTR.
DomainName	Indicates the Maintenance Domain Name.
AssociationName	Indicates the Maintenance Association Name.
Forwarded	Indicates if a LTM was forwarded by the responding MP, as returned in the FwdYes flag of the flags field.
TerminalMep	Displays a boolean value stating whether the forwarded LTM reached a MEP enclosing its MA, as returned in the Terminal MEP flag of the Flags field.
LastEgressIdentifier	Displays an octet field holding the Last Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Last Egress Identifier identifies the MEP Linktrace Indicator that originated, or the Linktrace Responder that forwarded, the LTM to which this LTR is the response. This is the same value as the Egress Identifier TLV of that LTM.
NextEgressIdentifier	Displays an octet field holding the Next Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Next Egress Identifier Identifies the Linktrace Responder that transmitted this LTR, and can forward the LTM to the next hop. This is the same value as the Egress Identifier TLV of the forwarded LTM, if any. If the FwdYes bit of the Flags field is false, the contents of this field are undefined, and the field is ignored by the receiver.
RelayAction	Indicates the value returned in the RelayAction field.
SrcMac	Displays the MAC address of the MP that responded to the LTM request for this LTR.
IngressAction	Displays the value returned in the IngressAction Field of the LTM. The value ingNoTlv indicates that no Reply Ingress TLV was returned in the LTM.
IngressMac	Displays the MAC address returned in the ingress MAC address field. If the rcCfmLtrReplyIngress object contains the value ingNoTlv(5), then the contents of this field are meaningless.
EgressAction	Displays the value returned in the Egress Action Field of the LTM. The value egrNoTlv(5) indicates that no Reply Egress TLV was returned in the LTM.

Name	Description
EgressMac	Displays the MAC address returned in the egress MAC address field. If the rcCfmLtrReplyEgress object contains the value egrNoTlv(5), then the contents of this field are meaningless.

Configuring Layer 2 tracetree

Use this procedure to configure a Layer 2 Tracetree. This feature enables CFM to debug Layer 2. Layer 2 Tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.

Note:

VSP 9000 only supports this command on SPBM B-VLANs only, not C-VLANs.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. From the **L2 Traceroute/TraceTree** tab, configure the Layer 2 tracetree properties.
4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetree on the multicast tree.
5. Click **Apply**.
6. Click the **Refresh** button to update the results.

L2Tracetree field descriptions

Use the data in the following table to use the **L2Tracetree** tab.

Name	Description
VlanId	Identifies the Backbone VLAN.
Priority	Specifies a 3–bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the target MAC address.

Name	Description
HostName	Specifies the target host name.
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Tracetree transmission.
Isid	Specifies the service instance identifier (I-SID).
IsTraceTree	Specifies whether the multicast tree or unicast path is traced. If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.
Status	Specifies the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Tracetree messages. • abort: the service aborted or is about to abort the L2Tracetree messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time. The default is ready.
ResultOk	Indicates the result of the operation: • true: the L2Tracetree Messages will be (or have been) sent. • false: the L2Tracetree Messages will not be sent The default is true.
Ttl	Specifies the Time-to-Live value. Indicates the number of hops remaining to this L2Tracetree. The tracetree is decremented by one by each bridge that handles the Layer 2 tracetree and the decremented value is returned to the tracetree. If the output is 0, then the L2Tracetree is not transmitted to the next hop. The value of the TTL field in the L2Tracetree is transmitted by the originating MEP is controlled by a managed object. The default is 64.
SourceMode	Specifies the source modes of the transmit loopback service: • nodal • noVlanMac — Use the noVlanMac option with C-VLANs only. • smltVirtual — Use the smltVirtual option with B-VLANs only.

Name	Description
	The default is nodal.
SeqNumber	The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
Flag	Specifies the L2Tracetree result flag, which indicates the L2Tracetree status or error code. Each sum represents a status or error: • 1 — No error • 2 — L2Tracetree internal error • 3 — Invalid MAC address • 4 — MEP must be enabled to perform L2Tracetree • 5 — No L2Tracetree response received • 6 — L2Tracetree to own MEP MAC is not sent • 7 — L2Tracetree completed • 8 — Lookup failure for L2Tracetree • 9 — On a leaf node in the I-SID tree • 10 — Not in the I-SID tree • 11 — Requested SMLT source from non-primary node

Viewing Layer 2 tracetree results

Use this procedure to view Layer 2 Tracetree results. The Layer 2 Tracetree command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 Traceroute/TraceTree** tab.
4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetree on the multicast tree.
5. Click **Apply**.
6. Click the **Refresh** button to update the results.
7. To view the tracetree results, highlight an entry, and then click **Result**.

L2 Tracetree Result field descriptions

Use the data in the following table to use the **L2 Tracetree Result** tab.

Name	Description
VlanId	A value that uniquely identifies the Backbone VLAN (B-VLAN).
SeqNumber	The transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which response of the L2Tracetree is going to be returned. The default is 0.
Hop	The number of hops away from L2Tracetree initiator.
ReceiveOrder	An index to distinguish among multiple L2Tracetree responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Time-to-Live (TTL) field value for a returned L2Tracetree response.
SrcMac	MAC address of the MP that responds to the L2Tracetree request for this L2tractreeReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Tracetree to the responding node.
LastHostName	The host name of the node that forwarded the L2Tracetree to the responding node.

Configuring Layer 2 trace multicast route on a VLAN

Use this procedure to configure the Layer 2 tracemroute on the VLAN (Layer 2). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID, and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

Note:

If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context. See the following procedure to perform a Layer 3 tracemroute on a VRF, Configuring Layer 2 tracemroute on a VRF on page 382.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics** > **L2Ping/L2Trace Route**.
2. Click the **L2MCAST Traceroute** tab.
3. Click **Insert** to insert the L2 MCAST Traceroute.
4. Type the **SrcIpAddr**.
5. Type the **GroupIpAddr**.
6. Enter the **ServiceType**. If you want to perform a Layer 2 tracemroute on a VLAN, select **vlan**. If you want to perform a Layer 2 tracemroute on a Layer 3 GRT, select **vrfid**.

   Note:

   If you want to perform a Layer 2 tracemroute on a Layer 2 or a Layer 3 VRF, review the following procedure Configuring Layer 2 tracemroute on a VRF on page 382.
7. In the **ServiceId**, enter the VLAN ID.
8. Enter the **Priority**.
9. Enter the **Ttl** value.
10. Click **Insert**.
11. Click **Apply** to save your changes.
12. To start the Layer 2 tracemoute, set the Status to transmit and click the **Start** button.
13. Update the Layer 2 tracemroute by clicking the **Refresh** button.
14. To stop the Layer 2 tracemroute, click the **Stop** button.
15. To see the result, click the **Result** button.

L2 MCAST Traceroute field descriptions

Use the data in the following table to use the **L2MCAST Traceroute** tab.

Name	Description
SrcIpAddrType	Specifies the source IP address type as IPv4.
SrcIpAddr	Specifies the source IP address of the flow where the multicast trace tree originates.
GroupIpAddrType	Specifies the group IP address type as IPv4.
GroupIpAddr	Specifies the group IP address.
ServiceType	Specifies where you configure the Layer 2 tracemroute. This is either VLAN or VRF.
VRFName	Specifies the VRF name.
Priority	Specifies the priority value. The value is between 0 and 7.
Ttl	Specifies the returned trace response. The TTL value is between 1 and 255.

Name	Description
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	Specifies the status of the transmit loopback service: • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit the trace messages. • abort: Specifies the services is aborted or about to abort the trace messages. The column will also be used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.
ResultOK	Specifies the result of the operation: • true: The trace messages will be or have been sent. • false: The trace messages will not be sent.
Flag	Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case. • 1 — No error • 2 — Internal Error • 3 — Mep must be enabled to perform the trace • 4 — No response received • 5 — Trace completed • 6 — On a leaf node in the I-SID tree • 7 — No data I-SID was found for S, G

Configuring Layer 2 tracemroute on a VRF

Use this procedure to configure the Layer 2 tracemroute on the VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

Note:

If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context.

See the following procedure to perform a Layer 3 tracemroute on a VLAN Configuring Layer 2 tracemroute on a VLAN on page 380.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **VRF Context View** > **Set VRF Context View**
2. Select a VRF and click the **Launch VRF Context View** tab.
3. From the navigation tree, expand the following folders:**Configuration** > **Edit** > **Diagnostics** > **L2Ping/L2Trace Route**.
4. Click the **L2MCAST Traceroute** tab.
5. Click **Insert** to insert the L2 MCAST traceroute.
6. Type the **SrcIpAddr**.
7. Type the **GroupIpAddr**.
8. Enter the **ServiceType**. If you want to perform a Layer 2 tracemroute on a Layer 2 VRF, select **vlan**. If you want to perform a Layer 2 tracemroute on a Layer 3 VRF, select **vrfid**.
9. In the **ServiceId**, enter the VLAN ID.
10. Enter the **Priority**.
11. Enter the **Ttl** value.
12. Click **Insert**.
13. Click **Apply** to save your changes.
14. To start the Layer 2 tracemoute, set the Status to transmit and click the **Start** button.
15. Update the Layer 2 tracemroute by clicking the **Refresh** button.
16. To stop the Layer 2 tracemroute, click the **Stop** button.
17. To see the result, click the **Result** button.

L2 MCAST Traceroute field descriptions

Use the data in the following table to use the **L2MCAST Traceroute** tab.

Name	Description
SrcIpAddrType	Specifies the source IP address type as IPv4.
SrcIpAddr	Specifies the source IP address of the flow where the multicast trace tree originates.
GroupIpAddrType	Specifies the group IP address type as IPv4.
GroupIpAddr	Specifies the group IP address.
ServiceType	Specifies where you configure the Layer 2 tracemroute. This is either VLAN or VRF.
VRFName	Specifies the VRF name.

Name	Description
Priority	Specifies the priority value. The value is between 0 and 7.
Ttl	Specifies the returned trace response. The TTL value is between 1 and 255.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	Specifies the status of the transmit loopback service: • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit the trace messages. • abort: Specifies the services is aborted or about to abort the trace messages. The column will also be used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.
ResultOK	Specifies the result of the operation: • true: The trace messages will be or have been sent. • false: The trace messages will not be sent.
Flag	Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case. • 1 — No error • 2 — Internal Error • 3 — Mep must be enabled to perform the trace • 4 — No response received • 5 — Trace completed • 6 — On a leaf node in the I-SID tree • 7 — No data I-SID was found for S, G

Viewing Layer 2 trace multicast route results

Use this procedure to view Layer 2 tracemroute results.

Procedure

1. From the navigation tree, expand the following folders: **Configuration** > **Edit** > **Diagnostics** > **L2Ping/L2Trace Route**
2. Click the **L2 MCAST Traceroute** tab.
3. To view the CFMl2 trace multicast route results, highlight an entry and click the **Result** button.

L2tracemroute Result field descriptions

Use the data in the following table to use the **L2tracemroute Result** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the C-VLAN.
SeqNumber	Specifies the transaction identifier/sequence number returned by a previous transmit linktrace message command. Indicates which l2 tracemroute response is going to be returned.
Hop	Specifies the number of hops away from the l2 tracemroute initiator.
ReceiveOrder	Specifies an index to distinguish among multiple l2 tracemroute responses with the same transaction identifier field value. This value is assigned sequentially from 1, in the order that the linktrace initiator received the responses.
Ttl	Specifies the TTL value for a returned l2 tracemroute response.
SrcMac	Specifies the MAC address of the MP that responds to the l2 tracemroute request for this l2 tracemrouteReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the l2 tracemroute to the responding node.
LastHostName	Specifies the host name of the node that forwarded the l2 tracemroute to the responding node.

CFM configuration example

This section provides a configuration example for Connectivity Fault Management (CFM).

CFM configuration example

The following sections show the steps required to explicitly configure CFM.

VSP9000 A

```
MAINTENANCE-DOMAIN CONFIGURATION

cfm maintenance-domain "spbm" index 1 maintenance-level 6

MAINTENANCE-ASSOCIATION CONFIGURATION

cfm maintenance-association "spbm" "2" index 1
cfm maintenance-association "spbm" "3" index 2

MAINTENANCE-ENDPOINT CONFIGUARTION

cfm maintenance-endpoint "spbm" "2" 1 state enable
```

```
cfm maintenance-endpoint "spbm" "3" 1 state enable

VLAN NODAL MEP/MIP CONFIGURATION

vlan nodal-mep 2 spbm 2 1
vlan nodal-mip-level 2 6
vlan nodal-mep 3 spbm 3 1
vlan nodal-mip-level 3 6
```

VSP9000 B

```
MAINTENANCE-DOMAIN CONFIGURATION

cfm maintenance-domain spbm index 1 maintenance-level 6

MAINTENANCE-ASSOCIATION CONFIGURATION

cfm maintenance-association "spbm" "2" index 1
cfm maintenance-association "spbm" "3" index 2

MAINTENANCE-ENDPOINT CONFIGUARTION

cfm maintenance-endpoint "spbm" "2" 2 state enable
cfm maintenance-endpoint "spbm" "3" 2 state enable

VLAN NODAL MEP/MIP CONFIGURATION

vlan nodal-mep 2 spbm 2 2
vlan nodal-mip-level 2 6
vlan nodal-mep 3 spbm 3 2
vlan nodal-mip-level 3 6
```

CFM sample output

The following sections show sample CFM output.

L2ping can use the system ID or the router name. The example below shows a case where the VLAN and MAC are given.

show isis adjacencies

```
VSP-9012:1# show isis adjacencies
 ================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE     IP ADDR            L STATE        UPTIME         PRI     HOLDTIME
SYSID
--------------------------------------------------------------------------------
Port3/3      44.17.10.33        1 UP         00:37:37         127   19
0014.0dbf.a3df
Port3/19     44.17.10.36        1 UP      1d 05:09:16         127   21
0014.0da2.b3df
--------------------------------------------------------------------------------
2 out of 2 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

l2 ping

```
VSP-9012:1# l2 ping vlan 500 mac 00.14.0d.bf.a3.df


Please wait for l2ping to complete or press any key to abort
```

```
----00:14:0d:bf:a3:df    L2 PING Statistics----  0(64) bytes of data
1 packets transmitted, 0 packets received,   100.00% packet loss
```

l2 ping

```
VSP-9012:1# l2 ping vlan 500 routernodename VSP-MONTI0


Please wait for l2ping to complete or press any key to abort

----00:14:0d:a2:b3:df    L2 PING Statistics----  0(64) bytes of data
1 packets transmitted, 1 packets received,  0.00% packet loss
  round-trip (us)        min/max/ave/stdv =  26895/26895/26895.00/ 0.00
```

l2 traceroute

```
VSP-9012:1# l2 traceroute vlan 500 routernodename VSP-MONTI0


Please wait for l2traceroute to complete or press any key to abort

l2traceroute to VSP-MONTI0  (00:14:0d:a2:b3:df),  vlan 500
0    VSP-PETER4         (00:15:9b:11:33:df)
1    VSP-MONTI0          (00:14:0d:a2:b3:df)
```

l2 tracetree

```
VSP-9012:1# l2 tracetree 500 1

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10 hops 64
1   VSP-PETER4        00:15:9b:11:33:df -> VSP-MONTI0        00:14:0d:a2:b3:df
2   VSP-MONTI0        00:14:0d:a2:b3:df -> VSP-LEE2          00:15:e8:b8:a3:df
```

l2 tracemroute

```
VSP-9012:1(config)#l2 tracemroute source 192.0.2.81 233.252.0.1 vlan 201

Please wait for l2 tracemroute to complete or press any key to abort.

Source 192.0.2.81

Group: 233.252.0.1

VLAN:201

BMAC: 03:00:03:f4:24:01

B-VLAN: 10

I-SID: 16000001

================================================================
1 VSP-PETER4 00:03:00:00:00:00 -> VSP-LEE1 00:14:0d:bf:a3:df
2 VSP-LEE1 00:14:0d:bf:a3:df -> VSP-LEE2 00:15:e8:b8:a3:df
```

L2ping and L2traceroute can also be used with an IP address. The following outputs show examples using an IP address.

l2 ping

```
VSP-9012:1# l2 ping ip-address 10.1.1.1
```

```
Please wait for l2ping to complete or press any key to abort

L2 PING  Statistics : IP 10.1.1.1, paths found 1, paths attempted 1
================================================================================
TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                              PKTS  PKTS  LOSS     MIN/MAX/AVE (us)
================================================================================
500  VSP-SHAMIM        (00:1a:8f:08:53:df)  1     0     100.00%  0/0/0.00
```

l2 traceroute

```
VSP-9012:1# l2 traceroute ip-address 10.1.1.1


Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 10.1.1.1, paths found 1
================================================================
VSP-SHAMIM  (00:1a:8f:08:53:df),  vlan 500
0    VSP-PETER4          (00:15:9b:11:33:df)
1    VSP-MONTI0          (00:14:0d:a2:b3:df)
```

show cfm maintenance-domain

```
VSP-9012:1#show cfm maintenance-domain
================================================================================
                          Maintenance Domain
================================================================================
Domain Name             Domain Index    Level Domain Type
--------------------------------------------------------------------------------
md1                     99              3     NONE

 Total number of Maintenance Domain entries: 1.
```

show cfm maintenance-association

```
VSP-9012:1#show cfm maintenance-association
================================================================================
                          Maintenance Association Status

================================================================================
Domain Name             Assn Name              Domain Idx  Assn Idx
--------------------------------------------------------------------------------
md1                     ma1                    1           98

 Total number of Maintenance Association entries: 1.

================================================================================
                          Maintenance Association config
================================================================================
Domain Name             Assn Name
--------------------------------------------------------------------------------
md1                     ma1

 Total number of MA entries: 1.
```

show cfm maintenance-endpoint

```
VSP-9012:1#show cfm maintenance-endpoint
================================================================================
                          Maintenance Endpoint Config
================================================================================
DOMAIN                  ASSOCIATION            MEP  ADMIN
NAME                    NAME                   ID
--------------------------------------------------------------------------------
```

```
md1                     ma1                     1    enable

 Total number of MEP entries: 1.

================================================================================
                          Maintenance Endpoint Service
================================================================================
DOMAIN_NAME             ASSN_NAME               MEP_ID TYPE    SERVICE_DESCRIPTION
--------------------------------------------------------------------------------
md1                     ma1                     1      unused

 Total number of MEP entries: 1.
```

show vlan nodal-mep

```
VSP-9012:1#show vlan nodal-mep

================================================================================
                                  Vlan Nodal Mep
================================================================================
VLAN_ID    DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
--------------------------------------------------------------------------------
1
2
3
4          md1.ma1.1
5
6
7
8
9
10
11
12
13
14
```

show vlan nodal-mip-level

```
VSP-9012:1#show vlan nodal-mip-level

================================================================================
                              Vlan Nodal Mip Level
================================================================================
VLAN_ID    NODAL_MIP_LEVEL_LIST
--------------------------------------------------------------------------------
1
2
3
4          6
5
6
7
8
9
10
11
12
13
14
```

show cfm spbm

```
VSP-9012:1(config)#show cfm spbm

LEVEL ADMIN    MEPID    MAC
```

```
=======================================
6         enable         4         00:15:e8:b8:a3:df
```

show cfm cmac

```
VSP-9012:1(config)#show cfm cmac

LEVEL ADMIN    MEPID    MAC
=======================================
0         enable         4         00:15:e8:b8:a3:de
```

Glossary

Autonomous System (AS)	A set of routers under a single technical administration, using a single IGP and common metrics to route packets within the AS, and using an EGP to route packets to other ASs.
autonomous system border router (ASBR)	A router attached at the edge of an OSPF network. An ASBR uses one or more interfaces that run an interdomain routing protocol such as BGP. In addition, a router distributing static routes or Routing Information Protocol (RIP) routes into OSPF is considered an ASBR.
Autonomous System Number (ASN)	A two-byte number that is used to identify a specific AS.
Avaya command line interface (ACLI)	A textual user interface. When you use ACLI, you respond to a prompt by typing a command. After you enter the command, you receive a system response.
Backbone Core Bridge (BCB)	Backbone Core Bridges (BCBs) form the core of the SPBM network. The BCBs are SPBM nodes that do not terminate the VSN services. BCBs forward encapsulated VSN traffic based on the Backbone MAC Destination Address (B-MAC-DA). A BCB can access information to send that traffic to any Backbone Edge Bridges (BEBs) in the SPBM backbone.
Backbone Edge Bridge (BEB)	Backbone Edge Bridges (BEBs) are SPBM nodes where Virtual Services Networks (VSNs) terminate. BEBs handle the boundary between the core MAC-in-MAC Shortest Bath Bridging MAC (SPBM) domain and the edge customer 802.1Q domain. A BEB node performs 802.1ah MAC-in-MAC encapsulation and decapsulation for the Virtual Services Network (VSN).
Backbone MAC (B-MAC)	Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation encapsulates customer MAC addresses in Backbone MAC (B-MAC) addresses. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that SPBM uses for delivery from end to end. As the MAC header stays the same across the network, no need exists to swap a label or perform a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end. In Shortest Path Bridging MAC (SPBM), each node has a System ID, which is used in the topology announcement. This same System ID also serves as the switch

	Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network.
Backbone VLAN identifier (B-VID)	The Backbone VLAN identifier (B-VID) indicates the Shortest Path Bridging MAC (SPBM) B-VLAN associated with the SPBM instance.
Border Gateway Protocol (BGP)	An inter-domain routing protocol that provides loop-free inter-domain routing between Autonomous Systems (AS) or within an AS.
Circuitless IP (CLIP)	A CLIP is often called a loopback and is a virtual interface that does not map to any physical interface.
Complete Sequence Number Packets (CSNP)	Complete Sequence Number Packets (CSNP) contain the most recent sequence numbers of all Link State Packets (LSPs) in the database. When all routers update their LSP database, synchronization is complete.
Connectivity Fault Management (CFM)	Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) network. CFM operates at Layer 2 and provides the equivalent of ping and traceroute. IEEE 802.1ag Connectivity Fault Management (CFM) divides or separates a network into administrative domains called Maintenance Domains (MD).
Control Processor (CP) module	The Control Processor module runs all high level protocols (BGP, OSPF) and distributes the results (routing updates) to the rest of the system. The CP manages and configures the IO and Switch Fabric modules, and maintains and monitors the health of the chassis.
Customer MAC (C-MAC)	For customer MAC (C-MAC) addresses, which is customer traffic, to forward across the service provider back, SPBM uses IEEE 802.1ah Provider Backbone Bridging MAC-in-MAC encapsulation. The system encapsulates C-MAC addresses within a backbone MAC (B-MAC) address pair made up of a BMAC destination address (BMAC-DA) and a BMAC source address (BMAC-SA).
Customer VLAN (C-VLAN)	A traditional VLAN with MAC learning and flooding, where user devices connect to the network. In SPBM, C-VLANs are mapped to a Service Instance Identifier (I-SID) at the Backbone Edge Bridges (BEBs).
Data I-SID	In SPBM, the data I-SID is allocated by the Backbone Edge Bridge (BEB) when the multicst stream reaches the BEB. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID. Data is transported from the sender to the receiver across the SPBM cloud using the data I-SID.
Designated Intermediate System (DIS)	A Designated Intermediate System (DIS) is the designated router in Intermediate System to Intermediate System (IS-IS) terminology. You can modify the priority to affect the likelihood of a router being elected the designated router. The higher the priority, the more likely the router is to be

	elected as the DIS. If two routers have the same priority, the router with the highest MAC address (Sequence Number Packet [SNP] address) is elected as the DIS.
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.
Enterprise Device Manager (EDM)	A Web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.
equal cost multipath (ECMP)	Distributes routing traffic among multiple equal-cost routes.
Global routing engine (GRE)	The base router or routing instance 0 in the Virtual Routing and Forwarding (VRF).
Global Routing Table (GRT)	The Global Routing Table (GRT) is a table that maintains the information needed to forward an IP packet along the best route.
graphical user interface (GUI)	A graphical (rather than textual) computer interface.
IEEE 802.1aq	IEEE 802.1aq is the standard for Shortest Path Bridging MAC (SPBM). SPBM makes network virtualization much easier to deploy within, reducing the complexity of the network while at the same time providing greater scalability. 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.
IGMP querier	The Internet Group Management Protocol (IGMP) querier is a multicast router that must exist on the network to generate IGMP queries. IGMP queries are sent to all switches by IGMP snooping to determine which multicast receivers want the multicast sender multicast stream. Without a querier the tables for IGMP snooping cannot be created and snooping does not work.

IGMP queries	IGMP queries are sent by a router, acting as an IGMP querier, to learn the existence of host group members. The router sends IGMP queries and hosts respond by issuing IGMP reports.
IGMP report	An IGMP report is a message sent by a host letting an IGMP querier know that the host is joining, maintaining or leaving its membership in the multicast group.
IGMP snooping	Internet Group Management Protocol (IGMP) Snooping is widely used on Layer 2 access switches to prune multicast traffic. IGMP snooping is the process of the network switch listening to a conversation between hosts and routers. By listening, the switch maintains a table for links that need IP multicast streams.
Institute of Electrical and Electronics Engineers (IEEE)	An international professional society that issues standards and is a member of the American National Standards Institute, the International Standards Institute, and the International Standards Organization.
Interior Gateway Protocol (IGP)	Distributes routing information between routers that belong to a single Autonomous System (AS).
Intermediate System to Intermediate System (IS-IS)	Intermediate System to Intermediate System(IS-IS) is a link-state, interior gateway protocol. ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System to Intermediate System (IS-IS). IS-IS operation is similar to Open Shortest Path First (OSPF). In Shortest Path Bridging MAC (SPBM) networks, IS-IS discovers network topology and builds shortest path trees between network nodes that IS-IS uses for forwarding unicast traffic and determining the forwarding table for multicast traffic. SPBM employs IS-IS as the interior gateway protocol and implements additional Type-Length-Values (TLVs) to support additional functionality.
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.
Internet Protocol multicast (IPMC)	The technology foundation for audio and video streaming, push applications, software distribution, multipoint conferencing, and proxy and caching solutions.
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 SPBM	With IP multicast over SPBM, 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.

IS-IS Hello packets

Intermediate System to Intermediate System (IS-IS) uses Hello packets to initialize and maintain adjacencies between neighboring routers. IS-IS Hello packets contain the IP address of the interface over which the Hello transmits. These packets are broadcast to discover the identities of neighboring IS-IS systems and to determine whether the neighbor is a Level 1 router.

Layer 1

Layer 1 is the Physical Layer of the Open System Interconnection (OSI) model. Layer 1 interacts with the MAC sublayer of Layer 2, and performs character encoding, transmission, reception, and character decoding.

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).

Layer 4

The Transport Layer of the OSI model. An example of a Layer 4 protocol is Transmission Control Protocol (TCP).

Link Layer Discovery Protocol (LLDP)

Link Layer Discovery Protocol is used by network devices to advertise their identities. Devices send LLDP information at fixed intervals in the form of Ethernet frames, with each frame having one Link Layer Discovery Protocol Data Unit.

Link State Packets (LSP)

Link State Packets (LSP) contain information about the state of adjacencies or defined and distributed static routes. Intermediate System to Intermediate System (IS-IS) exchanges this information with neighboring

	IS-IS routers at periodic intervals. Every router in the domain has an identical link state database and each runs shortest path first to calculate routes.
Link State Protocol Data Unit (LSPDUs)	Link State Protocol Data Unit is similar to a Link State Advertisement in Open Shortest Path First (OSPF). Intermediate System to Intermediate System (IS-IS) runs on all nodes of Shortest Path Bridging-MAC (SPBM). Since IS-IS is the basis of SPBM, the device must first form the IS-IS adjacency by first sending out hellos and then Link State Protocol Data Units. After the hellos are confirmed both nodes sends Link State Protocol Data Units (LSPDUs) that contain connectivity information for the SPBM node. These nodes also send copies of all other LSPDUs they have in their databases. This establishes a network of connectivity providing the necessary information for each node to find the best and proper path to all destinations in the network.
link trace message	The link trace message (LTM) is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP, which is the SPBM system id or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR. LTM contains: • Time to live (TTL) • Transaction Identifier • Originator MAC address • Target MAC address
link-state database (LSDB)	A database built by each OSPF router to store LSA information. The router uses the LSDB to calculate the shortest path to each destination in the autonomous system (AS), with itself at the root of each path.
Local Area Network (LAN)	A data communications system that lies within a limited spatial area, uses a specific user group and topology, and can connect to a public switched telecommunications network (but is not one).
Loopback Messages (LBM)	A Loopback Message (LBM) is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a Maintenance End Point (MEP) or Maintenance Intermediate Point (MIP), but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR. The system achieves fault verification through the use of Loopback Messages (LBM).
Loopback Response (LBR)	Loopback Response (LBR) is the response from a Maintenance Point (MP).

MAC-in-MAC encapsulation	MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that the device uses for delivery from end to end. As the MAC header stays the same across the network, there is no need to swap a label or do a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end.
Maintenance Associations (MA)	Maintenance Associations (MA) are administrative associations in a network that is divided by the 802.1ag Connectivity Fault Management (CFM) feature. CFM groups MAs within Maintenance Domains. Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.
Maintenance Domains (MD)	Maintenance Domains (MD) are administrative domains that divides a network by the 802.1ag Connectivity Fault Management (CFM) feature. Each MD is further subdivided into logical groupings called Maintenance Associations (MA). Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.
Maintenance Points (MP)	Maintenance Points (MP) are a demarcation point on an interface that participates in Connectivity Fault Management (CFM) within a Maintenance Domain (MD). There are two types of MP: Maintenance End Point (MEP) and Maintenance Intermediate Point (MIP). Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.
MD5 Authentication	MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum. There is an optional key ID.
Media Access Control (MAC)	Arbitrates access to and from a shared medium.
multicast group ID (MGID)	The multicast group ID (MGID) is a hardware mechanism the switch uses to send data to several ports simultaneously. Instead of sending the data to a specific port number, the switch directs the data to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs.
MultiLink Trunking (MLT)	A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.

Network Entity Title (NET)	The Network Entity Title (NET) is the combination of all three global parameters: Manual area, System ID and NSEL. • Manual area — The manual area or area ID is up to 13 bytes long. The first byte of the area number (for example, 49) is the Authority and Format Indicator (AFI). The next bytes are the assigned domain (area) identifier, which is up to 12 bytes (for example, 49.0102.0304.0506.0708.0910.1112). • System ID — The system ID is any 6 bytes that are unique in a given area or level. The system ID defaults to the node BMAC. • NSEL — The last byte (00) is the n-selector. In the Avaya Ethernet Routing Switch 8800/8600 implementation, this part is automatically attached. There is no user input accepted.
Open Shortest Path First (OSPF)	A link-state routing protocol used as an Interior Gateway Protocol (IGP).
Packet Capture Tool (PCAP)	A data packet capture tool that captures ingress and egress (on Ethernet modules only) packets on selected ports. You can analyze captured packets for troubleshooting purposes.
Partial Sequence Number Packets (PSNP)	Partial Sequence Number Packets (PSNP) are requests for missing Link State Packets (LSPs). When a receiving router detects a missing LSP, it sends a PSNP to the router that sent the Complete Sequence Number Packets (CSNP).
port	A physical interface that transmits and receives data.
Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
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.
Provider Backbone Bridge (PBB)	To forward customer traffic across the service-provider backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses.
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 checking (RPC)	Prevents packet forwarding for incoming IP packets with incorrect or forged (spoofed) IP addresses.
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.
Routing Information Protocol (RIP)	A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.
scope level	In IP Multicast over SPBM, the scope level is the level in which the multicast stream is constrained. For instance, if a sender sends a multicast stream to a BEB on a Layer 2 Virtual Services Network (VSN) only receivers that are part of a Layer 2 VSN can receive that stream. Similarly, if a sender sends a multicast stream to a BEB on a Layer 3 VSN only receivers that are part of a Layer 3 VSN can receive that stream.
Service Instance Identifier (I-SID)	The SPBM B-MAC header includes a Service Instance Identifier (I-SID) with a length of 24 bits. SPBM uses this I-SID to identify and transmit any virtualized traffic in an encapsulated SPBM frame. SPBM uses I-SIDs to virtualize VLANs (Layer 2 Virtual Services Network [VSN]) or VRFs (Layer 3 Virtual Services Network [VSN]) across the MAC-in-MAC backbone. With Layer 2 VSNs, you associate the I-SID with a customer VLAN, which is then virtualized across the backbone. With Layer 3 VSNs, you associate the I-SID with a customer VRF, which is also virtualized across the backbone.
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 first (SPF)	A class of routing protocols that use Djikstra's algorithm to compute the shortest path through a network, according to specified metrics, for efficient transmission of packet data.
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.
Split MultiLink Trunking (SMLT)	An Avaya extension to IEEE 802.1AX (link aggregation), provides nodal and link failure protection and flexible bandwidth scaling to improve on the level of Layer 2 resiliency.
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.
Top of Rack (TOR)	A Top of Rack (TOR) switch refers to a switch that sits at the top or near the top of a rack often found in data centers.
Virtual Enterprise Network Architecture (VENA)	Virtual Enterprise Network Architecture (VENA) is a virtualization architecture for next generation Enterprise data networks. VENA provides the data infrastructure for the private cloud by leveraging an open and interoperable IEEE technology, called Shortest Path Bridging (SPB) standard (802.1aq). VENA allows for the aggregation of multiple independent virtual servers to exist on a physical server and decouples the physical infrastructure from the connectivity services making the network adaptive and dynamic with simple one-touch provisioning.
Virtual Link Aggregation Control Protocol (VLACP)	Virtual Link Aggregation Control Protocol (VLACP) is a Layer 2 handshaking protocol that can detect end-to-end failure between two physical Ethernet interfaces.

Virtual Local Area Network (VLAN)	A Virtual Local Area Network is a group of hosts that communicate as if they are attached to the same broadcast domain regardless of their physical location. VLANs are layer 2 constructs.
Virtual Private Network (VPN)	A Virtual Private Network (VPN) requires remote users to be authenticated and ensures private information is not accessible to unauthorized parties. A VPN can allow users to access network resources or to share data.
VLAN Identifier (VID)	VLAN Identifier (VID) is a data field in IEEE 802.1Q VLAN tagging.