



Configuring Fabric Connect on VSP Operating System Software

Release 6.0.1
NN47227-510
Issue 10.02
January 2017

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 in varying mediums which may include product information, operating instructions and performance specifications that are generally made available to users of products. 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 or on the express behalf of 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: <https://support.avaya.com/helpcenter/getGenericDetails?detailId=C20091120112456651010> under the link "Warranty & Product Lifecycle" 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.

"Hosted Service" means an Avaya hosted service subscription that You acquire from either Avaya or an authorized Avaya Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Avaya or Avaya Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN AVAYA HOSTED SERVICE SUBSCRIPTION FROM AVAYA OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE AVAYA WEBSITE, [HTTPS://SUPPORT.AVAYA.COM/LICENSEINFO](https://support.avaya.com/LICENSEINFO) UNDER THE LINK "Avaya Terms of Use for Hosted Services" OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR

IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

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

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

Licence types

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

Heritage Nortel Software

"Heritage Nortel Software" means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software is the software contained within the list of Heritage Nortel Products located at <https://support.avaya.com/LicenseInfo> under the link "Heritage Nortel Products" or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants Customer 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.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Avaya Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

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: <https://support.avaya.com/Copyright> or such successor site as designated by Avaya. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE [HTTP://WWW.MPEGLA.COM](http://www.mpegla.com).

Service Provider

THE FOLLOWING APPLIES TO AVAYA CHANNEL PARTNER’S HOSTING OF AVAYA PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN AVAYA CHANNEL PARTNER’S HOSTING OF AVAYA PRODUCTS MUST BE AUTHORIZED IN WRITING BY AVAYA AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT SOFTWARE OR CODECS, THE AVAYA CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS, AT THE AVAYA CHANNEL PARTNER’S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER.

WITH RESPECT TO CODECS, IF THE AVAYA CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE AVAYA CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE AVAYA CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE

[WWW.SIPRO.COM/CONTACT.HTML](http://www.sipro.com/contact.html). THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE [HTTP://WWW.MPEGLA.COM](http://www.mpegla.com).

Compliance with Laws

You acknowledge and agree that it is Your responsibility for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Avaya product is used.

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: <https://support.avaya.com> or such successor site as designated by Avaya.

Security Vulnerabilities

Information about Avaya’s security support policies can be found in the Security Policies and Support section of <https://support.avaya.com/security>.

Suspected Avaya product security vulnerabilities are handled per the Avaya Product Security Support Flow (<https://support.avaya.com/css/P8/documents/100161515>).

Downloading Documentation

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

Contact Avaya Support

See the Avaya Support website: <https://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: <https://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, its licensors, its suppliers, 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	9
Introduction.....	9
Purpose.....	9
Chapter 2: New in this document	10
Release 6.0.1.....	10
Release 6.0.....	10
Chapter 3: SPBM and IS-IS infrastructure configuration	14
SPBM and IS-IS infrastructure fundamentals.....	14
spbm-config-mode boot flag.....	17
vxlan-gw-full-interworking-mode boot flag.....	18
MAC-in-MAC encapsulation.....	18
I-SID.....	18
BCBs and BEBs.....	19
VLANs without member ports.....	20
Basic SPBM network topology.....	20
E-Tree and Private VLAN topology.....	21
IS-IS.....	23
Standard TLVs.....	24
IS-IS hierarchies.....	26
IS-IS PDUs.....	26
IS-IS configuration parameters.....	27
SPBM B-VLAN.....	30
Pre-populated FIB.....	30
RPFC.....	31
SPBM FIB.....	31
SPBM restrictions and limitations.....	32
IP Multicast over Fabric Connect.....	34
How IP Multicast over Fabric Connect works.....	35
BEB as IGMP Querier.....	37
Network Load Balancing (NLB)	38
Switch clustering at the edge of the SPBM network.....	38
Considerations when you connect an IP Multicast over Fabric Connect network to a PIM network.....	41
IP Multicast over Fabric Connect restrictions.....	42
SPBM script.....	43
Run vms endura script.....	44
Fabric Extend.....	46
Fabric Attach.....	63
IS-IS external metric.....	80

SPB Ethernets	80
SPBM and IS-IS infrastructure configuration using CLI.....	81
Running the SPBM script.....	81
Removing existing SPBM configuration.....	83
Configuring the IS-IS port interfaces using SPBM script.....	84
Removing specific IS-IS and MLT interfaces.....	85
Configuring minimum SPBM and IS-IS parameters.....	86
Configuring I-SIDs for private VLANs.....	90
Displaying global SPBM parameters.....	91
Displaying global IS-IS parameters.....	93
Enabling IP Multicast over Fabric Connect globally.....	95
Displaying IP Multicast over Fabric Connect information.....	97
Displaying IS-IS areas.....	100
Configuring SMLT parameters for SPBM.....	101
Configuring optional SPBM parameters.....	103
Configuring optional IS-IS global parameters.....	105
Configuring optional IS-IS interface parameters.....	110
Displaying IS-IS interface parameters.....	112
Displaying the IP unicast FIB, multicast FIB, unicast FIB, and unicast tree.....	114
Displaying IS-IS LSDB and adjacencies.....	119
Displaying IS-IS statistics and counters.....	123
Running the vms endura script.....	125
Fabric Extend configuration using the CLI.....	127
Fabric Attach configuration using the CLI.....	135
IS-IS external metric configuration using the CLI.....	165
SPBM and IS-IS infrastructure configuration using EDM.....	169
Configuring required SPBM and IS-IS parameters.....	169
Configuring IP Multicast over Fabric Connect globally.....	174
Modifying IP Multicast over Fabric Connect globally.....	175
Displaying IP Multicast over Fabric Connect routes.....	176
Displaying the UNI ports for IP multicast routes.....	177
Displaying SPBM and IS-IS summary information.....	178
Displaying the SPBM I-SID information.....	179
Displaying Level 1 Area information.....	180
Configuring SMLT parameters for SPBM.....	180
Enabling or disabling SPBM at the global level.....	181
Configuring SPBM parameters.....	182
Displaying SPBM nicknames.....	183
Configuring interface SPBM parameters.....	184
Configuring SPBM on an interface.....	184
Displaying the IP unicast FIB.....	185
Displaying the IPv6 unicast FIB.....	186
Displaying the unicast FIB.....	187

Displaying the multicast FIB.....	188
Displaying LSP summary information.....	189
Displaying IS-IS adjacencies.....	189
Configuring IS-IS global parameters.....	190
Configuring system-level IS-IS parameters.....	192
Displaying IS-IS system statistics.....	193
Configuring IS-IS interfaces.....	194
Configuring IS-IS interface level parameters.....	195
Displaying IS-IS interface counters.....	196
Displaying IS-IS interface control packets.....	197
Graphing IS-IS interface counters.....	198
Graphing IS-IS interface sending control packet statistics.....	199
Graphing IS-IS interface receiving control packet statistics.....	200
Configuring an IS-IS Manual Area.....	201
Fabric Extend configuration using EDM.....	201
Fabric Attach configuration using the EDM.....	206
SPBM configuration examples.....	223
Basic SPBM configuration example.....	224
Ethernet and MLT configuration.....	224
IS-IS SPBM global configuration.....	225
IS-IS SPBM Interface Configuration.....	226
IP multicast over Fabric Connect global configuration.....	227
Verifying SPBM operations.....	228
Fabric Extend configuration examples.....	230
Fabric Extend over IP using the GRT.....	230
Fabric Extend over IP using a VRF.....	233
Fabric Extend over VPLS.....	236
Fabric Extend over Layer 2 Pseudowire.....	239
Fabric Extend with ONAs in the core and branches.....	241
Fabric Attach configuration examples.....	244
Configuring a Fabric Attach solution.....	244
Configuring Fabric Attach in an SMLT.....	249
Chapter 4: SPBM and IS-IS services configuration.....	258
Fabric Connect Service Types.....	258
Layer 2 VSN configuration.....	259
Layer 2 VSN configuration fundamentals.....	259
Layer 2 VSN configuration using the CLI.....	271
Layer 2 VSN configuration using EDM.....	314
Layer 2 VSN configuration examples.....	325
IP Shortcuts configuration.....	330
IP Shortcuts configuration fundamentals.....	330
IP Shortcuts configuration using the CLI.....	349
IP Shortcuts configuration using EDM.....	385

IP Shortcuts SPBM configuration example.....	400
IP multicast over SPBM within the GRT configuration example.....	403
Layer 3 VSN configuration.....	403
Layer 3 VSN configuration fundamentals.....	403
Layer 3 VSN configuration using the CLI.....	418
Layer 3 VSN configuration using EDM.....	457
Layer 3 VSN configuration example.....	472
Enable/disable ICMP Response on VRFs/L3 VSNs.....	480
Inter-VSN routing configuration.....	481
Inter-VSN routing configuration fundamentals.....	481
Inter-VSN routing configuration using the CLI.....	482
Inter-VSN routing configuration using EDM.....	484
Inter-VSN routing configuration example.....	494
Chapter 5: Operations and Management.....	498
CFM fundamentals.....	498
Maintenance Domain (MD).....	500
Maintenance Association (MA).....	500
Maintenance association endpoints (MEP).....	501
Fault verification.....	502
LBM message.....	502
I2ping.....	502
Fault isolation.....	503
Link trace message.....	503
I2tracert.....	504
I2 tracertree.....	505
Maintenance domain intermediate points (MIP).....	505
Layer 2 tracemroute.....	506
Nodal MPs.....	506
Configuration considerations.....	507
CFM configuration using CLI.....	508
Autogenerated CFM.....	509
Configuring explicit mode CFM.....	514
Triggering a loopback test (LBM).....	520
Triggering linktrace (LTM).....	521
Triggering a Layer 2 ping.....	523
Triggering a Layer 2 traceroute.....	525
Triggering a Layer 2 tracertree.....	526
Triggering a Layer 2 tracemroute.....	527
Using trace CFM to diagnose problems.....	530
Using trace SPBM to diagnose problems.....	533
CFM configuration using EDM.....	535
Autogenerated CFM.....	536
Configuring explicit CFM.....	539

Configuring Layer 2 ping.....	543
Initiating a Layer 2 traceroute.....	545
Viewing Layer 2 traceroute results.....	548
Configuring Layer 2 IP ping.....	549
Viewing Layer 2 IP Ping results.....	551
Configuring Layer 2 IP traceroute.....	552
Viewing Layer 2 IP traceroute results.....	554
Triggering a loopback test.....	556
Triggering linktrace.....	558
Viewing linktrace results.....	560
Configuring Layer 2 tracetable.....	562
Viewing Layer 2 tracetable results.....	565
Configuring Layer 2 trace multicast route on a VLAN.....	566
Configuring Layer 2 tracemroute on a VRF.....	568
Viewing Layer 2 trace multicast route results.....	570
CFM configuration example.....	571
CFM configuration example.....	571
CFM sample output.....	572
Chapter 6: Resources	576
Support.....	576
Documentation.....	576
Training.....	576
Viewing Avaya Mentor videos.....	576
Searching a documentation collection.....	577
Subscribing to e-notifications.....	578
Appendix A: SPBM reference architectures	581
Reference architectures.....	581
Solution-specific reference architectures.....	590
Glossary	596

Chapter 1: Introduction

Introduction

Purpose

This document provides information on features in VSP Operating System Software (VOSS). VOSS runs on the following product families:

- Avaya Virtual Services Platform 4000 Series
- Avaya Virtual Services Platform 7200 Series
- Avaya Virtual Services Platform 8000 Series

This document provides instructions to configure the Avaya Fabric Connect operations on the switch. Operations include Shortest Path Bridging MAC (SPBM), Intermediate System to Intermediate System (IS-IS), and Connectivity Fault Management (CFM).

Using this document

This document is organized into three main sections:

1. Infrastructure configuration — You must first configure your base SPBM and IS-IS architecture described in [SPBM and IS-IS infrastructure configuration](#) on page 14. This chapter includes initial steps to configure the minimum SPBM and IS-IS parameters to enable SPBM on your network. For more information, see [Configuring minimum SPBM and IS-IS parameters](#) on page 86.
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 include: Layer 2 VSNs, IP Shortcuts, Layer 3 VSNs, *Transparent Port UNI* (T-UNI), and Inter-VSN routing described in [SPBM and IS-IS infrastructure configuration](#) on page 14
3. Operations and management — [Operations and Management](#) on page 498 provides tools to monitor and troubleshoot your SPBM network.

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

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Chapter 2: New in this document

The following sections detail what is new in *Configuring Fabric Connect*.

! **Important:**

The features in this document might not apply to all hardware platforms. For more information about feature support, see *Release Notes*.

Release 6.0.1

Fabric Attach system ID

In an SMLT configuration, you must delete the FA on SMLT before deleting the SPBM instance. For more information, see [Fabric Attach considerations](#) on page 79.

SPB path tiebreakers

See [Link metric](#) on page 29 for tiebreaking information on how SPB paths are chosen.

Release 6.0

Distributed Virtual Routing (DvR)

Distributed Virtual Routing (DvR) is a technology for optimizing traffic flows in a distributed switching and routing architecture. DvR optimizes traffic flows to avoid traffic tromboning due to inefficient routing, thereby increasing the total routing throughput.

DvR also simplifies large scale data center deployments by introducing a Controller-Leaf architecture. In this architecture, Layer 3 configuration is required only on the Controller nodes, whereas the Leaf nodes require only Layer 2 configuration. All Layer 3 configuration is automatically distributed to the Leaf nodes by the Controller nodes. The Controller and Leaf nodes form a logical group called the DvR domain.

DvR and IS-IS accept policies

The DvR backbone is automatically established among the Controller nodes from all DvR domains. You can configure accept policies on the Controller nodes or a non-DvR BEBs, as a filter to determine which DvR host routes can be learned from the DvR backbone.

For information on IS-IS accept policies, see [IS-IS accept policies](#) on page 343.

For CLI configuration of IS-IS accept policies to accept host routes from the DvR backbone, see [Configuring IS-IS accept policies](#) on page 356.

For EDM configuration, see:

- [Configuring an IS-IS accept policy for a specific advertising BEB](#) on page 391
- [Configuring an IS-IS accept policy to apply for a specific I-SID](#) on page 392
- [Configuring an IS-IS accept policy for a specific advertising BEB and I-SID](#) on page 393
- [Configuring an IS-IS accept policy for a specific I-SID list](#) on page 395
- [Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list](#) on page 396
- [Applying IS-IS accept policies globally](#) on page 390

DvR and IPv4 Multicast over Fabric Connect

IPv4 Multicast must be enabled globally on all DvR enabled nodes in a DvR domain.

Also, when you configure IP Multicast over Fabric Connect within the GRT or on a Layer 3 VSN (VRF) on the DvR Controllers, the configuration information is automatically pushed to all DvR Leaf nodes within the domain.

For more information, see:

- [IP Multicast over Fabric Connect](#) on page 34
- [IP Multicast over Fabric Connect within the GRT](#) on page 348
- [Layer 3 VSN with IP Multicast over Fabric Connect](#) on page 416

The corresponding CLI and EDM configuration procedures in this document are updated.

For more information on DvR, see *Configuring IPv4 Routing*.

Fabric Attach Zero Touch Client Attachment

Fabric Attach Zero Touch Client Attachment provides the ability for an FA client to automatically attach to an existing Shortest Path Bridging Network and provide for automatic configuration of the service identifier (I-SID) and virtual LAN based on FA client element type. FA clients must signal the desire to join an SPB network through the use of specific LLDP TLVs.

For more information, see [FA Zero Touch Client Attachment](#) on page 64.

Fabric Extend IP over ELAN/VPLS enhancement

This release removes the single next hop / ARP restriction on VSP 7200 and VSP 8000 series switches. This feature allows multiple switches running Fabric Extend IP to be directly connected over a Layer 2 broadcast domain without the need for loopback VRFs.

Increased VRF and Layer 3 VSN scaling

You can now use a boot config flag to increase the number of Virtual Routing and Forwarding (VRF) instances on the switch from the previous maximum of 24. This enhancement also impacts the number of Layer 3 Virtual Services Networks (VSN). The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform.

For more information about maximum scaling numbers, see *Release Notes*.

! Important:

If you use the boot config flag to increase the number of VRFs and Layer 3 VSNs, and the switch operates in SPBM mode, the switch reduces the number of configurable VLANs.

A Premier or Premier + MACsec license is required to use more than 24 VRFs.

For more information, see [Configuring IS-IS accept policies](#) on page 356.

For information about how to increase the number of supported VRFs, see *Configuring IPv4 Routing*.

Network Load Balancing (NLB) Multicast operation

When you enable NLB multicast mode on a VLAN, the routed traffic destined to the NLB cluster is flooded by default on all ports of the VLAN. All VLANs support multiple cluster IPs by default.

Multicast MAC flooding and static multicast ARP entries are not supported for NLB Unicast or NLB Multicast in this release.

***** Note:

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

For more information, see:

- [Network Load Balancing \(NLB\)](#) on page 38
- [IP Multicast over Fabric Connect restrictions](#) on page 42
- [Configuring BCBs for Inter-VSN routing](#) on page 489
- [Configuring BEBs for Inter-VSN routing](#) on page 485

NNI-MSTP Boot Flag

The default value for the nni-mstp flag is false. In previous releases, MSTP was enabled for the CIST and all MSTIs other than MSTI-62. In the current release, the default behavior of the MSTP on SPBM NNI ports is that CIST is disabled automatically on the NNI and the NNI ports cannot be members of any VLANs other than B-VLANs. You can override this by setting the nni-mstp flag to true, which disables MSTP on MSTI-62, and allows any VLAN to be configured on NNI ports.

On SPBM NNI links, MSTP is disabled and no VLAN, except SPBM B-VLANs can be added. When nni-mstp flag is set to true, it only disables MSTI 62 and additional VLANs on other MSTIs can be added to NNI links.

For more information, see [SPBM restrictions and limitations](#) on page 32.

Resources

Information about related resources is moved to the last chapter in this document.

SPB Ethertype — change in behavior on NNI

SPB switches now follow the configured Ethertype on egress from NNI interfaces. On ingress the switches will honor Ethertype of either 0x8100 and 0x88a8.

For more information, see [SPB Ethertype](#) on page 80.

VXLAN Gateway

VXLAN Gateway is a hardware-based virtual tunnel end point (VTEP) that terminates virtual extensible LAN (VXLAN) tunnels. The VXLAN tunnels “stretch” emulated Layer 2 segments over an IP network. Each VTEP has at least one segment ID called a VXLAN Network Identifier (VNID). This VNID mechanism allows up to *16 million* VXLAN segments to coexist within the same administrative domain. Each VTEP can support multiple VNIDs.

For information about maximum scaling numbers, see *Release Notes*.

For more information, see [vxlan-gw-full-interworking-mode boot flag](#) on page 18.

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 the enterprise data center. 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 (BVLANS) 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 B-MAC reachability.

- **Multicast**

- With Layer 2 VSN with IP multicast over Fabric Connect, 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 Fabric Connect, 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 Fabric Connect, 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 Fabric Connect is not supported.

The switch supports the IEEE 802.1aq standard of SPBM, which allows for larger Layer 2 topologies and permits faster convergence.

Multiple tenants using different SPBM services

The following figure shows multiple tenants using different services within an SPBM metro network. In this network, you can use some or all of the SPBM implementation options to meet the needs of the community while maintaining the security of information within VLAN members.

Large Campus/
Multi-tenant

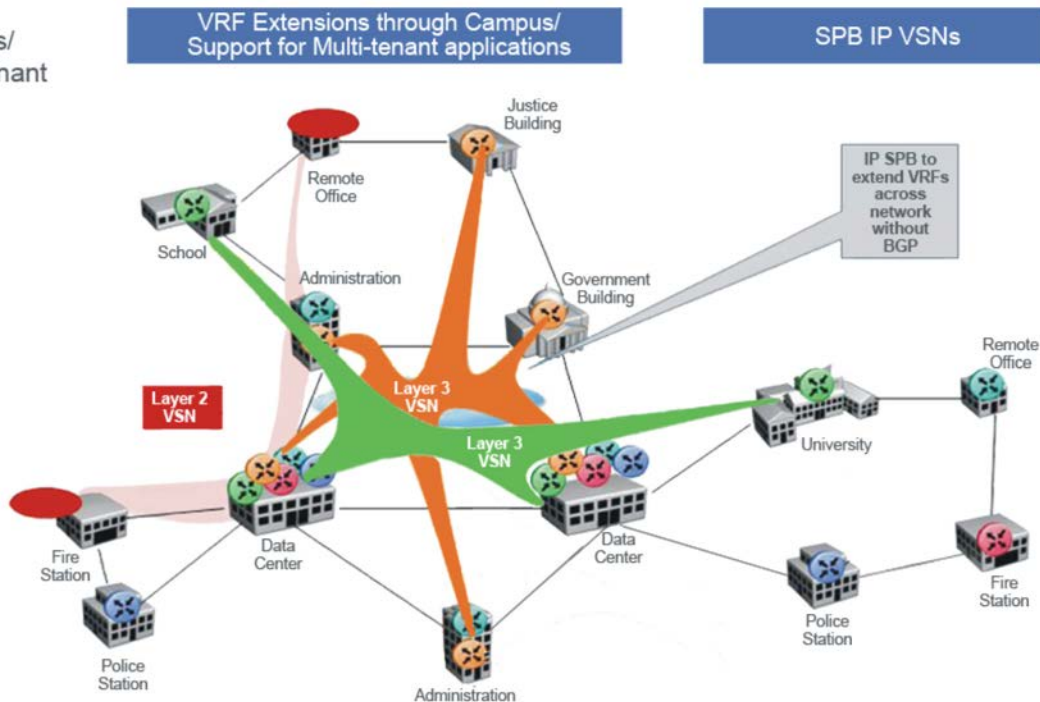


Figure 1: Multi-tenant SPBM metro network

To illustrate the versatility and robustness of SPBM even further, the following figure shows a logical view of multiple tenants in a ring topology. In this architecture, each tenant has its own domain where some users have VLAN requirements and are using Layer 2 VSNs and others have VRF requirements and are using Layer 3 VSNs. In all three domains, they can share data center resources across the SPBM network.




Figure 2: SPBM ring topology with shared data centers

spbm-config-mode boot flag

Shortest Path Bridging (SPB) and Protocol Independent Multicast (PIM) cannot interoperate with each other on the switch at the same time. The boot flag called `spbm-config-mode` ensures that SPB and PIM stay mutually exclusive.

- The `spbm-config-mode` boot flag is enabled by default. This enables you to configure SPB and IS-IS, but you cannot configure PIM either globally or on an interface.
- If you disable the boot flag, save the configuration and reboot with the saved configuration. After you enable the flag, you can configure PIM and IGMP Snooping, but you cannot configure SPB or IS-IS.

! Important:

- Any change to the `spbm-config-mode` boot flag requires a reboot for the change to take effect.
- If you disable the boot flag, save the configuration and reboot with the saved configuration. After you disable the flag, you can configure PIM and IGMP Snooping, but you cannot configure SPB or IS-IS.

For more information, see *Configuring IP Multicast Routing Protocols*.

vxlan-gw-full-interworking-mode boot flag

The VXLAN Gateway implementation is available in the following modes:

- **Base Interworking Mode** – This is the default mode. In this mode, VXLAN Gateway supports Layer 2 gateway communication between VXLAN and traditional VLAN environments.
- **Full Interworking Mode** – This mode supports the Base mode communication between VXLAN and traditional VLAN environments as well as VXLAN-to-VXLAN communication and all SPB functionality including vIST and SMLT. To enter this mode, you must enable the `vxlan-gw-full-interworking-mode` boot configuration flag.

*** Note:**

Changing the mode requires a reboot for the change to take effect.

For complete information about this feature, see *Configuring VLANs, Spanning Tree, and NLB*.

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 B-MAC source address (BMAC-SA) and a B-MAC destination address (BMAC-DA) 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.

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 in access networks do not impact forwarding results in 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 Backbone Edge Bridges (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 Backbone Core Bridges (BCBs). BCBs forward encapsulated traffic based on the BMAC-DA.

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 a service instance for 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:

- 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 Fabric Connect, 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 *Transparent Port UNI*, the I-SID is associated with a port or MLT, which is then virtualized across the backbone. *Transparent Port UNI* associates multiple ports or MLT to an I-SID.
- 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 Fabric Connect, 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 Fabric Connect, 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.

*** Note:**

I-SID to VLAN binding is used to automatically determine the path between client and server in order to attach network devices to FA Zero touch services.

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 B-MAC-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 (VSN). 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.

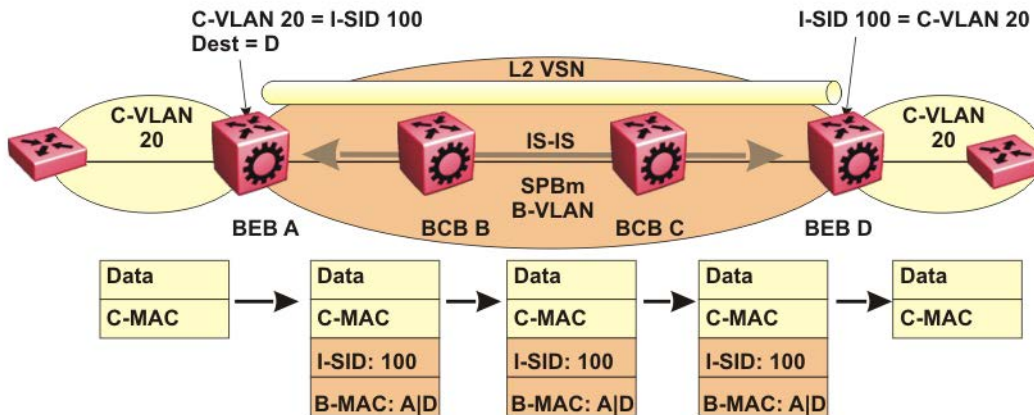


Figure 3: SPBM L2 VSN

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.

E-Tree and Private VLAN topology

Ethernet Private Tree (E-Tree) extends Shortest Path Bridging MAC (SPBM) to Private VLANs (PVLAN).

Transport within the SPBM network is achieved by associating the private VLAN with an I-SID. Flooded traffic from both promiscuous and isolated devices is transported over the same I-SID multicast tree and suppression for spoke-to-spoke traffic is done on the egress SPB Backbone Edge Bridge (BEB). This means the Private VLAN IDs are globally significant and must be the same on all BEBs

The following list provides details for E-Tree and Private VLAN topology:

- E-Tree associates a Private VLAN with an I-SID.

*** Note:**

The same I-SID could be attached to a regular VLAN. In that case, all ports on the regular VLAN behave like Promiscuous ports on the PVLAN.

- Other SPB BEBs can associate a regular CVLAN to the same I-SID that E-Tree uses.

*** Note:**

The CVLAN ID must match the primary PVLAN ID.

- CVLAN devices assigned to the same I-SID that E-Tree uses have Promiscuous connectivity within the segment.

The following figure shows a basic E-Tree network topology consisting of groups of private VLANs connected by the SPBM core network.




Figure 4: Sample E-Tree configuration

Private VLAN port types

The private VLAN port type is isolated, promiscuous, or trunk. If the port is a member of an MLT, then the port inherits the private VLAN type of the MLT.

In terms of network topology, the isolated port is considered a spoke. The isolated port, or spoke, does not communicate with any other isolated port in the network. The isolated port only communicates with the promiscuous ports, or hubs.

E-Tree and Private VLAN limitations

The following limitations apply to E-Tree and Private VLAN topology:

- A port that is of Private VLAN type trunk must be tagged. Isolated and Promiscuous Private VLAN ports can be either tagged or untagged.
- When a port or MLT that has a Private VLAN type set to Isolated or Promiscuous is added to a private VLAN, if that port is used by other non private VLANs, then those non private VLANs are removed.
- A port which is Private VLAN type Isolated and is tagged can belong to only one Private VLAN.

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 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).

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. The switch also supports Sub-TLVs. TLVs exist inside IS-IS packets and Sub-TLVs exist as additional information in TLVs.

The switch supports and is in full compliance with standard 802.1aq 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. The following table lists all the TLVs that the switch supports.

Table 1: 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	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

Table continues...

TLV	Description	Usage
	<p>environment. This TLV is carried within IS-IS Hello Packets (IIH).</p>	<p>(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.</p> <ul style="list-style-type: none"> • MCID Sub TLV: The MCID is a digest of the VLANs and MSTI. Neighboring SPBM nodes must agree on the MCID to form an adjacency. The MCID is set to all zeros (0). <p>After the switch receives a non-zero MCID Sub TLV, it reflects content back to the neighbor.</p>
144	<p>Multi-topology Capability (MT-Capability) TLV.</p> <p>This TLV carries the SPB instance ID in a multiple SPB instance environment. This TLV is carried within LSPs.</p> <p>In multicast over Fabric Connect, 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.</p>	<p>TLV 144 is the service identifier TLV. TLV 144 advertizes B-MAC and I-SID information.</p> <p>This TLV carries the following Sub TLVs:</p> <p>Sub-TLV 1: SPB instance Sub TLV contains a unique SPSourceID (nickname) to identify the SPBM node within this SPB topology.</p> <p>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.</p>
184	<p>SPBM IP VPN reachability — IS-IS TLV 184 is used to advertise SPBM L3 VSN route information across the SPBM cloud.</p>	<p>IP reachability for Layer 3 VSNs</p>
185	<p>IPVPN multicast TLV with IPMC sub TLV — The IPVPN multicast TLV contains information about the scope I-SID.</p>	<p>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-</p>

Table continues...

TLV	Description	Usage
		<p>SID that maps to the source and group addresses.</p> <p>As part of the IPVPN TLV, sub-TLVs define IPv4 unicast, IPv6 unicast and IPv4 multicast information.</p> <p>Layer 2 VSN IP multicast over Fabric Connect and Layer 3 VSN IP multicast over Fabric Connect (using VRF) use TLV 185.</p>
186	<p>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.</p>	<p>IP Shortcuts with IP multicast over Fabric Connect use TLV 186.</p> <p>All multicast streams are constrained within the level in which they originate, which is called the scope level.</p>
236	<p>IPv6 Reachability — The IPv6 reachability TLV 236 is used to distribute IPv6 network reachability between IS-IS peers.</p>	<p>SPBM uses the existing IS-IS TLV to carry IPv6 shortcut routes through the SPBM core.</p>

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.

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. 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 switch software only supports one manual area.
- **NSEL** — The last byte (00) is the n-selector. In this 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-lsp-interval` 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-lsp-interval`, 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. The switch 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

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

The switch software supports only interface level authentication. The switch software 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 switch 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.

- The switch only supports the wide metric.
- The total cost of a path equals the sum of the cost of each link.
- The default value for wide metrics is 10.

* Note:

When multiple paths exist to reach a node, the path with the lowest sum of metrics of the individual links is chosen. If the sum of the paths are the same, the one with the lowest number of hops is chosen. If the number of hops is the same as well, then the tie-breaking is done by the system ID.

For the primary BVLAN, the path that has a node with the lowest system ID is chosen. Whereas, for the secondary BVLAN, the path that has a node with the highest system ID is chosen.

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

The overload bit is sent by a node in LSP updates to inform other devices, whether to use that node to pass transit traffic. For example, when an LSP with an overload bit is received, the device ignores that LSP in its SPF calculation to avoid sending transit traffic through the overloaded node; however the overloaded node can still receive traffic destined to itself.

The overload bit is turned on by default on bootup, and cleared after 20 seconds. You can use the `overload-on-startup` parameter to control the time before the overload bit is cleared after bootup, as this setting is user configurable.

You can permanently set the overload bit using the `overload` parameter. If this is configured, the overload bit will not be cleared after bootup and it will be sent in all LSP updates. If the overload bit is set permanently, other devices do not include this node for use as a transit node in IS-IS computations. By default, the `overload` parameter is set to false.

The `overload` and `overload-on-startup` parameters are two independent settings and are configured under the `router isis` configuration mode in the CLI.

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:**

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 disabled

You cannot add ports to a B-VLAN manually, IS-IS enabled ports are automatically added 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.

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

```
=====
                        SPBM UNICAST FIB ENTRY INFO
=====
```

DESTINATION ADDRESS	BVLAN	SYSID	HOST-NAME	OUTGOING INTERFACE	COST
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	1/12	0
00:e0:7b:84:57:df	11	00e0.7b84.57df	86-30	1/12	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 SPBM Node Nickname and the I-SID ID converted to hexadecimal format to form the multicast MAC address.

```
|-----|-----|
      nickname|0x30000                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(config)#show isis spbm multicast-fib

=====
                        SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA                ISID      BVLAN SYSID      HOST-NAME      OUTGOING-INTERFACES      INCOMING
                           INTERFACE
-----
03:00:07:e4:e2:02      15000066 1001  0077.0077.0077  Switch-25      1/33                MLT-2
03:00:08:e4:e2:02      15000066 1001  0088.0088.0088  Switch-33      1/50,1/33          40.40.40.40
03:00:41:00:04:4d      1101      4058  00bb.0000.4100  Switch-1(*)    1/3,1/49,0.0.0.0  Tunnel_to_HQ
03:00:41:00:04:4f      1103      4058  00bb.0000.4100  Switch-1(*)    1/3,1/49,0.0.0.0  cpp
=====
Total number of SPBM MULTICAST FIB entries 4
=====
```

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. You should always use Simple Loop Prevention Protocol (SLPP) in an SMLT environment.

*** Note:**

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 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.
- You must configure SPBM B-VLANs on all devices in the same MSTP region. MSTP requires this configuration to generate the correct digest.
- Configure the SPBM B-VLANs to use matching VLAN IDs.

Best practices for SPB regarding MSTP

Use NNI ports exclusively to transport traffic for SPB-based services and not be configured as members of any VLANs other than SPB B-VLANs. In releases that do not support `nni-mstp`, when an SPBM IS-IS interface is created on an NNI port or an MLT, MSTP is automatically disabled for MSTI-62 on the port/MLT. However, MSTP is not automatically disabled on NNI ports for the CIST (default MSTI). In releases that support the `boot config flags nni-mstp` command, the default behavior of the MSTP NNI ports is that CIST is disabled automatically on the NNI and the NNI ports cannot be members of any VLANs other than B-VLANs. The default `boot config flags nni-mstp` must be set to `false` (which is the default). The following example shows the command to disable the MSTP on the NNI ports.

```
Switch:1(config)#interface gigabitEthernet 1/8
Switch:1(config-if)#no spanning-tree mstp
```

Coexistence of MSTP and SPB based services on NNI ports:

In releases that do not support `nni-mstp` boot configuration, you can support the coexistence of non-SPB based services on the NNI ports, by adding NNI ports as members of VLANs, except for B-VLANs. These other VLANs rely on the use of MSTP for Loop prevention. The network operator must carefully consider the implications of keeping MSTP enabled on the NNI ports because any MSTP topology changes detected on the NNI ports impacts all services and causes most dynamically learned information on the UNI side to be flushed and relearned. This includes, but is not limited to, all customer MAC and ARP records. This can also cause all the UNI ports on a BEB to be temporarily put into a spanning-tree blocking state before transitioning to a forwarding state again. The net result is that MSTP topology changes on the NNI ports adversely impact traffic for SPB-based services. Therefore, it is recommended that the NNI ports be used exclusively for SPB traffic.

SPBM IS-IS

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

- The switch does not support IP over IS-IS as defined by RFC 1195. IS-IS protocol is only to facilitate SPBM.
- The switch uses level 1 IS-IS. The switch does not support level 2 IS-IS. The CLI command `show isis int-12-cont1-pkts` is not supported 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 switch supports the wide metric.
- IS-IS enabled ports cannot be added to an MLT. The switch does not support this configuration.

SPBM NNI SMLT

The switch does not support NNI on SMLT links.

VLACP

VLACP is generally used when a repeater or switch exists between connected 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.

System MTU

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

IP multicast over Fabric Connect

IP multicast over Fabric Connect cannot connect to existing Protocol Independent Multicast (PIM) networks that connect to SPB originated streams or that add PIM network streams into the SPB network. SPB-PIM Gateway (SPB-PIM GW), however, provides multicast interdomain communication between an SPB network and a PIM network. SPB-PIM GW accomplishes this interdomain communication across a special Gateway VLAN. The Gateway VLAN communicates with the PIM network through the PIM protocol messaging and translates the PIM network requirements into SPB language, and vice versa. For more information about SPB-PIM GW, see *Configuring SPB-PIM Gateway*.

Other

The following list identifies other restrictions and limitations:

- The software does not support I-SID filters.
- You cannot enable C-VLAN and B-VLAN on the same port.

IP Multicast over Fabric Connect

Avaya leads the industry with a new approach to transporting IP multicast using IP Multicast over Fabric Connect. IP Multicast over Fabric Connect 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 Fabric Connect 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 Fabric Connect 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.

IP Multicast over Fabric Connect introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. IP Multicast over Fabric Connect 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 Fabric Connect, 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 Fabric Connect). 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 Fabric Connect). Example applications: Video surveillance, TV/Video/Ticker/Image distribution, VX-LAN.
- Layer 3 Virtual Services Network with VRF based routing support for IP Multicast over Fabric Connect in the core and IGMP on the access (Layer 3 VSN with IP Multicast over Fabric Connect). Example applications: Video surveillance, TV/Video/Ticker/Image Distribution, VX-LAN, Multi-tenant IP multicast.

IP Multicast over Fabric Connect and DvR

You must enable IP multicast over Fabric Connect globally on all DvR enabled nodes (Controllers and Leaf nodes) in a DvR domain.

For more information on DvR, see *Configuring IPv4 Routing*.

How IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect 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 5: IP Multicast over Fabric Connect 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 Fabric Connect 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 Fabric Connect 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 Fabric Connect 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 IP Multicast over Fabric Connect, 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, 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.

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 conceptual and configuration information on IGMP, see *Configuring IP Multicast Routing Protocols*.

Network Load Balancing (NLB)

SPBM supports Network Load Balancing (NLB) Unicast and Multicast modes.

*** Note:**

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

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, Spanning Tree, and NLB*.

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 Fabric Connect 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 6: IP Multicast over Fabric Connect 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 vIST 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 Fabric Connect, 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 vIST 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 vIST 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 vIST 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 Fabric Connect and SMLT

The following section summarizes the IP Multicast over Fabric Connect 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 ERS 8800, VSP 4000 Series, VSP 7200 Series, and VSP 8000 Series, 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 vIST peers do not see the membership queries of each other:

- The vIST peer with the higher IP address sends the queries out all SMLT and non-SMLT ports on SPBM access links.
- The vIST 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 vIST peer with a higher IP address and an vIST 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 vIST peer with the higher IP address as the querier, and non-SMLT access devices see the directly connected vIST peer as the querier. Non-SMLT access devices that connect on either side of the vIST peers can talk to each other using the SPBM cloud.

Considerations when you connect an IP Multicast over Fabric Connect network to a PIM network

IP Multicast over Fabric Connect 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 support a configuration option to accept streams with non-local sources.

You must configure static IGMP receivers on the PIM interface that face the IP Multicast over Fabric Connect 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 Fabric Connect restrictions

Review the following restrictions for the IP Multicast over Fabric Connect 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 Fabric Connect.

SPBM supports Network Load Balancing (NLB) unicast and multicast modes. SPBM does not support NLB Multicast operation with IGMP.

*** Note:**

The NLB Multicast operation feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

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 232.122.122.128 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 Fabric Connect 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 IP Multicast over Fabric Connect 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.

Supported services

The switch does not support IP Multicast over Fabric Connect routing on inter-VSN routing interfaces.

The switch supports the following modes of IP Multicast over Fabric Connect:

- 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 Fabric Connect enabled.

SPBM script

You can use a CLI script to quickly configure the SPB and IS-IS infrastructure to enable Fabric Connect on a switch. You can use the SPB script, rather than manually configure the minimum SPBM and IS-IS parameters.

You can use the command `run spbm` to quickly configure the following:

- Configure the SPB Ethertype.
- Create an SPB instance.
- Create an SPBM backbone VLAN and associate it to the SPB instance.
- Create an SPBM secondary backbone VLAN and associate it to the SPB instance.
- Add an SPB nickname.
- Create a manual area.
- Enable IS-IS on one of the switch interfaces.
- Enable IS-IS globally.

- Configure the IS-IS system name.
- Configure the IS-IS system ID.

The following table displays the default values applied if you use the `run spbm` command. The SPB script creates some of the default values based on the MAC address of the switch, including the nickname and System ID value.

Parameter	Default values
Ethertype	0x8100
Primary BVLAN	4051
Secondary BVLAN	4052
Manual area	49.0000
Nickname	Derived from the chassis MAC
System name	Derived from the command line prompt
System ID value	Derived from the chassis MAC, using a different algorithm from that for the Nickname

*** Note:**

The SPB script only creates the SPBM instance, VLAN, or other parameters if they do not already exist. For example, if the SPBM instance and VLAN already exist, the SPB script does not create them. If the SPB script cannot create one of the parameters because the parameter is already configured, the script stops and an error message displays.

Run vms endura script

*** Note:**

Only Avaya Virtual Services Platform 4000 Series supports this feature in the current release.

The `run vms endura switch` CLI command executes a script that pre-configures basic and common configuration parameters to quickly and easily deploy a Pelco Endura Video Surveillance network in accordance with best practices using networking equipment.

The `run vms endura` CLI script is specifically targeted for use with the VSP 4450GSX-PWR+ and VSP 4850GTS/GTS-PWR+ products. The script creates a switch configuration specifically tailored to the default IP address subnet ranges and DHCP services of a Pelco Endura Systems Manager 5000 (SM5000) Video Management System (VMS). This enables customers to use a single CLI command on a VSP 4450 series switch to setup the core switch configuration where the VMS/video surveillance management and operation systems reside, and also use of a single CLI command to setup each edge VSP 4850 series switch where IP cameras are connected.

A VSP 4000 series switch should be in a factory default state before the `run vms endura` CLI command is used to ensure correct operation of the configuration. A VSP 4450GSX-PWR+ switch should be used as the core switch and each VSP 4850GTS or VSP 4850GTS-PWR+ switch must be used as edge/access switches in the VMS solution. Typically, the edge/access switch will be a VSP 4850GTS-PWR+ PoE switch to connect and power IP video surveillance cameras.

Essentially, the run vms endura script creates an SPB network core solution with IP shortcuts to connect IP subnet “zones” between the core and edge IP subnets. All network edge IP subnet areas containing IP cameras are configured with an IP gateway address that is redistributed over the SPB fabric, so the fabric core acts as a single IP routing entity for the solution. DHCP services are relayed between each IP subnet area and the central server for IP camera address allocation.

IP multicast over Fabric Connect virtualization is also enabled to support and allow efficient IP multicast communication over the fabric core from IP cameras to central VMS servers for viewing and recording video streams.

CLI Command ‘switch x’ value

The ‘switch’ value used in the CLI command must be between 5 and 99 inclusively, where the value 5 must be used in the script executed on a switch located in the core where the Endura VMS core systems are connected.

It is recommended that a VSP 4450GSX-PWR+ be used as the core switch.

Subsequent ‘switch’ values used when the script is executed on switches located at the edge/access layer must be between 6 and 99, and be unique for each additional switch that is part of the solution.

For example, the first edge/access switch with IP Cameras connected would use a value of ‘switch 6’. For additional edge/access switch you would then use ‘switch 7’, ‘switch 8’ and so on for each IP subnet and IP camera zone. Up to 48 IP cameras can be connected to a switch within an IP subnet zone.

Switch parameters configured by the script

The following major parameters are configured by the run vms endura command and utilize the switch “#” value in the command to set up parameters:

- SNMP-Server switch hostname
- SPB parameters such as; System ID, Nickname, SPB Area ID, Backbone VLAN ID’s (4051 and 4052), Multicast virtualization and CFM.
- IP loopback interface addresses
- IP redistribution over IS-IS (IP Shortcuts)
- All SFP ports as SPB NNI ports
- All copper RJ-45 ports as end device ports with Spanning Tree enabled
- Spanning Tree mstrstp mode
- VLAN port memberships
- VLAN IP address (Gateway IP for VLAN)
- DHCP Relay

Note:

DHCP Relay parameters are only configured when the script is run on VSP 4850GTS and VSP 4850GTS-PWR+ switches.

Configuration file

Upon successful completion of the run vms endura script, the switch configuration is saved with a filename based on the 'switch' value used when the script is executed. The switch primary boot config file flags are updated with the new filename.

For example, executing the command `run vms endura switch 5` will result in a switch configuration filename of `spb-switch-5.cfg`.

Fabric Extend

Fabric Extend enables Enterprises to extend the Fabric Connect technology over Layer 2 or Layer 3 core networks. The *logical IS-IS interface*, which is discussed in detail later in this chapter, is the mechanism that enables Fabric Extend to connect SPB fabric nodes. Logical IS-IS interfaces create virtual tunnels and encapsulate SPB traffic by adding a Virtual Extensible LAN (VXLAN) header to SPB packets.

The following figure illustrates two Fabric Connect “islands” separated by a third-party core IP network. The IP network could be third-party equipment in an enterprise or a service provider’s infrastructure such as an MPLS VPN service.

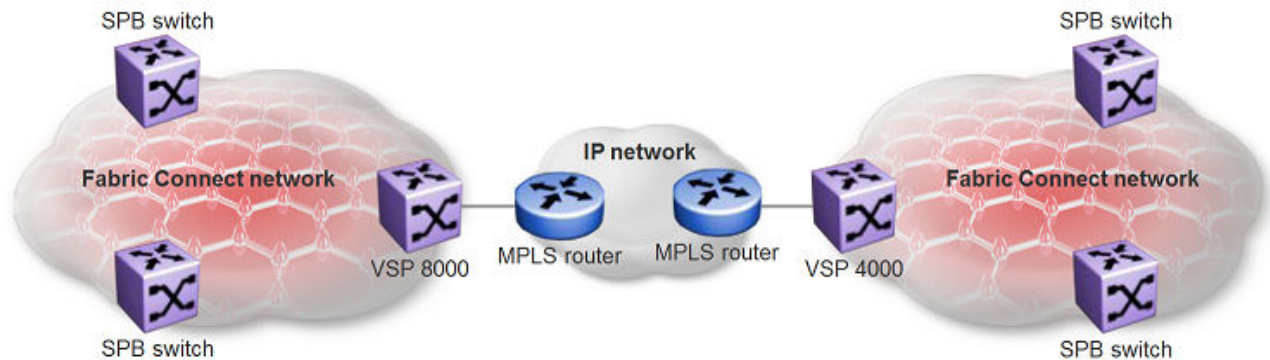


Figure 7: Fabric Connect networks connected by an IP network

The following figure illustrates how Fabric Extend enables you to connect the fabric islands to create ONE Fabric Connect network. This figure shows a layer 3 core network where Fabric Extend uses IP tunneling by adding a VXLAN header to the SPBM packets. This can be over a third party IPv4 transport network such as MPLS IP-VPN or in a Campus IP backbones.




Figure 8: Single Fabric Connect Domain realized using Fabric Extend

The following figure shows a layer 2 core network where Fabric Extend can transport SPBM packets over a layer 2 MPLS VPLS or PBB E-LINE service by creating layer 3 tunnels over a layer 2 third party network.



Figure 9: Fabric Extend over VLAN tunnels

What are the advantages of connecting Fabric Connect networks?

Fabric Connect is an Ethernet-based, industry-standard (IEEE 802.1aq) networking virtualization solution. With Fabric Connect, you can have thousands of virtualized service instances at any point in the network. Other Fabric Connect advantages include rapid time to service, Layer 2 and Layer 3 Unicast and IP Multicast virtualization, and scalable IP multicast. But the most significant advantage of Fabric Connect is that you provision services at the network edge **only**, not the core.

The **Fabric Extend** feature enables you to *extend* the Fabric Connect model. This allows Enterprises to extend Fabric Connect technology over Layer 2 and Layer 3 core networks. The interconnection of Fabric Connect deployments can be over any IP-based network whether it's a campus backbone, Data Center, or a MAN/WAN IP MPLS network.

What devices support Fabric Extend?

The VSP 7200 and the VSP 8000 Series support Fabric Extend natively. You can use these switches in a main office of a hub and spoke deployment or to connect one Data Center to another Data Center.

The VSP 4000 also supports Fabric Extend, but the switch must be connected to an Open Networking Adapter (ONA) because the VSP 4000 does not support Fabric Extend natively. The ONA enables the VSP 4000 to support Fabric Extend. The VSP 4000 uses the ONA to encapsulate Fabric Connect traffic. For example, you can use the VSP 4000 in a branch office of a hub and spoke deployment.

*** Note:**

In a Layer 2 core Fabric Extend solution, the VSP 4000 does not require an ONA because the tunnels are point-to-point VLAN connections, not VXLAN. Therefore, there is no need for an ONA to encapsulate a VXLAN header to SPB packets.

For more information on which switches support ONAs, see the feature overview in *Release Notes*.

Fabric Extend licensing

The Fabric Extend solution requires a Premier License. If a Premier License is not present, you cannot configure Fabric Extend and logical IS-IS interfaces.

! Important:

The VSP 7200 and the VSP 8000 Series support PLDS licenses **only** and share the same order codes.

The VSP 4000 supports licenses generated from either the Avaya Data Licensing Portal or PLDS. This means that the VSP 4000 can have either a GENLIC Premier License or a PLDS Premier License.

For more information about licensing, see *Administering*.

Fabric Extend and ONA

The VSP 7200 and the VSP 8000 Series support Fabric Extend natively on any of its physical ports. However, the VSP 4000 requires an Avaya Open Networking Adapter (ONA) to enable this functionality. The ONA is the Fabric Extend packet encapsulation engine for the VSP 4000. The ONA / VSP 4000 combination can also provide enhanced features such as IP fragmentation and reassembly on Fabric Extend tunnels.

*** Note:**

In a Layer 2 core Fabric Extend solution, the VSP 4000 does not require an ONA because the tunnels are point-to-point VLAN connections, not VXLAN. Therefore, there is no need for an ONA to encapsulate a VXLAN header to SPB packets.

The VSP 4000 manages the ONA in the following ways:

- Controls and provisions the ONA.
- If PoE capable, the VSP 4000 supplies power to the ONA. (The ONA also supports an optional wall unit power adapter.)
- Transports traffic to and from the ONA over 1GbE ports and sets QoS appropriately to the ONA's.
 - The ONA 1101GT can support basic Avaya Fabric Extend at line rate 1G traffic from the VSP 4000 at 1500 byte packet sizes.

- Oversubscription of the ONA's packet engine may result if packets are smaller than 1500 bytes or if you enable enhanced features such as fragmentation and reassembly of packets. This results in packet drop starting with lower QoS queued packets consistent with PCP and DSCP markings on packets received from the VSP 4000. For more details on the ONA 1101GT forwarding performance, see [ONA considerations](#) on page 58.

The ONA can operate in different modes. Fabric Extend is Operational Mode 1. To enable Fabric Extend, use the ONA's Manual Configuration menu to change the Operational Mode parameter to 1. For more information, refer to the manual that ships with the ONA.

In the following figure, the VSP 8000 is in a Fabric Connect network and is configured with Fabric Extend (FE). The VSP 4000 is also in a Fabric Connect network and is configured with SPB. The VSP 8000 and the VSP 4000 use industry-standard VXLAN tunnels to create a flow for FE traffic between the VSP 8000 and the ONA attached to the VSP 4000.

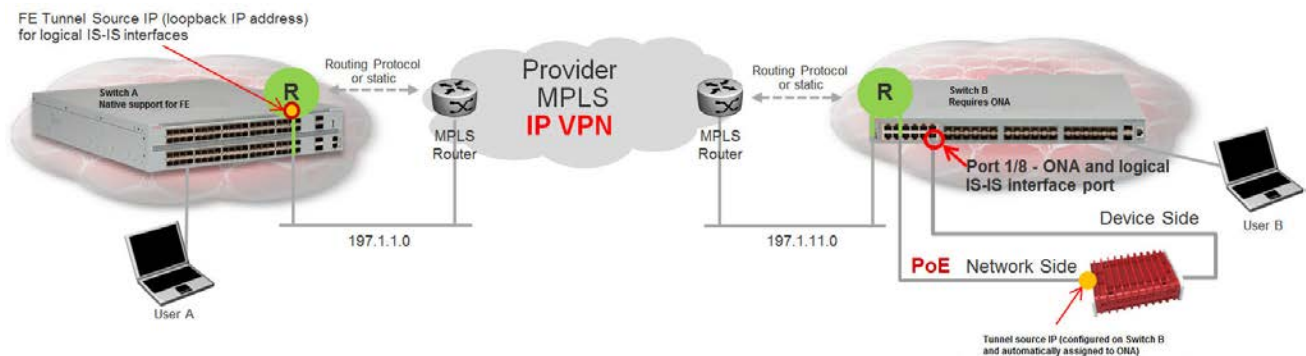


Figure 10: Fabric Extend traffic flow

The following flow occurs when User A sends a packet to User B:

- The VSP 8000 receives the packet and encapsulates it with a MAC-in-MAC header.
- The VSP 8000 sends the MAC-in-MAC-encapsulated packet over the VXLAN tunnel to the VSP 4000.
- The VSP 4000 receives the packet and sends it to the ONA network port.
- The ONA decapsulates the packet by removing the VXLAN header and sends the MAC-in-MAC packet header out the ONA device port back to the VSP 4000.
- The VSP 4000 decapsulates the MAC-in-MAC header and forwards the packet to User B.

The following flow occurs when User B sends a packet to User A:

- The VSP 4000 receives the packet and sends it to the ONA over the ONA device port with MAC-in-MAC encapsulation.
- The ONA encapsulates the packet with a VXLAN header.
- The ONA then sends the packet out the ONA network port and back to the VSP 4000.
- The VSP 4000 sends the VXLAN-encapsulated packet over the Routed IP network to the VSP 8000.

*** Note:**

To interoperate with the VSP 8000, you must set the MTU on the VSP 4000/ONA combination to 1950 bytes.

- The VSP 8000 decapsulates the packet by removing the VXLAN header and the MAC-in-MAC header, and then forwards it to User A.

*** Note:**

Connect the ONA as shown with two ports to the VSP 4000. You cannot connect the ONA directly to the IP core infrastructure.

Logical IS-IS interface

The *logical IS-IS interface* is the mechanism that enables Fabric Extend to connect SPB fabric nodes.

Logical IS-IS interfaces perform the following functions depending on the type of core network:

- In a Layer 3 core network, logical IS-IS interfaces create virtual IP tunnels and encapsulate SPB traffic by adding a Virtual Extensible LAN (VXLAN) header to SPB packets.
- In a Layer 2 core network, logical IS-IS interfaces do not use VXLAN. The tunnels are point-to-point VLAN connections so there is no need to encapsulate a VXLAN header to SPB packets. The logical IS-IS interfaces translate the Backbone VLAN IDs (B-VIDs) and maps them to each of the branch provider VIDs.

Fabric Extend uses virtual tunnels in Layer 3 core solutions to connect SPB fabric nodes. These nodes can stretch over IP routed campus networks, service provider Layer 2 core networks, or service provider Layer 3 core networks such as IP MPLS VPNs.

*** Note:**

VLACP cannot be used on logical IS-IS interface connections.

Layer 2 core network

If the service provider has a Layer 2 core network, note the following points:

- The syntax for configuring a logical interface is:

```
logical-intf isis <id> vid <list of vlans> primary-vid <vlanId> port  
<slot/port> Mlt <mltId> [name <name>]
```

- vid <list of vlans> should have two VLANs, not more than two or less than two. The VID range is <2-4059>. You do not have to configure the VIDs as platform VLANs.
- primary-vid should be included in vid <list of vlans>.
- Each logical interface must have a unique set of VIDs for each port or MLT. The same VIDs however, can be reused across a different set of ports or MLTs.
- Logical interface VIDs and BVLANS cannot be the same.
- Configuring the same VIDs as primary and secondary is not allowed.
- The port/MLT on which the Layer 2 core IS-IS logical interface is configured cannot be part of any other user configured VLANs.

- Cannot delete an MLT that is configured as a logical interface tunnel MLT.
- A logical interface consists of a port/MLT and a list of VLANs, where port/MLT is the physical connectivity to the Layer 2 core network and VLANs are the list of VLANs used to transport/bridge IS-IS control packets and Mac-in-Mac data traffic.
- VXLAN headers are not used in Layer 2 core Fabric Extend solutions.
- IS-IS control packets are not encapsulated before they are sent over a logical interface. Instead, the VLAN in the outer Ethernet header (SPB primary bvid) is replaced by the user configured logical interface VLAN.
- Spanning tree is disabled by default on **port/MLT** on which a Layer 2 core logical IS-IS interface is configured.

Layer 3 core network

If the service provider has a Layer 3 core network, note the following points:

- The syntax for configuring a logical interface is:

```
logical-intf isis <id> dest-ip <destIpAddr> [name <name>]
```
- A logical IS-IS interface points to a remote BEB destination IP address.
- Port and VlanId are not needed to create a logical IS-IS interface, instead they can be retrieved from the next hop of destination IP address.
- IS-IS control packets (IS-IS hello, LSDB, CSNP, PSNP) are encapsulated with a VXLAN header and sent over a logical IS-IS interface.

Types of Fabric Extend deployments

As the number of Fabric Connect networks increased, the need to connect those networks became more and more desirable. Fabric Extend solves the problem of going beyond the Ethernet Fabric Connect connections to include the following IP routed wide area network (WAN) and campus solutions:

1. SPB Fabric over an MPLS IP-VPN provider WAN
2. SPB Fabric over an MPLS Virtual Private LAN Service (VPLS) or Provider Backbone Bridging (PBB) Ethernet LAN (ELAN) provider network
3. SPB Fabric over an IP campus network
4. SPB Fabric over an MPLS Pseudo-Wire or Ethernet Virtual Private Line (E-Line) provider network

SPB Fabric over an MPLS IP-VPN provider WAN

The most common Fabric Extend deployment is a hub and spoke topology that connects the Main office over a service provider's MPLS IP VPN to multiple Branch offices. The following figure illustrates how the hub device on the main site establishes virtual tunnels with all of the spoke devices in the same domain. In this scenario, the traffic flows are bidirectional: from hub-to-spoke and spoke-to-hub.

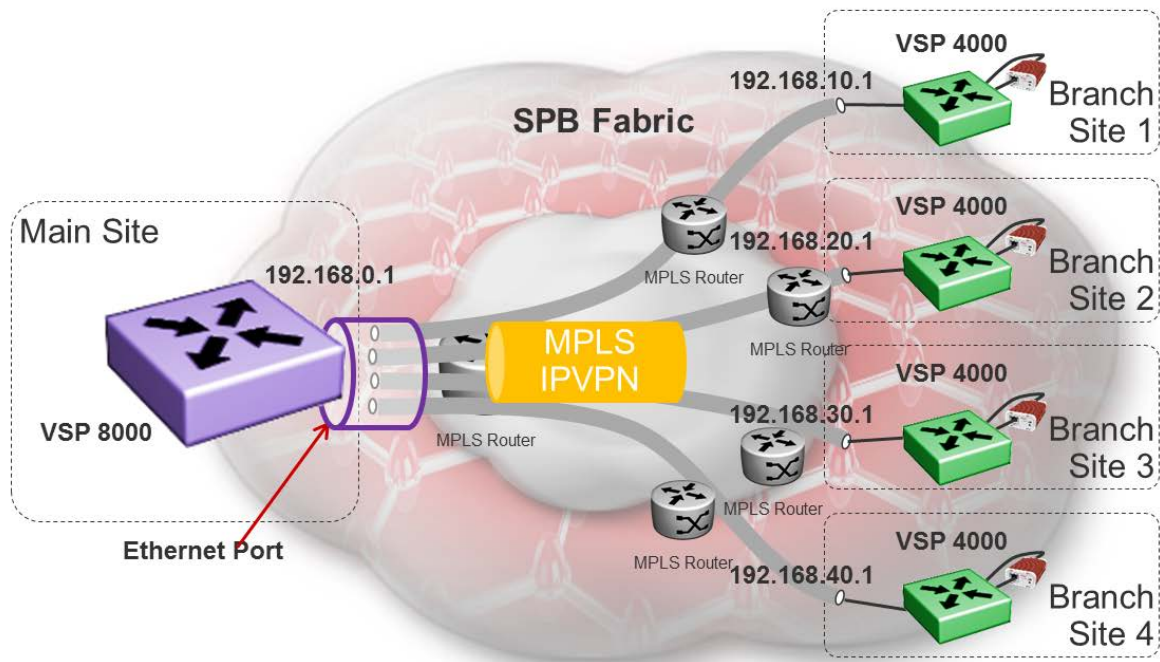


Figure 11: Fabric Extend IP VPN Deployment Option

*** Note:**

If fragmentation and reassembly is required, you must have a VSP 4000/ONA combination on the main site as well.

SPB Fabric over an MPLS VPLS/P2P-VPLS/E-LINE/P2P-VLAN provider network

Where the above hub and spoke deployment is over a Layer 3 MPLS IP-VPN, the following VPLS deployment is over a Layer 2 segment. This type of hub and spoke deployment extends the fabric over an MPLS Virtual Private LAN Service (VPLS) or Provider Backbone Bridging (PBB) Ethernet LAN (E-LINE) network. In this scenario, the SPB nodes are connected with a point-to-point Ethernet link.

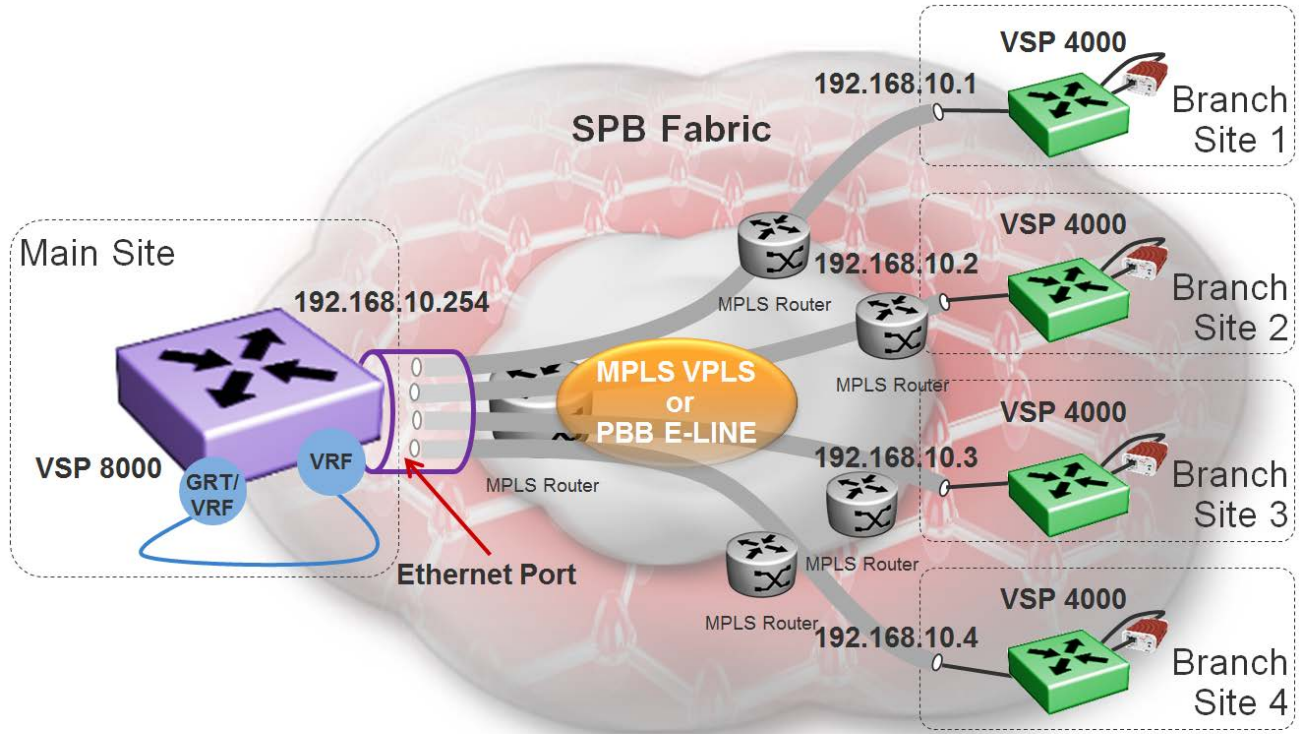


Figure 12: Fabric Extend VPLS Deployment Option

SPB Fabric over an IP campus network

Some customers do not want to migrate their infrastructures to SPB immediately. They want to keep their existing IP core network and deploy SPB on the edge. In this scenario, Fabric Extend supports a fabric overlay on top of the existing campus infrastructure.

The following figure illustrates how this deployment supports any-to-any traffic with full-mesh tunnels between fabric nodes. The fabric nodes serve as campus switches, support routing into the IP infrastructure, and provide an overlay fabric that enables all fabric benefits.




Figure 13: Fabric Extend Full Mesh Campus Deployment Option

SPB Fabric over an MPLS PWE3/E-Line provider network

The following hub and spoke deployment over an MPLS Pseudowire or Ethernet Virtual Private Line (E-Line) uses service provider VLAN tunnels. Because you can map many (VID, port/mlt list) sets to an I-SID, this gives Service Providers the flexibility to let more than one customer use the same VLAN with different I-SIDs.

*** Note:**

The VSP 4000s in this type of deployment do not require an ONA because the tunnels are point-to-point VLAN connections, not VXLAN. Therefore, there is no need for an ONA to encapsulate a VXLAN header to SPB packets.

The following figure illustrates how two dedicated Backbone VLAN IDs (B-VIDs) are mapped from the hub to spoke sites. Logical IS-IS interfaces translate the B-VIDs and maps them to each of the branch provider VIDs.

For a detailed configuration example showing logical interfaces using B-VID translation to two different logical VLAN IDs, see *Shortest Path Bridging (802.1aq) Technical Configuration Guide*, NN48500–617.

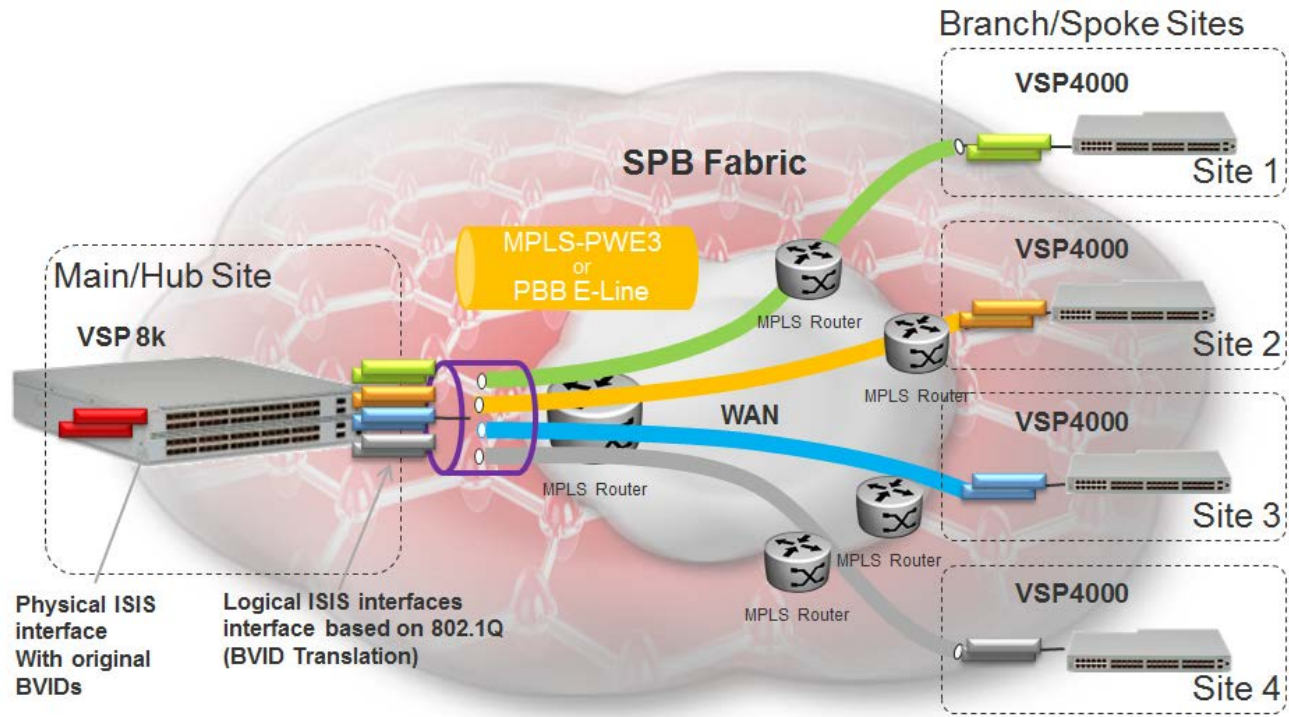


Figure 14: Fabric Extend Pseudowire Deployment Option

Fabric Extend view in AFO

The Fabric Extend view within Avaya Fabric Orchestrator (AFO) provides a graphical management interface for administrators to configure and monitor fabric extensions.

* Note:

Although the Fabric Extend view is not required, Avaya strongly recommends that you include this element in your Fabric Extend solution.

Every Fabric Extend network deployment involves creating numerous bidirectional tunnels. The Fabric Extend view in AFO automates the provisioning of these tunnels by creating *Fabric Extend domains*. When you add VSP nodes to a Fabric Extend domain, the Fabric Extend view automatically creates tunnels between the nodes belonging to the same domain. The Fabric Extend view also ensures error-free bidirectional tunnel provisioning and decommissioning, if required.

Fabric Extend view functions

Fabric Extend view provides the following functions:

- Graphically represents the Fabric Connect “islands”
- Identifies Fabric Extend capable switches
- Graphically represents virtual Fabric Extend links and status
- Provides an easy way to group a set of switches into a Fabric Extend domain
- Provides an easy way to configure point-to-point fabric extensions

Fabric Extend domains

There are two types of Fabric Extend domains:

- **Mesh** – This type of domain creates full-mesh tunnels between all nodes. If you add a switch to a mesh domain, the Fabric Extend view automatically builds Fabric Extend tunnels to all the other switches in the domain.
- **Hub-and-Spoke** – This type of domain identifies each node as either a hub or a spoke.
 - Hub nodes automatically establish bidirectional tunnels with all spoke nodes in the domain.
 - Spoke nodes automatically establish bidirectional tunnels only with the hub nodes in the domain.

Point-to-Point tunnels

You can use Fabric Extend view to provision your own tunnels between Fabric Extend-capable nodes. You must specify the tunnel configuration for both ends of the tunnels.

Fabric Extend considerations

Review the following restrictions, limitations, and behavioral characteristics that are associated with Fabric Extend.

* Note:

If your Fabric Extend configuration includes a VSP 4000/ONA combination, see [ONA considerations](#) on page 58 for more information.

- **Premier License**—Fabric Extend requires a Premier License.
- **Avaya Fabric Orchestrator (AFO)**—The Fabric Extend view within Avaya Fabric Orchestrator (AFO) is not required, but it is highly recommended.
- **Tunnel failover time**—With IS-IS interface default values, tunnel failure detection can take up to 27 seconds. You can reduce the IS-IS interface hello timers to speed up logical link failure detection, but be careful to avoid link flapping due to values that are too low.

* Note:

If the number of IS-IS interfaces on a node is greater than 100, it is a good practice to set the hello timer not lower than 5 seconds.

- **ACL Filters over VXLAN**—IP filters configured to match IP header fields in the headers of VXLAN encapsulated packets, work only when the switch acts as a transit router and does not participate in the initiation or termination of VXLAN traffic.
- **VLACP**—VLACP is not supported over logical IS-IS interfaces.
- **CFM CCM**—CFM Continuity Check Messages are not supported over logical IS-IS interfaces.
- **CFM traceroute and tracemroute**—If CFM packets transit over a layer 3 tunnel (that is the CFM packets ingress a Fabric Extend layer 3 core tunnel and egress through another layer 3 core tunnel), the transit SPBM nodes do not display as intermediate hops in the output for CFM `12 traceroute` and `12 tracemroute`.

This is because the CFM packets are encapsulated in the outer layer 3 header as part of VXLAN encapsulation, and the transit SPBM nodes cannot look into the payload of the VXLAN packet and send a copy of the CFM packet to local CPU for processing.

- **CFM L2 ping**—CFM L2 ping to MCoSPB source mac is not supported and may fail if they are reachable via Fabric Extend tunnel.
- **MACsec**—Switch-based MAC Security (MACsec) encryption is Layer 2 so it cannot be used with Fabric Extend IP, which is Layer 3.
- **MTU minimum in Layer 2 Pseudowire core networks**—Service provider Layer 2 connections must be at least 1544 bytes. In this type of deployment the tunnels are point-to-point VLAN connections that do not require VXLAN encapsulation. The default MTU value is 1950.
- **Logical IS-IS interfaces**—Layer 2 core and Layer 3 core logical IS-IS interfaces are not supported on the same switch at the same time.
- **Fragmentation/reassembly**—There is no fragmentation/reassembly support in Layer 2 core solutions.

If a tunnel was initially UP between a VSP 4000 and a VSP 7200 or VSP 8000 with MTU 1950 and then the VSP 4000 was later configured for fragmentation, the following behavior occurs:

- If the ONA MTU is less than 1594, the tunnel to the VSP 7200 (or VSP 8000) will go DOWN.
- If the ONA MTU is 1594 and above, the tunnel will stay UP, but any fragmented packets received from the VSP 4000 will be lost at the VSP 7200 (or VSP 8000) site.

Fragmented traffic can only be sent with a VSP 4000/ONA combination on both ends with the same MTU configured on each end.

- **RFC4963 and RFC4459 considerations:**

The ONA 1101GT provides for the IP MTU of the Network port to be reduced from the default setting of 1950 bytes to 1500 bytes or lower. The MTU reduction feature with Fabric Extend is provided to facilitate the connection of two Fabric Connect networks over an IP network with any MTU without requiring end stations on the networks to reduce their MTU. The ONA 1101GT with the IP MTU of the network port set to 1500 bytes will fragment Fabric Extend VXLAN tunnel packets exceeding 1500 bytes. The ONA 1101GT will also reassemble fragmented Fabric Extend VXLAN tunnel packets at the tunnel termination point. The IP fragmentation and reassembly RFC 791 describes the procedure for IP fragmentation, and transmission and reassembly of datagrams and RFC4963 and RFC4459 detail limitations and network design considerations when using fragmentation to avoid out of order packets and performance degradation.

Factors that can impact performance are —

- The link speed per VXLAN IP address should be slower than 1G to avoid reassembly context exhaustion.
- ECMP and link aggregation algorithms in the IP core should be configured not to use UDP port hashing that could send IP fragments after the first fragment on different paths causing out of order packets. This is due to the fact that subsequent fragments do not have UDP port information.

 **Important:**

Different MTU sizes on each end can result in traffic drops.

- **Layer 2 logical IS-IS interfaces**—Layer 2 logical IS-IS interfaces are created using VLANs. Different Layer 2 network Service Providers can share the same VLAN as long as they use different ports or MLT IDs.
- **MTU minimum in Layer 3 core networks**—Service provider IP connections must be at least 1594 bytes to establish IS-IS adjacency over FE tunnels. The 1594 bytes includes the actual maximum frame size with MAC-in-MAC and VXLAN headers. If this required MTU size is not available, a log message reports that the IS-IS adjacency was not established. MTU cannot be auto-discovered over an IP tunnel so the tunnel MTU will not be automatically set. The default MTU value is 1950.

If the maximum MTU size has to be fewer than 1594 bytes, then you require fragmentation and reassembly of packets. The VSP 4000/ONA combination supports fragmentation and reassembly, but you must have VSP 4000s with ONAs at BOTH ends of the IP WAN connection.

- **IP Shortcuts**—The tunnel destination IP cannot be reachable through an IP Shortcuts route.

 **Important:**

If you enable IP Shortcuts and you are using the GRT as the tunnel source VRF, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

If you enable IP Shortcuts and you are using a VRF as the tunnel source VRF, this is not an issue.

- You cannot establish a Virtual IST (vIST) session over a logical IS-IS interface. IST hellos cannot be processed or sent over a logical IS-IS interface if that is the only interface to reach BEBs in vIST pairs.

Assume that vIST is established over a regular NNI interface and the NNI interface goes down. If the vIST pairs are reachable through a logical IS-IS interface, then the vIST session goes down in up to 240 seconds (based on the IST hold down timer). During this time, the error message `IST packets cannot be sent over Fabric Extend tunnels, vist session may go down` is logged.

 **Caution:**

Expect traffic loss when the vIST session is down or when the error message is being logged.

ONA considerations

 **Note:**

Review the following restrictions, limitations, and behavioral characteristics that are associated with the ONA.

For installation information, see the *Avaya Open Networking Adapter 1101GT Installation Job Aid* (NN48800-300).

For additional information such as software upgrade procedures, see the *Avaya Open Networking Adapter 1101GT Release Notes* (NN48800-400).

ONA Network port requirements

The following are **Network** port mandatory requirements for configuring Fabric Extend on the VSP 4000:

- The ONA Network port should not be part of any static/LACP MLT configurations.
- The ONA Network port should be part of a VLAN that belongs to the GRT.
- The ONA Network port that is configured on the switch cannot be tagged. It must be an Access port.

ONA Device port requirements

The following are **Device** port mandatory requirements for configuring Fabric Extend on the VSP 4000:

- The ONA Device port should not be part of any static/LACP MLT, VLAN, or brouter configurations.
- The ONA Device port should not be configured as an access port. It is automatically configured as a trunk port when the `ip-tunnel-source-address` command is configured.
- The ONA Device port has to be connected directly to the VSP 4000 node where the FE tunnels originate.

Layer 3 and Layer 2 ONA requirements

An ONA is required for Fabric Extend Layer 3 core solutions. An ONA is *not* required in Layer 2 core solutions because the tunnels are point-to-point VLAN connections, not VXLAN. Therefore, there is no need for an ONA to encapsulate a VXLAN header to SPB packets.

DHCP server

ONAs require access to a local DHCP server to automatically configure IP addresses. Configure an untagged ONA management VLAN to where the ONA is connected with its network side interface. If DHCP is used, a DHCP relay configuration needs to be added to the ONA network side port in order for the ONA to get an IP address assigned from a DHCP server. Alternatively, you can manually configure its IP address and other required settings with the ONA Manual Configuration menu.

IP tunnel source address

Before the ONA can get an IP tunnel source address from the VSP 4000, the following steps must be taken:

- Connect the Device and Network ports on the ONA to the VSP 4000.
- Make sure that the ONA is connected to a DHCP server. If a DHCP server is unavailable, statically configure an IP tunnel source address on the ONA.
- Create a Management VLAN on the ONA that includes the Network port.
- Designate the Device port for the IP tunnel source address in the configuration file.

The syntax for the IP tunnel source address is: `ip-tunnel-source-address <A.B.C.D> port <slot/port> [mtu <750-1950>] [vrf WORD<1-16>]`.

Automatic routing of VXLAN packets on the VSP 4000

If you configure an IP tunnel source address in a VRF instead of a GRT, then the VSP 4000 automatically routes VXLAN packets from the ONA network port into the VRF configured as part of the IP tunnel source. Although the ONA network port is a part of the management VLAN that is in the GRT, for VXLAN encapsulated packets, the VSP 4000 automatically routes the packets into the VRF in which the tunnel source IP address is configured. This is done using a filter rule that the VSP 4000 software automatically sets up that filters based on whether the incoming port is equal to the ONA network port and the packet has a VXLAN header.

The Management VLAN on the VSP 4000 that is used to communicate with the ONA must always be in a GRT and must not be a part of the IP tunnel source VRF.

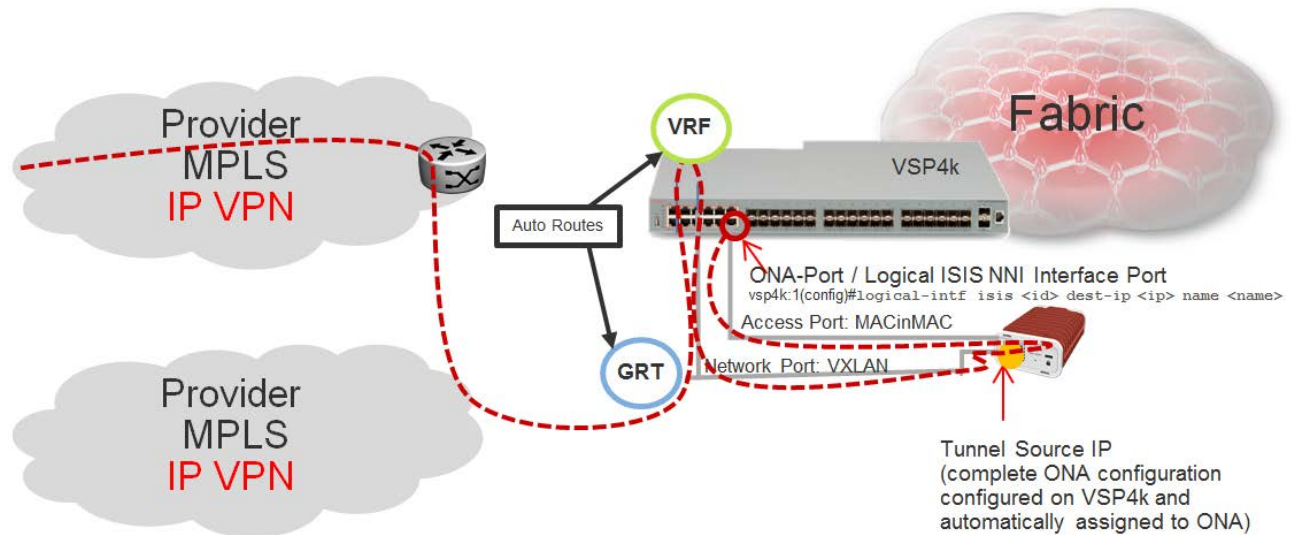


Figure 15: Autorouting between GRT and VRF

ONA Gateway

The ONA gateway has to be a local IP address on the ONA Management VLAN. The ONA gateway IP address must be the same as the local IP address of the VSP 4000 connected to the ONA.

* Note:

Avaya does not support ONA gateway IP addresses that are not local to the VSP 4000. For example, you cannot use a VRRP IP address configured in a switch cluster for the ONA gateway.

Maximum MTU

The ONA supports a maximum transmission unit (MTU) size of 1950 bytes. For the VSP 4000 to work with a VSP 8000 or VSP 7200, the MTU size must be left at the default setting of 1950. If the core network does not support jumbo frames, the VSP 4000 with ONA must be used on all sites.

Fragmentation and reassembly

If the maximum MTU size has to be fewer than 1594 bytes, then you require fragmentation and reassembly of packets. The VSP 4000s with ONAs support fragmentation and reassembly, but you must have VSP 4000s with ONAs at BOTH ends of the IP WAN connection.

QoS priority queues

The ONA 1101GT implements both Layer 2 and Layer 3 QoS. Specifically, it implements IEEE 802.1Q VLAN TCI PCP (Priority Code Point) and IETF IPv4 DSCP (Differentiated Services Code Point). These are implemented in hardware with the limitation that there are four Weighted Random Early Detection (WRED) priority queues, numbered 4 (highest) to 7 (lowest). The following tables show the mappings from the PCP and DSCP values in the packet to the priority queue.

The hardware puts each packet in 1 of the 4 HW queues in the following order:

1. If a packet is a tagged VLAN packet, the PCP field determines the priority queue. (Ethertypes 0x8100 and 0x88a8 identify tagged VLAN packets.)
2. If the packet is an IPv4 packet, the DSCP field determines the priority queue.
3. Use the highest priority queue (4).

The HW QoS is always enabled, and the CP to priority queue mappings are static.

The following table defines the 3 bit VLAN PCP value to queue number mapping. The queues are numbered 4..7 with 4 being the highest priority and 7 the lowest priority.

Table 2: VLAN PCP to queue mapping

VLAN PCP	Queue Number
0	7
1	7
2	6
3	6
4	5
5	5
6	4
7	4

The following table defines the 6 bit IPv4 DSCP value to queue number mapping. The queues are numbered 4..7 with 4 being the highest priority and 7 the lowest.

Table 3: IPv4 DSCP to queue mapping

IPv4 DSCP	VLAN PCP	Queue Number
0	1	7
1	1	7
2	1	7
3	1	7
4	1	7
5	1	7
6	1	7

Table continues...

IPv4 DSCP	VLAN PCP	Queue Number
7	1	7
8	2	6
9	1	7
10	2	6
11	1	7
12	2	6
13	1	7
14	2	6
15	1	7
16	3	6
17	1	7
18	3	6
19	1	7
20	3	6
21	1	7
22	3	6
23	1	7
24	4	5
25	1	7
26	4	5
27	4	5
28	4	5
29	1	7
30	4	5
31	1	7
32	5	5
33	1	7
34	5	5
35	5	5
36	5	5
37	1	7
38	5	5
39	1	7
40	6	4
41	5	5

Table continues...

IPv4 DSCP	VLAN PCP	Queue Number
42	1	7
43	1	7
44	1	7
45	1	7
46	6	4
47	6	4
48	7	4
49	1	7
50	1	7
51	1	7
52	1	7
53	1	7
54	1	7
55	1	7
56	7	4
57	1	7
58	1	7
59	1	7
60	1	7
61	1	7
62	1	7
63	1	7

Fabric Attach

With Fabric Attach, network edge devices that do not support Shortest Path Bridging (SPB), MAC-in-MAC encapsulation (802.1ah) or service identifiers (I-SIDs) can take advantage of SPB infrastructure. To attach to an SPB network, edge devices signal an SPB-aware FA Server to automatically configure the I-SIDs. The edge devices can then utilize existing SPB features across the fabric and leverage SPB infrastructure capabilities without manual configuration. Fabric Attach uses the IEEE 802.1AB Logical Link Discovery Protocol (LLDP) to signal a desire to join the SPB network.

FA uses the client-server model. An initial handshake occurs between the FA Server and the FA Client. After the discovery phase is complete, the FA Server accepts requests (from FA Clients) to add the C-VID (VLAN ID) and I-SID elements in the SPB network, and also automatically configures the necessary C-VID and I-SID. The FA Server then responds with an acknowledgement of whether the request succeeded. FA Clients can also be aggregated into a proxy device that handles the

handshakes and requests on behalf of many clients, to the server. All of the discovery handshakes and I-SID mapping requests are then transferred using LLDP Type, Length, Value (TLV) fields.

FA leverages LLDP to discover directly connected FA peers and to exchange information associated with FA between those peers. Based on the LLDP standard, FA information is transmitted using organizational TLVs within LLDP Protocol Data Units (PDU).

FA Zero Touch Client Attachment

FA Zero Touch Client Attachment eases the configuration process on FA-capable devices by automating specific configuration tasks required for FA functionality.

Note:

Only the base functionality of Zero Touch Client Attachment is supported.

After you initially configure Zero Touch Client Attachment on the FA Server, the settings are exported to receiving FA devices, where the required configuration tasks are automatically performed.

Base Zero Touch Client Attachment operation is tightly coupled with FA operation. Although you can enable or disable Zero Touch Client Attachment separately from FA, the feature is dependant on data that is only available during exchanges between the FA Server and FA Proxies, after a primary FA Server has been selected. By default, base Zero Touch Client Attachment support is *enabled*.

Base Zero Touch Client Attachment operation, when enabled, extracts management VLAN data from the primary FA Server advertisements and uses this data to update the in-use management VLAN if applicable. An FA Client can also utilize FA-provided management VLAN data after the FA Proxy or Server is discovered.

Zero Touch is active when the following criteria are met:

- On an FA Proxy:
 - Zero Touch Client Attachment is enabled
 - Fabric Attach is enabled
 - A primary FA Server is discovered and selected
- On an FA Server:
 - Zero Touch Client Attachment is enabled
 - FA is enabled
 - FA Proxies or FA Clients are discovered

The switch supports configurable VLANs in the range of 1 to 4059. VLAN 0 is invalid. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. VLAN IDs on the switch range from 2 to 4084 but, by default, the system reserves VLAN IDs 4060 to 4094 for internal use. If you enable both the VRF scaling and the SPBM mode boot configuration flags, the system also reserves VLAN IDs 3500 to 3999.

*** Note:**

You must enable Base Zero Touch **auto-client attach** and define the target Fabric Attach client in order to initiate Zero Touch Client Attachment processing.

FA Signaling generated by an FA Proxy or Server contains management VLAN data. If the management VLAN advertised by the primary FA Server differs from the management VLAN currently configured on the FA Proxy, Zero Touch Client Attachment initiates the following:

- VLAN creation — If the FA Server-specified management VLAN does not exist on the FA Proxy, Zero Touch Client Attachment creates a port-based VLAN.
- Management VLAN update — The created port-based VLAN becomes the designated management VLAN for the FA Proxy. No operations related to the previous management VLAN, such as port membership updates or VLAN deletion, are performed.
- Port VLAN membership update (FA Proxy/Server) — If required, Zero Touch Client Attachment updates the port VLAN membership to ensure that the uplink port through which the primary FA Server is accessed is a member of the management VLAN, for network accessibility.
- Port Default VLAN (PVID) update — The port-based PVID is automatically updated based on the VLAN ID value.
- Port Default Priority update — The default 802.1p user priority for the port is updated based on the specified port priority value of the Zero Touch client (range is 0–7).
- Zero Touch Client Specification removal — All Zero Touch client-related settings are updated based on the FA client discovery. Deleting a Zero Touch client specification or disabling any related Zero Touch option does not result in the immediate removal of any previously applied settings.

*** Note:**

The FA Proxy does not update the acquired management VLAN if the primary FA Server is lost. This data is updated if the management VLAN advertised by the current primary FA Server changes or if another primary FA Server is selected and new management VLAN data is advertised by the server.

Management VLAN and port membership updates performed by Zero Touch are maintained in non-volatile storage and are restored following a system reset. You must remove or update these configuration settings if they are deemed unnecessary at a later time.

- IP Address Source Mode Update — Updates the IP address source mode of the receiving device to *DHCP-When-Needed*, to initiate DHCP-based IP address acquisition if necessary.
- Automation of the FA Client Port Mode — Automates the configuration of EAP port modes based on the type of discovered FA Clients. Applies to FA Proxy and FA Server devices. Automated configuration is applied only to FA-enabled ports.
- ZTC Installation — Initiates ZTC installation on applicable ports on the receiving device. Applies to FA Proxy and, in a limited manner, to FA Server devices. Automated configuration is applied only to FA-enabled ports.

- Auto Trusted FA Client Port Mode — Initiates automatic QoS interface class update based on the type of discovered FA clients. Applies to FA Proxy and FA Server devices. Automated configuration is applied only to FA-enabled ports.
- Auto PVID FA Client Port Mode — Initiates automatic port PVID, port management VLAN membership and port tagging mode based on the type of discovered FA device. Applies to FA Proxy and FA Server devices. Automated configuration is applied only to FA-enabled ports. This configuration is incompatible with the automatic FA Client Port Mode and ZTC Automatic attach options.

Fabric Attach licensing

The Fabric Attach solution operates with a Base License.

Important:

The VSP 7200 and the VSP 8000 Series support PLDS licenses *only*, and share the same order codes.

The VSP 4000 supports licenses generated from either the Avaya Data Licensing Portal or PLDS.

For more information about licensing, see *Administering*.

Fabric Attach components

FA components dynamically communicate with each other using FA signaling.

FA signaling

FA has defined organizational specific TLVs within the standard LLDP protocol, to exchange messages and data amongst components of an FA solution. The FA TLVs facilitate handshaking and authentication, processing of requests for the creation of services, and providing responses on whether the requests succeeded. In addition, these services are deleted when the service requests are terminated, or when the authentication criteria are no longer valid. All components that participate in FA must be able to send, receive, and interpret the FA TLVs.

FA components

FA includes the following network elements as components:

- FA Server:

An SPB-capable switch at the edge of a Fabric Connect cloud.

An FA Server receives requests from FA Clients or FA Proxies to create services with specific I-SID-to-VLAN bindings. The FA Server completes the association between conventional networks and fabric-based virtual service networks. For more details on the operation of an FA Server, see [Fabric Attach Server](#) on page 67.

- FA Proxy:

A network switch that supports the definition of I-SID-to-VLAN assignments and has the ability to advertise these assignments for possible use by an FA Server. FA Proxy switches also support the client mode for directly attached users or end devices. Typically, FA Proxies support downstream FA Client devices, while being directly connected to an upstream FA Server device.

- FA Client:

A network attached end-point device that advertises I-SID-to-VLAN binding requests for service creation, to an FA Proxy or an FA Server. FA Clients use FA signaling to automatically attach to fabric services.

What network devices support Fabric Attach?

FA component	Network devices
FA Server	Avaya Virtual Services Platform 4000, 7200 Series and 8000 Series switches.
FA Proxy	Access switches such as the Avaya ERS 4800 and ERS 5900.
FA Client	<ul style="list-style-type: none"> • WLAN 9100 Access Points • IP Phones, Hypervisors supporting FA Client functionality on an Open vSwitch, or other third party devices

Fabric Attach Server

FA Server operation

In an FA solution, the FA Server performs the role of connecting FA Clients and FA Proxies to the SPB fabric, with minimal configuration. As part of the discovery handshake between the FA Server and client or proxy devices, LLDP PDUs are exchanged. Using standard LLDP, the FA Server learns neighbors, that include the proxy and client devices. In addition, the FA Server transmits organizational-specific element-discovery TLVs that are used by the client or proxy device to recognize its attachment to the FA Server.




Figure 16: Fabric Attach Server connecting client or proxy devices to the Fabric network

After the initial discovery handshake is complete, the client or proxy device transmits I-SID-to-VLAN assignment mapping requests to the FA Server to join the SPB fabric. These requests include the C-VID (VLAN ID) and the I-SID that the client or proxy device needs to join. The FA Server then creates the requested C-VID and I-SID on its device. It then responds with a PDU (containing the FA-specific TLV) to indicate whether the request succeeded. The I-SID thus created is a ELAN I-SID with endpoints of type Switched UNI. After I-SID creation, the I-SID is also advertised to the SPB network by IS-IS.

The traffic that is sent to or received from the SPB cloud is MAC-in-MAC (MiM) encapsulated. The FA Server, being SPB-capable, decapsulates the MiM traffic. If the I-SID matches the I-SID created

on behalf of the client or proxy, the FA Server sends the traffic to that client or proxy and passes it on the C-VID that it expects.

FA Server configuration

An FA Server can be configured at two levels—global and interface.

Configuration at the global level enables or disables FA on the entire switch. However, for attachment of clients or proxy devices, you must also configure FA at the interface level. Interfaces can be ports (including channelized ports), MLTs, SMLT or LACP MLTs. Enabling FA on an interface also enables transmission of LLDP packets that contain the FA-specific TLVs.

When you disable FA on an interface, LLDP transmission automatically stops on that interface.

Caution:

Disabling FA or IS-IS triggers a flush of FA information on the switch. Disabling FA at the global level flushes *all* FA element-discovery information and mappings. Disabling at the interface level flushes element-discovery information and mappings associated with that interface.

Important:

The only provisioning mode supported on the FA Server is SPB.

FA Proxies and FA Clients

The configuration mode of FA Proxies and FA Clients is not supported. However, in an FA solution, the FA Server interacts with FA Proxies and FA Clients by accepting LLDP PDUs (containing FA TLVs) and using them to automatically create Switched UNI I-SIDs and endpoints, based on the mapping requests contained in those TLVs. For more information, see [FA TLVs](#) on page 68.

Fabric Attach operation

The following sections detail FA operation.

FA TLVs

FA leverages LLDP to discover directly connected FA peers and to exchange information associated with FA amongst those peers. FA information is transmitted using company-specific proprietary organizational Type, Length, Value (TLV) fields within LLDP Protocol Data Units (PDU). The following section describes the TLVs for FA.

FA uses two TLVs:

- FA Element TLV
- FA Assignment TLV

FA Element TLV

The FA Element TLV is used by FA elements to advertise Fabric Attach capabilities. This data forms the basis for FA element discovery and is used in the initial handshake between the FA Server and a client or proxy device.

TLV Type [127]	TLV Length [50 octets]	Avaya OUI [00-04-0D]	Subtype [11]	HMAC-SHA Digest	Element Type	State	Mgmt VLAN	Rsvd	System ID
7 bits	9 bits	3 octets	1 octet	32 octets	6 bits	6 bits	12 bits	1 octet	10 octets

Figure 17: FA Element TLV format

Table 4: FA Element TLV field descriptions

Field	Description
TLV Type	Indicates whether the discovered element is a client or a proxy device.
OUI and Subtype	The information in these fields is used in LLDP packet handling.
HMAC-SHA Digest	<p>Data integrity and source validation is supported through the use of the HMAC-SHA256 message authentication. This field supports a digest exchange between the source and destination devices. Symmetric private keys are used for digest generation. The HMAC-SHA256 generated digest size is 32 octets.</p> <p>The HMAC-SHA256 digest is computed starting with the Element Type data, that is, it starts at zero-based byte 38 of the TLV. The digest is then placed in the HMAC-SHA256 Digest field in the TLV prior to transmission. Upon receipt, the digest is again computed and the resulting digest is compared against the received digest. If the received digest is the same as the newly computed digest, the TLV is considered valid and processing commences. If the comparison fails, the TLV is discarded and processing is terminated.</p> <p>⚠ Caution:</p> <p>If FA communication occurs between non-secure systems, the HMAC-SHA256 Digest data must always be zero. If one system operates in secure mode and the other operates in non-secure mode, the FA Element TLV is discarded before it is processed by the system operating in secure mode.</p>
Element Type	<p>Indicates the supported element type. The primary element types are the FA Server, FA Proxy and FA Client.</p> <p>An FA Server is an SPB capable device that accepts externally generated I-SID-to-VLAN assignments. An FA Proxy is a non-SPBM device that supports I-SID-to-VLAN assignment definitions and advertises these assignments for possible use by an FA Server. An FA Client, also a non-SPBM device, advertises I-SID-to-VLAN assignments to a directly connected FA Proxy or an FA Server. Both tagged and untagged FA Client connections are supported.</p> <p>The list of supported element types and their values are:</p> <ul style="list-style-type: none"> • FA Element Type - Other (1) • FA Server (2) • FA Proxy (3) • FA Server No Authentication (4) • FA Proxy No Authentication (5) • FA Client - Wireless Access Point Type 1, which directly attaches to the SPBM network.

Table continues...

Field	Description
	<ul style="list-style-type: none"> FA Client - Wireless Access Point Type 2, which is tunneled to a controller. FA Client - Switch (8) FA Client - Router (9) FA Client - IP Phone (10) FA Client - IP Camera (11) FA Client - IP Video (12) FA Client - Security Device (13) FA Client – Virtual Switch (14) FA Client – Server/Endpoint (15)
State	<p>Indicates the link tagging requirements in FA Client-sourced frames. This field also indicates the current provisioning mode.</p> <p>The Link VLAN Tagging bit (bit 1) has one of the following values:</p> <ul style="list-style-type: none"> 0 — indicates that all traffic on the link is tagged. In this case, all discovered FA Clients are treated as tagged. 1 — indicates that traffic on the link is either tagged or untagged. Here, all discovered FA Clients are treated as untagged. <p>The automatic provisioning mode bits (bits 2 and 3) always have the value 1 for SPB provisioning. The switch only supports the SPB provisioning mode.</p>
Mgmt VLAN	When you configure a management VLAN on the FA Server, it is included in this field in FA Server or FA Proxy sourced frames, and is used to support management VLAN auto-configuration on the downstream proxy and client devices.
System ID	<p>This field contains connection information that a TLV recipient can use to enforce connectivity restrictions.</p> <p>It contains the system MAC address (6 octets) for MLT configurations and the virtual BMAC address for vIST and SMLT configurations. It also contains information on the connection type such as MLT or SMLT.</p>

Limitations

- The FA Element TLV exists only once in an LLDP PDU and is included in all PDUs when the FA service is enabled.
- The maximum length of the FA Element TLV is 56 bytes.

FA I-SID-to-VLAN Assignment TLV



The FA I-SID-to-VLAN Assignment TLV is used by FA Clients to distribute I-SID-to-VLAN assignments that need to be supported by an FA Proxy or an FA Server.

TLV Type [127]	TLV Length [41-506 octets]	Avaya OUI [00-04-0D]	Subtype [12]	HMAC-SHA Digest	Assignment Status	VLAN	I-SID
7 bits	9 bits	3 octets	1 octet	32 octets	4 bits	12 bits	3 octets

Figure 18: FA Assignment TLV format

FA I-SID-to-VLAN Assignment TLV fields

Some fields are common to both the FA Element and FA Assignment TLVs. The following fields are specific only to the FA Assignment TLV.

TLV Field	Description
HMAC-SHA Digest	<p>The HMAC-SHA256 digest is computed for the series 1 to 94 of I-SID-to-VLAN assignments, that is, the data for the digest computation starts at zero-based byte 38 of the TLV. The digest is then placed in the HMAC-SHA256 Digest field in the TLV prior to transmission. Upon receipt, the digest is again computed for the series 1 to 94 of I-SID-to-VLAN assignments in the received TLV and the resulting digest is compared against the received digest. If the received digest is the same as the newly computed digest, the TLV is considered valid and processing can commence. If the comparison fails, the TLV is discarded and processing is terminated.</p> <p> Caution:</p> <p>If FA communication occurs between non-secure systems, the HMAC-SHA256 Digest data must always be zero. If one system operates in secure mode and the other operates in non-secure mode, the FA I-SID-to-VLAN Assignment TLV is discarded before it is processed by the system operating in secure mode.</p>
Assignment status	Indicates whether the FA Server accepted or rejected the I-SID-to-VLAN mapping request from a client or proxy device.
VLAN	Indicates the C-VID value advertised by the client or proxy device in the FA I-SID-to-VLAN mapping request.
I-SID	<p>Indicates the I-SID that is advertised by a client or proxy device in the FA I-SID-to-VLAN mapping request. This I-SID is used to create a Switched UNI (ELAN) I-SID.</p> <p> Note:</p> <p>This I-SID <i>cannot</i> be used by IPVPN, MVPN, SPBM dynamic multicast range, or Transparent Port UNI.</p>

Limitations

- The FA I-SID-to-VLAN Assignment TLV is included in an LLDP PDU only if the FA Server and proxy or client devices are directly connected to each other.
- This TLV can exist only once in an LLDP PDU.
- The size limit of this TLV is 511 bytes. This limits the maximum number of I-SID-to-VLAN assignments supported in an LLDP PDU to 94.
- For an FA I-SID-to-VLAN Assignment TLV to be processed, the FA Element TLV must also be present in the LLDP PDU.

FA Element Discovery

The first stage of establishing FA connectivity is element discovery.

On an FA Server, FA is enabled globally by default. However, you must explicitly enable FA on a desired port or MLT interface. After FA is enabled, the FA Server begins transmitting LLDP PDUs that contain the element discovery TLVs. This information is received by FA Client and FA Proxy devices which in turn also transmit their FA capabilities and settings. After the element handshake completes, the FA Server receives I-SID-to-VLAN assignment mappings from the connected client or proxy devices.

An FA Server can communicate with multiple different FA Client and FA Proxy devices.

FA data processing

In the following FA deployment, a client device (Client 1) attaches to the FA Server (FA Server 1) using a proxy device. Another client device (Client 2) attaches to the FA Server (FA Server 2) at the other edge of the network. The following section describes how data is processed when data traffic is transmitted from Client 1 to Client 2.

When Client 1 successfully attaches to FA Server 1, FA Server 1 creates a unique I-SID-to-VLAN mapping for Client 1 on its device. This mapping contains the I-SID and C-VID advertised by Client 1, using the FA Assignment TLV. For example, assume that Client 1 advertises I-SID 200 and C-VID 250.

Similarly, when Client 2 attaches to FA Server 2, FA Server 2 creates an I-SID-to-VLAN mapping for Client 2 on its device with, for example, I-SID 200 and C-VID 100. This is depicted in the following figure.



Figure 19: Learning of I-SID-to-VLAN mappings

When data traffic ingresses FA Server 1 at the FA-enabled port 1/1, it contains the C-VID of Client 1, which is, 250. The data is VLAN-encapsulated at this stage. As traffic egresses FA Server 1 into the SPB cloud, it is encapsulated with the ELAN I-SID created on FA Server 1 on behalf of Client 1, that is I-SID 200. The traffic is now MiM encapsulated with I-SID 200.

The following figure depicts VLAN encapsulation of data traffic from the FA Client to the FA Server (at either end of the SPB cloud) and its MiM encapsulation as it traverses the SPB cloud.

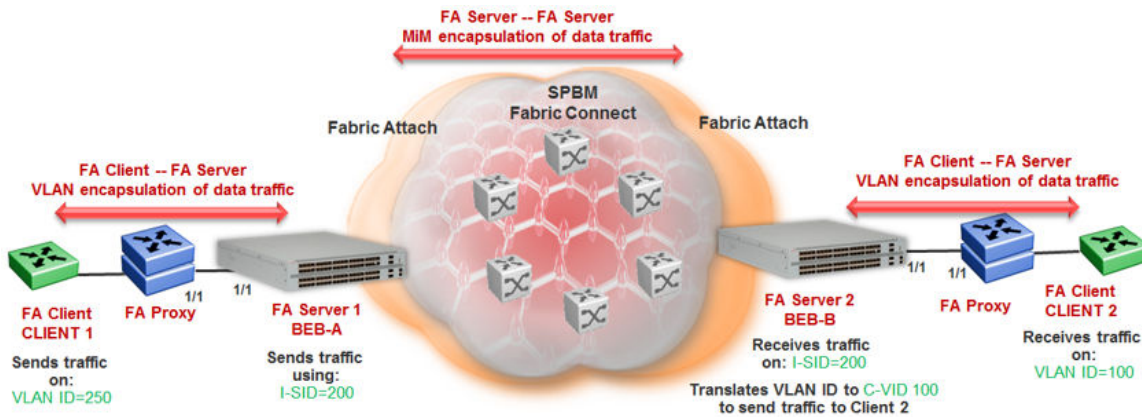


Figure 20: Data encapsulation — VLAN encapsulation and MiM encapsulation

As traffic exits the SPB cloud and ingresses the remote FA Server 2, it continues to be MiM encapsulated with I-SID 200.

At FA Server 2, the MiM traffic is decapsulated. Since the I-SID in the data packet matches the I-SID created on its device on behalf of Client 2, FA Server 2 prepares to send traffic to Client 2. At this stage, to successfully transmit the data traffic to Client 2, FA Server 2 must additionally know the C-VID that Client 2 expects traffic on. This information is obtained from the I-SID-to-VLAN mapping on FA Server 2 created on behalf of Client 2, which is C-VID 100. Thus FA Server 2 translates the C-VID in its data packets to this VLAN ID, and then passes it on to Client 2.

The following figure depicts the typical MiM encapsulation of a data packet. The B-DA and B-SA components indicated the system ID of the FA Server running SPB.

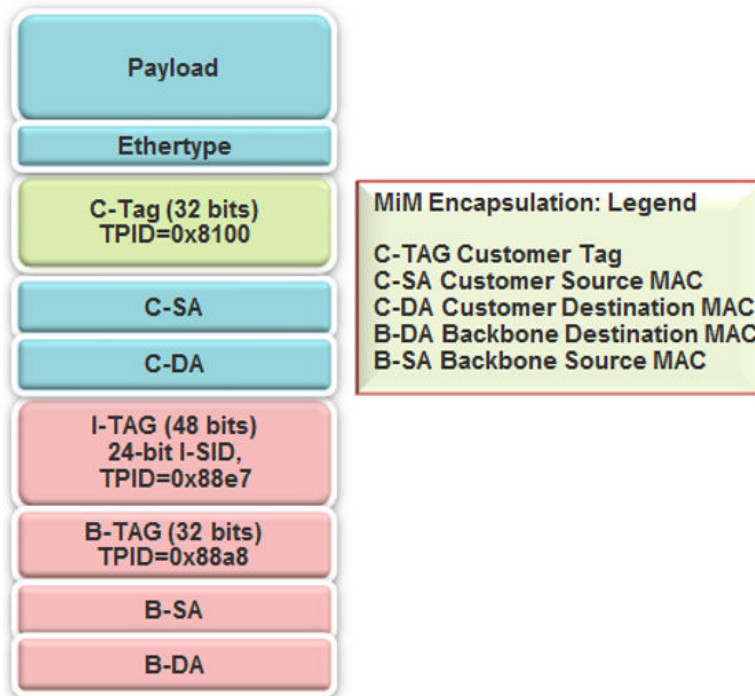


Figure 21: MiM encapsulation

FA Server and I-SID-to-VLAN assignments

FA Client or FA Proxy devices advertise I-SID-to-VLAN assignments to be supported on the FA Server. These assignments can be accepted or rejected by the FA Server. All communication between FA Proxies or Clients and the FA Server is using LLDP. Successful assignments result in the creation of a Switched UNI I-SIDs and endpoints based on the mapping requests.

The FA Server *rejects* I-SID-to-VLAN assignment requests if:

- FA is not configured properly on the port or MLT.
- Router IS-IS is disabled.

*** Note:**

For Fabric Attach to operate properly and for the FA Server to accept I-SID-to-VLAN assignment requests, IS-IS *must* be enabled.

The following error message is logged immediately after IS-IS is disabled, and appears *only once* in the log file. It does not appear again when an assignment request is made from the FA Proxy.

```
CP1 [12/04/15 00:33:49.733:UTC] 0x00374589 00000000 GlobalRouter FA
INFO Fabric Attach Assignments will be rejected since ISIS is
disabled.
```

- The C-VID and I-SID are not within the supported range.
The supported range of C-VIDs is from 1 to 4094. The value 4095 is not supported. The value 4096 indicates that the port is untagged. An I-SID value of 0 is not supported on the FA Server.
- The I-SID is already assigned to an IP VPN.
The system displays the error message `I-SID is already assigned to an IPVPN.`
- The I-SID is already in use for SPB multicast.
The system displays the error message `SPB Multicast is enabled, ISID 16000000 and greater reserved for dynamic data-isid's used to carry Multicast traffic over SPB.`
- The I-SID has a value that is reserved for internal use. This includes the range 16777213 to 16777215.
- The I-SID cannot be used in an IS-IS accept policy.
- The I-SID is associated with a platform VLAN and that VLAN is used as a private VLAN (that is, has a secondary VLAN specified).
- The I-SID is already in use for Transparent Port UNI.
- The port that receives the I-SID-to-VLAN assignment is a member of an MLT, but FA is not successfully enabled on that MLT interface.
- There is a resource error on the FA Server system, such as lack of memory.
- The number of I-SID-to-VLAN assignments on a port exceeds the maximum limit which is 94.
- The number of I-SIDs on the switch exceeds the maximum limit.
- The same endpoint is configured on more than one I-SID.
- The port or MLT is associated with more than one C-VID in the same I-SID.

When the FA Server rejects I-SID-to-VLAN assignments, aside from viewing the log file, you can use *trace* to troubleshoot the cause of rejection.

For an example on troubleshooting rejection of I-SID-to-VLAN assignments on the FA Server and for more information on using *trace*, see *Troubleshooting*.

FA management

You can configure a management I-SID on an FA-enabled port or MLT. This I-SID includes an optional C-VID parameter, which is a VLAN ID that is locally significant to the port or MLT and does **not** represent a platform VLAN.

Depending on whether the C-VID value is specified, the behavior is as follows:

- If the C-VID value is specified, the FA Server transmits this VLAN ID as the management VLAN in the FA Element TLV. A client or proxy receiving this TLV uses this VLAN ID for management traffic on the FA Server uplink.

The range for C-VID values is from 1 to 4094.

- If the C-VID value is *not* specified, the FA Server transmits a management VLAN with a VLAN ID value of 4095 in the FA Element TLV. A client or proxy receiving this TLV uses **untagged** traffic for network management on the FA Server uplink.

If you do not configure a management I-SID, the FA Server transmits a management VLAN ID value of 0 in the FA Element TLV. A client/proxy that receives the FA Element TLV retains the initial management configuration (if any) on its device.

Limitations of FA management I-SIDs

- A management I-SID value of 0 is not supported on the FA Server.
- You cannot enable BPDU on a management I-SID.

FA management configuration considerations

A Switched UNI I-SID that is created when an FA assignment is learned on a port or MLT, is uniquely identified by a tuple comprising of one of the combinations of (port, I-SID and C-VID) or (MLT ID, I-SID and C-VID). When you configure FA management, similar tuples are used. You can configure FA management on an FA-enabled port or MLT on which FA assignment mappings are learned, as long as the FA management tuple *exactly* matches the tuple created by the learned FA mapping.

The following scenarios describe the behavior when you configure FA management on a port or MLT that also receives learned FA mappings, but the tuples do not match.

- **Scenario 1:** You attempt to configure FA management on a port or MLT where an FA assignment mapping is already learned.

For example, consider an FA-enabled port 1/1 on which an assignment mapping is learned, with I-SID 100 and C-VID 20. You can configure FA management on port 1/1 as long as the I-SID and C-VID values exactly match that of the learned FA mapping. However, if you attempt to configure FA management on the port with a different I-SID and C-VID value, the configuration is not successful and an error message displays.

- **Scenario 2:** An FA assignment mapping is learned on a port or MLT that already has FA management configured.

For example, consider that FA management is configured on port 1/1. If an FA assignment mapping is learned on the port with the same I-SID and C-VID values as that of the FA management configuration, then the mapping is accepted. Otherwise the mapping is rejected.

FA message authentication and integrity protection

For the security of FA communication in terms of data integrity and authenticity, a keyed-hash message authentication code can be transmitted within every FA TLV.

It protects the I-SID-to-VLAN assignment exchanges between the FA Server and FA Proxy. The standard HMAC-SHA256 algorithm calculates the message authentication code (digest) involving a

cryptographic hash function (SHA-256) in combination with a shared secret key. The key is symmetric, that is, it is known by both the source and destination parties.

By default, on the FA Server, message authentication is enabled at the interface level and a default key is defined to provide secure communication.

You can configure a different authentication key on an interface (port or MLT) on the FA Server, to authenticate a client on that interface. The authentication key is stored in encrypted form when you save configuration on the FA Server. For an FA Client to authenticate and attach to the FA Server, the authentication key must match on both the client and the server. In general, the FA authentication key must match between two FA components exchanging FA TLVs through LLDP.

When you enable FA message authentication, the message authentication key (default or configured) generates a Hash-based Message Authentication Code (HMAC) digest that is included in FA I-SID-to-VLAN Assignment TLV. Upon receipt, the HMAC digest is recomputed for the TLV data and compared against the digest included in the TLV. If the digests are the same, the data is valid. If the digests are not the same, the data is considered invalid and is ignored.

The FA secure communication setting (enabled/disabled) and the symmetric key data are maintained across resets and restored during FA initialization.

Fabric Attach and Switched UNI

With the C-VLAN UNI feature, I-SID-to-VLAN mappings must be unique across the network. With the Transparent Port UNI (ELAN Transparent) feature, you can map an entire port or MLT to an I-SID.

With the Switched UNI feature, you can associate many different C-VID/port or C-VID/MLT list combinations to a single I-SID.

Switched UNI and FA

FA brings the capability of automatically creating Switched UNI I-SIDs on a switch, without manual intervention. The I-SIDs thus created are ELAN I-SIDs with endpoints of type Switched UNI, and are by default for Layer 2. MAC learning takes place and there is an any-to-any relationship. For Layer 3 participation, you must configure a platform VLAN with the same I-SID value as that of the I-SID in a learned FA mapping.

Note:

The number of Switched UNI I-SIDs created are different for different product families. For more information, see *Release Notes*.

Limitations of FA-created Switched UNI I-SIDs

- An FA-created Switched UNI I-SID is always ELAN.
- You cannot enable BPDU on an FA-created Switched UNI I-SID.
- The ELAN I-SIDs created are by default for Layer 2. For Layer 3 participation, you must *manually* configure a platform VLAN with the same I-SID value as that of the I-SID in a learned FA mapping. You can configure the platform VLAN with the same VLAN ID as that of the C-VID, or use a different value.
- The Switched UNI (ELAN) I-SID cannot be used by IPVPN, MVPN, SPBM dynamic multicast range, or a T-UNI.

- You cannot change from one UNI type to another dynamically. The I-SID must be deleted and created with the new UNI type (Customer VLAN (C-VLAN), Transparent Port user-network-interface (T-UNI), ELAN).
- If the port is a member of an MLT, you must add the entire MLT to the C-VID.
- The port is always in the forwarding state.
- You cannot associate a port or MLT with more than one C-VID in the same I-SID.
- The same C-VID, port or MLT cannot be a member of more than one I-SID. The supported range of C-VIDs is from 1 to 4094. The value 4095 is not supported and cannot be configured. The value 4096 indicates that the port is untagged.
- An I-SID value of 0 is not supported on the FA Server.

Fabric Attach deployment scenarios

Fabric Attach is typically deployed in the access layer(s) of a Fabric Connect network.

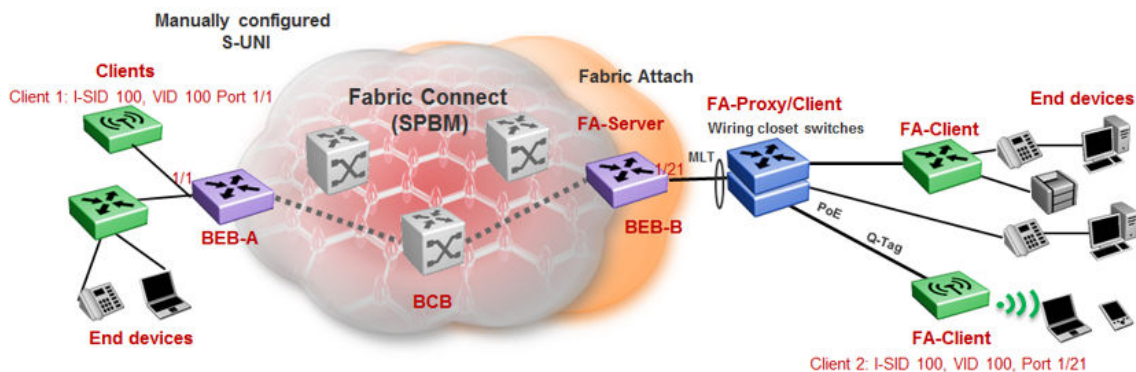
Fabric Attach, when used with a Fabric Connect solution, provides the same capabilities at the access layer, but those services and policies are now mapped across the entire network end-to-end. FA makes user and end device attachment simple and creates network configuration and sets up resources only when needed.

An FA Server can be connected to FA Client or FA Proxy devices on three types of interfaces, namely, a port, MLT or an SMLT. The following sections discuss FA in SMLT and non-SMLT deployments.

FA and Switched UNI in non-SMLT deployments

The following deployment shows an SPBM network in which one edge has *manually* configured Switched UNI I-SIDs and the other edge has Fabric Attach (FA). At the FA edge, the I-SIDs are learned using FA TLVs and are automatically created on the FA Server as ELAN I-SIDs with Switched UNI endpoints.

This deployment demonstrates that the FA-created I-SIDs can communicate with any other I-SID (manually created Switched UNI or a C-VLAN with an I-SID), on the local switch or across the SPBM fabric, as long as the I-SID values are the same.



BEB-B is a switch acting as the FA Server with an NNI interface to the SPBM cloud. FA Client and FA Proxy devices send I-SID-to-VLAN mapping requests to the FA Server on the respective FA-enabled ports, using LLDP TLVs. This enables the I-SID endpoints to communicate with the SPB cloud.

If several clients are aggregated in an MLT, at least one of the ports must send the mapping requests for the FA Server to create the I-SID endpoints for that MLT. For example, let Client 2 be a wireless FA Client (such as an Avaya WLAN 9100 AP device) on port 1/21, that sends an FA mapping request for I-SID 100 and C-VID (VLAN ID) 100. The FA Server (BEB-B) creates the requested I-SID 100 on its device, and advertises it to the SPB cloud.

BEB-A has manually configured Switched UNI endpoints, one of which is Client 1 (connected at port 1/1) using the *same* I-SID value 100.

With this setup, data traffic can freely flow between Client 1 and Client 2 through the two BEBs and the BCB.

Thus the Switched UNI I-SIDs learned using FA TLVs on one edge of the Fabric Connect (SPBM) network can communicate with the manually created I-SIDs on the other edge, as long as they both have the same value.

FA and Switched UNI in SMLT deployments

The following examples discuss FA in dual-homed and single-homed SMLT deployments.

Fabric Attach in a dual-homed SMLT deployment

The following section describes FA in a dual-homed SMLT deployment. A pair of switches that operate as IST peers act as the FA Server. An FA Proxy (typically a wiring closet switch or an access switch) is connected to FA Clients and in turn to end devices. The FA Clients or FA Proxies advertise I-SID-to-VLAN mappings namely the interface C-VID and the I-SID to the FA Server switches. Both switches receive the mapping information using LLDP TLVs. The switch that learns the mapping first from the LLDP TLV considers the I-SID endpoint to be discovered *locally*, and creates the I-SID on its device. It then sends the mapping information to its peer switch. When the peer switch receives the mapping across IST in a new SMLT message, it too creates the I-SID and endpoint on its device. This I-SID however, is considered to be discovered *remotely*, because the data was synchronized from its peer.

*** Note:**

- For the peer switches acting as the FA Server to transmit the same FA System ID (based on the virtual MAC), SMLT configuration *must* be the same on both peers.
- For successful FA operation, configuration of FA message authentication and the authentication key *must* be the same on both peers.
- For successful operation in Layer 3, a platform VLAN must be configured on both peers. This is necessary for proper MAC learning.




Figure 22: FA in a dual-homed SMLT deployment

In the above example deployment, BEB-A and BEB-B are IST peers collectively acting as the FA Server. FA TLVs sent from the clients (through the proxy) are learned on FA-enabled ports on BEB-A and BEB-B. When BEB-A learns the mapping for the first time on its port, it creates an I-SID on its device. This is considered *locally* discovered. In addition, it sends an SMLT message to its peer BEB-B, which also creates the I-SID on its device. This time, the I-SID is considered *remotely* discovered. Similarly, if BEB-B receives a mapping from a client for the first time, it creates an I-SID (locally discovered) and also sends an IST message to its peer to create an I-SID (remotely discovered).

Irrespective of whether the I-SID creation on the FA peers is triggered by a local TLV event or by messaging from the IST peer, they can both receive data traffic. Thus in a dual-homed SMLT deployment, any I-SID can be learned irrespective of whether it is discovered locally, discovered remotely or both.

*** Note:**

On the IST peers, if an FA TLV is learned on a port or normal MLT (instead of the admin SMLT), only the I-SID is sent to the peer switch.

Fabric Attach in a single-homed SMLT deployment

In the single-homed SMLT, as shown in the following deployment, the FA Server creates either a locally discovered I-SID (if received from a client using FA TLVs) or a remotely discovered I-SID (if synchronized from its IST peer), but not both.

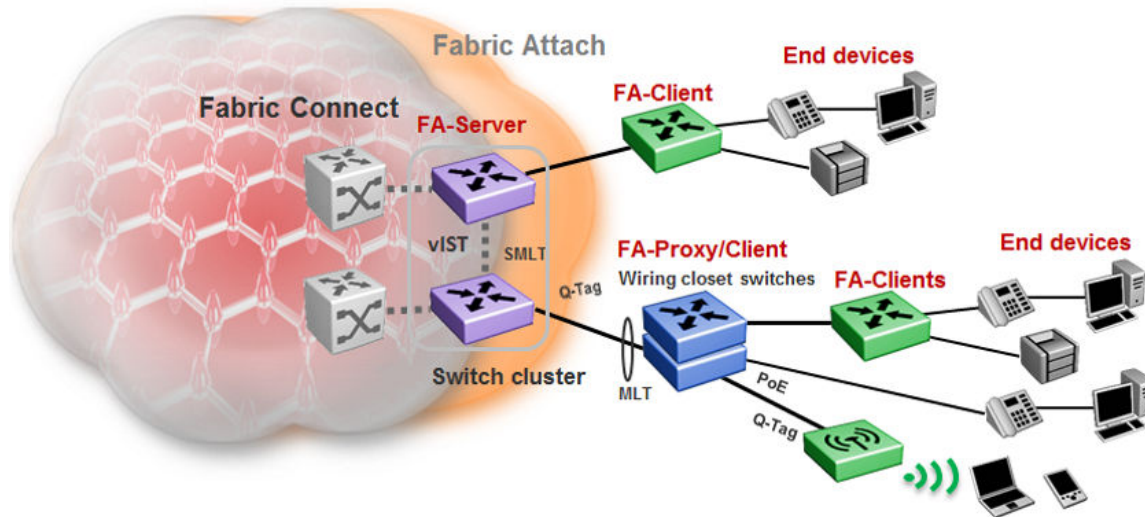


Figure 23: Fabric Attach in a single-homed SMLT deployment

Fabric Attach considerations

Review the following restrictions, limitations, and behavioral characteristics for Fabric Attach.

- IS-IS and FA must be globally enabled on the FA Server, for FA to operate successfully.
- Static MAC, Static ARP and configuration of a static IGMP group are not supported on FA-enabled ports.
- An FA port cannot be a BROUTER port.
- You cannot enable FA on an existing Flex UNI, Transparent Port UNI or a C-VLAN UNI port.

- FA I-SID-to-VLAN assignment mapping requests from a client or proxy device can be accepted or rejected by the FA Server.
- On an FA-enabled port or MLT, you must first disable LACP before you change the LACP key.
- On VLACP enabled ports, FA and LLDP signaling run independent of the VLACP state. Therefore, requests and responses are exchanged between the FA Server and client or proxy devices even if VLACP is operationally *down*. However, forwarding of data traffic is dependent on VLACP being operationally *up* on the port.

For example, if VLACP is enabled on the FA Server side of the link but not on the proxy or the client side, the FA Server learns the I-SID-to-VLAN assignment mappings and creates the required I-SIDs on its device. However, data traffic is not forwarded on the port until VLACP is operationally *up*.

- You cannot use a port designated as a Fabric Extend tunnel source (configured using the command `ip-tunnel-source-address`), for Fabric Attach. This limitation applies to the VSP 4000 Series only.
- FA uses the virtual MAC to create the FA system ID when the FA is on an SMLT. If you delete the SPBM instance, then this information is no longer available. Therefore, you must delete the FA on SMLT before deleting the SPBM instance.

IS-IS external metric

The software supports the IS-IS external metric to differentiate between internal and external routes with Accept Policies.

With this feature you can use IS-IS to:

- change the external metric-type of a route when redistributing it from another protocol to IS-IS through route redistribution using a route-map.
- change the external metric-type of a route when accepting a remote IS-IS route with the help of IS-IS accept policies using a route-map.
- match the external metric-type when redistributing IS-IS routes into other protocols using the match option in the route-map.
- match the external metric-type when accepting a remote IS-IS route with the help of IS-IS accept policies by using a route-map
- process the external metric-type in the route selection process.

The IS-IS metric type can also be set using the base redistribute command without using the route-map.

SPB Ethertype

The switch aligns the SPB ethertype to BCB's locally configured SPB ethertype. The BCBs mark the BTAG Ethertype of a transit MAC-in-MAC packet to match its locally configured value when it exits

on a different NNI port, even if the BTAG Ethertype on the incoming packet (CFM or SPB) does not match its configured value.

*** Note:**

ISIS Hello packets are always marked with 0x8100 ethertype, and do not change according to the BCB's locally configured values.

SPBM and IS-IS infrastructure configuration using CLI

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

! Important:

The `spbm-config-mode` boot flag must be enabled (default) before you can configure SPBM or IS-IS. To verify the setting, enter `show boot config flags` in Privileged EXEC mode.

Running the SPBM script

Use the following procedure to run the SPBM script to automate the minimum required SPBM and IS-IS parameters to allow Fabric Connect to operate on the switch.

Before you begin

- Enable SPBM before running the SPBM script.
- Delete existing IS-IS interfaces before running this script. See [Removing specific IS-IS and MLT interfaces](#) on page 85 for information on removing IS-IS interfaces.

About this task

You can use this procedure to quickly configure the minimum SPBM and IS-IS parameters. However, a manual procedure is available instead of using this script. The default values are given in square brackets. You may input your values at the prompt or if you wish to accept the default values, press `Enter`. This command first accepts all values and then removes existing SPBM configurations before configuring the entered values.

*** Note:**

This process causes the SPBM traffic to flap temporarily.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Run the SPBM script:

```
run spbm
```

*** Note:**

If the script causes a configuration conflict or cannot execute a command, an error message displays and the script stops.

Example

Run the SPBM script:

```
Switch:1(config)# run spbm
```

```
*****
*** This script will guide you through configuring the          ***
*** switch for optimal operation SPB.                          ***
*** -----                                                  ***
*** The values in [] are the default values, you can          ***
*** input alternative values at any of the prompts.           ***
*** If you wish to terminate or exit this script              ***
*** enter ^C <control-C> at any prompt.                       ***
*** NOTE: THE COMMAND WILL TEMPORARILY FLAP IS-IS,SPBM        ***
*****
SPB Ethertype <0x8100,0x88a8> [0x8100]:
SPB primary BVLAN 2-4059 [4051]:300
SPB secondary BVLAN 2-4059 [4052]:400
ISIS system id <xxxx.xxxx.xxxx> [a051.c6eb.7c65]:0200.0000.0100
SPB nickname <x.xx.xx> [b.7c.65]:0.02.02
SPB Manual Area <xx.xxxx.xxxx...xxxx> [49.0000]:50
ISIS System Name [Switch]:BEB1
Enable SPBM multicast (y/n) [n]:y
Enable IP shortcuts (y/n) [n]:y
Loopback interface ID <1-256> [1]:1
Loopback interface IP and subnet <a.b.c.d/x>:20.1.1.1/24
Configure SPBM SMLT? (y/n) [n]:y
Peer system id <xxxx.xxxx.xxxx>:0200.0000.0200
SMLT virtual BMAC <0x00:0x00:0x00:0x00:0x00:0x00>:02:00:00:10:00:10
ISIS MLT interface <MLT ID LIST>[:]:1
Enable CFM SPBM (y/n) [n]:y
Enter CFM SPBM MEPID <1-8191> [1]:2
Enter CFM SPBM level <0-7> [4]:4

****CONFIGURATION IN PROGRESS****
*SPBM enabled globally*
*SPBM instance 1 configured*
*SPBM BVLANS configured*
*SPBM SMLT configured*
*SPBM multicast enabled globally*
*IP shortcuts configured*
*SPBM SMLT configured*
*IS-IS enabled*
*IS-IS on port 1/5 configured*
*IS-IS on port 1/6 configured*
*IS-IS on MLT 1 configured*
*CFM SPBM configured*
****SCRIPT EXECUTION COMPLETE****
```


Removing existing SPBM configuration

Use the following procedure to remove existing SPBM configurations, disable CFM, and return the CFM MEP-ID and level configurations to default values.

Before you begin

Enable SPBM before running this script.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Run the script:

```
run spbm clean
```

* Note:

If the script causes a configuration conflict or cannot execute a command, an error message appears and the script stops.

Example

Run the script:

```
Switch:1(config)#run spbm clean
The following will delete all SPBM and interfaces and default the CFM configurations. Do
you want to continue? <y/n>[n]:y

Switch:1(config)#no router isis enable
Switch:1(config)#interface gigabitethernet 1/10
Switch:1(config-if)#no isis
Switch:1(config-if)#interface gigabitethernet 1/11
Switch:1(config-if)#no isis
Switch:1(config-if)#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#no vlan 4051
Switch:1(config)#no vlan 4052
Switch:1(config)#router isis
Switch:1(config-isis)#no spbm 1
Switch:1(config-isis)#router isis
Switch:1(config-isis)#no ip-source-address
Switch:1(config-isis)#no system-id
Switch:1(config-isis)#no manual-area 49.0000
Switch:1(config-isis)#no cfm spbm enable
Switch:1(config)#cfm spbm level 4
Switch:1(config)#cfm spbm mepid 1
Switch:1(config)#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#

**SPBM configurations have been removed**
```

Configuring the IS-IS port interfaces using SPBM script

Use the following procedure to run the SPBM script to configure the IS-IS port interfaces. As this command does not flap IS-IS or SPBM, it is particularly effective to use this command when SPBM is already configured and you require to configure additional ports or MLTs. Running the `run spbm interface` command does not alter existing IS-IS or SPBM configurations.

About this task

You can use this procedure to quickly configure the minimum SPBM and IS-IS parameters. However, a manual procedure is available instead of using this script.

* Note:

You must enable SPBM before running the SPBM script.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Run the SPBM script:

```
run spbm interface
```

* Note:

If the script causes a configuration conflict or cannot execute a command, an error message displays and the script stops.

Example

Run the SPBM script:

```
Switch:1(config)# run spbm interface
```

```
*****
*** This script will guide you through configuring the          ***
*** switch for optimal operation SPB.                          ***
*** -----                                                  ***
*** The values in [] are the default values, you can          ***
*** input alternative values at any of the prompts.           ***
*** If you wish to terminate or exit this script              ***
*** enter ^C <control-C> at any prompt.                       ***
*****
ISIS port interfaces <a/b,c/d> []:1/2,1/4,1/8
ISIS MLT interface <MLT ID LIST> []:1
*IS-IS on port 1/2 configured*
*IS-IS on port 1/4 configured*
*IS-IS on port 1/8 configured*
*IS-IS on MLT-1 configured*
```

Removing specific IS-IS and MLT interfaces

Use the following procedure to remove specific IS-IS ports and MLT interfaces when you get the error IS-IS SPBM interfaces have been configured. Please delete these interfaces.

About this task

This procedure removes existing IS-IS ports and MLT interfaces. You can choose which port and MLT interfaces need to be removed. This command does not alter the other SPBM or IS-IS configurations.

* Note:

You must enable SPBM before running the SPBM script.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Run the script:

```
run spbm interface clean
```

* Note:

If the script causes a configuration conflict or cannot execute a command, an error message displays and the script stops.

Example

Run the spbm interface clean script:

```
Switch:1(config)# run spbm interface clean
```

```
*****
*** This script will guide you through deleting the          ***
*** IS-IS SPBM interfaces.                                  ***
*** -----                                                ***
*** The values in [] are the default values.                ***
*** If you wish to terminate or exit this script           ***
*** enter ^C <control-C> at any prompt.                    ***
*****
ISIS port interfaces to be deleted <a/b,c/d>[:1/2,1/4,1/8
ISIS MLT interface <MLT ID LIST> [:1
IS-IS port 1/2 deleted
IS-IS port 1/4 deleted
IS-IS port 1/8 deleted
** 3 IS-IS port interfaces deleted **
MLT 1 deleted
** 1 IS-IS MLTs deleted **
```

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.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable SPBM globally:

```
spbm
```

3. Enter IS-IS Router Configuration mode:

```
router isis
```

4. Create the SPBM instance (only one SPBM instance is supported):

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

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

```
spbm <1-100> b-vid {<vlan-id [-vlan-id][,...]} [primary <1-4059>]
```

6. 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, you should 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.

7. Configure an IS-IS manual area (1-13 bytes in the format <xx.xxxx.xxxx...xxx>. Only one manual area is supported.):

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

8. Exit IS-IS Router Configuration mode to Global Configuration mode:

```
exit
```

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

```
vlan create <2-4059> type spbm-bvlan
```

10. Enter Interface Configuration mode, by specifying the ports or MLTs that are going to link to the SPBM network:

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

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

```
isis
```

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

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

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

```
isis enable
```

14. Enable interface.

15. Exit Interface Configuration mode:

```
exit
```

16. Enable IS-IS globally:

```
router isis enable
```

17. Display the SPBM configurations:

```
show isis spbm
```

18. Display the global IS-IS configuration:

```
show isis
```

19. Display the interface IS-IS configuration:

```
show isis interface
```

Example

```
Switch> enable
Switch# configure terminal
Switch(config)# spbm
Switch(config)# router isis
Switch(config-isis)# spbm 1
Switch(config-isis)# spbm 1 b-vid 10,20 primary 10
Switch(config-isis)# spbm 1 nick-name 1.11.16
Switch(config-isis)# manual-area c0.2000.000.00
Switch(config-isis)# exit
Switch(config)# interface GigabitEthernet 1/21
Switch(config-if)# isis
Switch(config-if)# isis spbm 1
```

SPBM and IS-IS infrastructure configuration

```
Switch(config-if)# isis enable
Switch(config-if)# exit
Switch(config)# vlan create 10 type spbm-vlan
Switch(config)# vlan create 20 type spbm-vlan
Switch(config)# router isis enable
Switch(config)# show isis spbm
```

```
Switch:1(config)#show isis spbm
```

```
=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY    NICK      LSDB      IP        IPV6      MULTICAST
INSTANCE  VLAN        VLAN      NAME      TRAP
-----
1         10,20      1000     1.11.16  disable  enable   enable   enable
=====
                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID
INSTANCE
-----
1         primary              00:00:0a:02:2e:03      0000.0a02.2e02
=====
Total Num of SPBM instances: 1
=====
```

```
Switch(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 : Switch1
ip source-address : 41.41.41.100
ipv6 source-address : 41:0:0:0:0:0:100
ip tunnel source-address : 11.11.12.11
Tunnel vrf : spboip
ONA Port   : 1/15
ip tunnel mtu : 1950
Num of Interfaces : 2
Num of Area Addresses : 1
=====
```


*** Note:**

The ONA Port : 1/15 parameter in the preceding example is applicable only to the Avaya Virtual Services Platform 4000 Series.

```
Switch(config)# show isis interface
```

```
Switch# 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
Port1/21   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>). 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. Only one SPBM instance is supported.
b-vid {<vlan-id [-vlan-id] [,...]}	Sets the IS-IS 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 <1-4059>	Sets the IS-IS instance primary data B-VLAN.

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

Variable	Value
<2-4059>	Specifies the VLAN ID. Creates an SPBM Backbone VLAN (B-VLAN). You can optionally specify a name for the SPBM B-VLAN.
type {port-mstprstp protocol-mstprstp spbm-bvlan}	Specifies the type of VLAN created. <ul style="list-style-type: none"> port-mstprstp — Create a VLAN by port. protocol-mstprstp — Create a VLAN by protocol. spbm-bvlan — Create an SPBM-BVLAN.

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. Wait up to 20 minutes for the LSPs with the original system ID to age out.

* Note:

To check the age out time, use the `show isis lsdb sysid <original-sys-id>` command on any of the other SPB nodes in the network. When there is no output from this command, proceed to the next step. The time left (in seconds) for the LSPs to age out is shown under the column LIFETIME.

6. Disable IS-IS.
7. Change the nickname to the original nickname.
8. Enable IS-IS.

Configuring I-SIDs for private VLANs

Before you begin

- A private VLAN must be created. For more information about creating private VLANs, see *Configuring VLANs, Spanning Tree, and NLB*.

About this task

There is one I-SID per private VLAN.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Assign the I-SID to the primary and secondary VLAN.

```
vlan i-sid <1-4059> <0-16777215>
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# vlan i-sid 5 75
```

Display private VLAN I-SID assignment:

```
Switch:1(config)# show vlan private-vlan
```

PRIVATE VLAN			
Primary VLAN	Primary ISID	Secondary VLAN	Secondary ISID
5	75	6	75

Variable definitions

Use the data in the following table to use the `vlan i-sid` command.

Variable	Value
<1-4059>	Primary VLAN ID. Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<0-16777215>	I-SID value, this value is same for primary and secondary VLANs.

Displaying global SPBM parameters

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

Procedure

1. Display the SPBM configuration:

```
show isis spbm
```

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

```
show vlan basic
```

Example

```
Switch# show isis spbm
=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY   NICK      LSDB      IP        IPV6      MULTICAST
INSTANCE  ID          VLAN      NAME      TRAP
-----
1         4086-4087  4086     3.03.01  disable  enable   enable   disable
=====
                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB   SMLT-VIRTUAL-BMAC   SMLT-PEER-SYSTEM-ID
INSTANCE
-----
1         primary          00:00:03:03:03:03   0000.0303.0302
=====
Total Num of SPBM instances: 1
=====
```

```
Switch# show vlan basic
=====
                        Vlan Basic
=====
VLAN      INST      ID  PROTOCOLID  SUBNETADDR      SUBNETMASK      VRFID
ID  NAME          TYPE
-----
1     Default      byPort      0  none        N/A             N/A             0
10    VLAN-10     spbm-bvlan  62  none        N/A             N/A             0
20    VLAN-20     spbm-bvlan  62  none        N/A             N/A             0
100   VLAN-100    byPort      0  none        N/A             N/A             0
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.

Table continues...

Parameter	Description
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.
IPv6	Indicates the status of SPBM IPv6 shortcuts on this SPBM instance. The default is disable.
MULTICAST	Indicates if SPBM multicast is enabled. The default is disabled.
SMLT-SPLIT-BEB	Specifies whether the switch is the primary or secondary vIST peer.
SMLT-VIRTUAL-BMAC	Specifies a virtual MAC address that can be used by both peers.
SMLT-PEER-SYSTEM-ID	Specifies the vIST peer system ID.

Displaying global IS-IS parameters

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

Procedure

1. Display IS-IS configuration information:

```
show isis
```

2. Display the IS-IS system-id:

```
show isis system-id
```

3. Display IS-IS net info:

```
show isis net
```

Example

```
Switch# 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  : Switch1
ip source-address : 41.41.41.100
```

```

        ipv6 source-address : 41:0:0:0:0:0:0:100
    ip tunnel source-address : 11.11.12.11
        Tunnel vrf : spboip
            ONA Port : 1/15
        ip tunnel mtu : 1950
        Num of Interfaces : 2
        Num of Area Addresses : 1
    
```

*** Note:**

The ONA Port : 1/15 parameter in the preceding example is applicable only to the Avaya Virtual Services Platform 4000 Series.

```

Switch# show isis system-id
=====
                        ISIS System-Id
=====
SYSTEM-ID
-----
0014.c7e1.33df

Switch# 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: I1, I2, or I1/2.
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 IS-IS overload-on-startup value in seconds. The overload-onstartup value is used as a timer to control when to send out Link State Packets (LSPs) with the overload bit cleared after IS-IS startup. The default value is 20 seconds.
Overload	Indicates if there is an overload condition.
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.

Table continues...

Parameter	Description
spf-delay	Indicates an SPF delay in milliseconds. The default value is 100 milliseconds.
Router Name	Indicates the IS-IS name of the router.
ip source-address	Indicates the IP source address used for SPBM IP shortcuts.
ipv6 source-address	Indicates the IPv6 source address used for SPBM IP shortcuts.
ip tunnel source-address	Indicates the IP tunnel source address used for SPBM Fabric Extend.
Tunnel vrf	Indicates the name of the vrf that contains the tunnel endpoints.
ONA Port	Indicates the port to which the ONA device is attached.  Note: The ONA port parameter is applicable only to the Avaya Virtual Services Platform 4000 Series.
ip tunnel mtu	Indicates the maximum size of a packet that can be transmitted through the IP tunnel.
Num of Interfaces	Indicates the number of interfaces on the router.
Num of Area Addresses	Indicates the number of area addresses on the router.
Num of Summary Address	Indicates the summary of the addresses on 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 Fabric Connect globally

Use this procedure to enable IP Multicast over Fabric Connect 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 Fabric Connect is disabled. There is no need to enable IP Multicast over Fabric Connect on the Backbone Core Bridges (BCBs).

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

*** Note:**

IP Multicast over Fabric Connect 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 Fabric Connect.

*** Note:**

You must enable IP multicast over Fabric Connect globally on all DvR enabled nodes (Controllers and Leaf nodes) in a DvR domain.

For more information on DvR, see *Configuring IPv4 Routing*.

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 IP Multicast over Fabric Connect globally:

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

*** Note:**

The switch only supports one SPBM instance.

5. **(Optional)** Disable IP Multicast over Fabric Connect globally:

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

Example

Enable IP Multicast over Fabric Connect 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 Fabric Connect globally. The default is disabled. Specifies the SPBM instance. The switch only supports one instance.

Displaying IP Multicast over Fabric Connect information

Use this procedure to display IP Multicast over Fabric Connect summary information.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the status of the global IP Multicast over Fabric Connect configuration:


```
show isis spbm multicast
```
3. Display IP Multicast over Fabric Connect 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 [vrf WORD<0-16>][vrfids WORD<0-255>]
```

Example

Display IP Multicast over Fabric Connect 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 <i>WORD</i> <0-255>	Displays the IP Multicast over Fabric Connect summary information for a specific host-name.
lspid <xxx.xxx.xxx.xx-xx>	Displays the IP Multicast over Fabric Connect summary information for the specified LSP ID that you enter in xxx.xxx.xxx.xx-xx — 8 byte format.

Use the data in the following table to use the `show ip mroute route` command.

Variable	Value
vrf <i>WORD</i> <1-32>	Specifies a VRF.
vrfids <i>WORD</i> <0-255>	Specifies the VRF ID

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 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.

Table continues...

Parameter	Description
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. <div style="background-color: #e0e0e0; padding: 5px;"> <p>* Note:</p> <p>The value you configure for fwd-cache-timeout applies only to the locally learned sender; it does not apply to SMLT synchronized sender records.</p> </div>
PROT	Indicates the multicast protocol through which the switch learned this route. The spb-access and spb-network values indicate the stream learned when IP Multicast over Fabric Connect 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

Use the following procedure to display IS-IS areas.

```
show isis manual-area
```

Example

```
Switch# 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:

- The assignment of primary and secondary roles to the vIST peers is automatic. The switch with the lower system ID (between the two vIST peers) is primary, and the switch with the higher system ID is secondary when default system-id values are being used.
- 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.
- When using the default hardware assigned system-id value, the SMLT Virtual BMAC is automatically derived by comparing the system-id values of the two vIST peers. A value of 0x01 plus the lower of the two system-id values is used as the SMLT Virtual BMAC.

When using a manually configured system-id value, the SMLT Virtual BMAC must also be manually configured.

- An I-SID must be assigned to every VLAN that is a member of a Layer 2 VSN. Also, if a Layer 2 VSN is created on one vIST Peer, it must also be created on the other vIST peer.

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 vIST peer, so that if it goes down, the local peer can take over forwarding for the failed peer:

SPBM and IS-IS infrastructure configuration

```
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

```
Switch:1> enable
```

```
Switch:1# configure terminal
```

Disable IS-IS on the switch:

```
Switch:1(config)# no router isis enable
```

Enter the IS-IS Router Configuration mode:

```
Switch:1(config)# router isis
```

```
Switch:1(config-isis)# spbm 1 smlt-peer-system-id 0018.b0bb.b3df
```

```
Switch:1(config-isis)# spbm 1 smlt-virtual-bmac 00:14:c7:e1:33:e0
```

```
Switch:1(config-isis)# router isis enable
```

```
Switch:1(config-isis)# show isis spbm
```

```
=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY    NICK      LSDB      IP        IPV6      MULTICAST
INSTANCE  VLAN
-----
1         4086-4087  4086      3.03.01  disable  enable  enable  disable
=====
                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID
INSTANCE
-----
1         primary              00:00:03:03:03:03      0000.0303.0302
-----
Total Num of SPBM instances: 1
=====
```

Variable definitions

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

Variable	Value
<code>smlt-peer-system-id <xxxx.xxxx.xxxx></code>	<p>Specifies the IS-IS SPBM peer system ID.</p> <p>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.</p>
<code>smlt-virtual-bmac <0x00:0x00:0x00:0x00:0x00:0x00></code>	<p>Specifies a virtual MAC address that can be used by both peers. SMLT virtual B-MAC is an optional configuration.</p> <p>* Note:</p> <ul style="list-style-type: none"> 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. 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:

- a. Disable IS-IS on the switch:

```
no router isis enable
```

- b. Enter IS-IS Router Configuration mode:

```
router isis
```

- c. Enable a trap when the SPBM LSDB changes:

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

- d. Enable IS-IS on the switch:

```
router isis enable
```

- e. Exit IS-IS Router Configuration mode:

```
exit
```

4. Configure the optional SPBM interface parameters. To configure these parameters, you must disable IS-IS on the interface:

- a. Specify an SPBM interface to configure:

```
interface {GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]} | mlt <mltid> }
```

- b. Disable IS-IS on the interface:

```
no isis enable
```

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

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

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

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

- e. Enable IS-IS on the switch:

```
isis enable
```

Example

```
Switch> enable
```

```
Switch# configure terminal
```

```
Switch(config)# spbm ethertype 0x8100
```

```
Switch(config-isis)# no router isis enable
```

```
Switch(config)# router isis
```

```
Switch(config-isis)# spbm 1 lsdb-trap enable
```

```
Switch(config-isis)# router isis enable
```

```
Switch(config-isis)# exit
```

```
Switch(config)# interface gigabitethernet 1/7
```

```
Switch(config-if)# no isis enable
```

```
Switch(config-if)# isis spbm 1 interface-type pt-pt
```

```
Switch(config-if)# isis spbm 1 l1-metric 500
```

```
Switch(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> lsdbs-trap enable	Configures whether to enable or disable a trap when the SPBM LSDB changes. 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 {narrow|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

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

```
Switch(config-isis)# spf-delay 200
Switch(config-isis)# default system-id
```

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 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: <ul style="list-style-type: none"> • l1: Level-1 router type • l12: Not valid. 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. 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.

Variable	Value
	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
<code>WORD<0-255></code>	<p>Specifies a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763.</p> <p>By default, the system name comes from the host name configured at the system level.</p> <p>Use the <code>no</code> or <code>default</code> options to set this parameter to the default value (host name).</p> <p> Note:</p> <p>No consistency checks appear when you edit <code>sys-name</code>.</p>

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

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

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. Wait up to 20 minutes for the LSPs with the original system ID to age out.

Note:

To check the age out time, use the `show isis lsdb sysid <original-sys-id>` command on any of the other SPB nodes in the network. When there is no output from this command, proceed to the next step. The time left (in seconds) for the LSPs to age out is shown under the column LIFETIME.

6. Disable IS-IS.
7. Change the nickname to the original nickname.

8. Enable IS-IS.

Configuring optional IS-IS interface parameters

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

! **Important:**

Save your configuration using `save config` for the updates to be available after reboot. Saving the configuration also ensures that any authentication keys (passwords) specified during the configuration are properly encrypted.

Procedure

1. Enter Interface Configuration mode:

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

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

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

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

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>]	<p>Specifies the authentication type used for IS-IS hello packets on the interface. type can be one of the following:</p> <ul style="list-style-type: none"> • 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. <p>The default is none. Use the no or default options to set the hello-auth type to none.</p>
l1-dr-priority <0-127>	<p>Configures the level 1 IS-IS designated router priority to the specified value. The default value is 64.</p> <p>Use the no or default options to set this parameter to the default value of 64.</p> <p> Note:</p> <p>This parameter is not used for SPBM because SPBM only runs on point-to-point interfaces. This parameter is for</p>

Table continues...

Variable	Value
	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. Enter Privileged EXEC mode:
`enable`
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

```
Switch: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
Port1/21   pt-pt     Level 1    UP        UP         1        1          10
=====
```

```
Switch:1# show isis int-auth
```

```
=====
                        ISIS Interface Auth
=====
IFIDX      AUTH-TYPE  AUTH-KEYID  AUTH-KEY
=====
```

```

-----
Mlt2      none      0
Port1/21  none      0

Switch:1# show isis int-timers
=====
                        ISIS Interface Timers
=====
IFIDX      LEVEL      HELLO      HELLO      HELLO
                        INTERVAL     MULTIPLIER  DR
-----
Mlt2      Level 1      9          3          3
Port1/21  Level 1      9          3          3

Switch:1# show isis int-ckt-level
=====
                        ISIS Circuit level parameters
=====
IFIDX      LEVEL      DIS      CKTID
-----
Mlt2      Level 1      1
Port1/21  Level 1      2

```

Variable definitions

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

Variable	Value
[I1, I2, I12]	Displays the interface information for the specified level: I1, I2, or I12.

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 (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: <ul style="list-style-type: none"> • none for no authentication. • simple for a simple password. • hmac-md5 for MD5 encryption.
AUTH-KEYID	Shows the authentication password configured for the interface. If the KeyId is not configured, the value is 0.
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-timers` 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.

In SPBM, Backbone MAC (B-MAC) addresses are carried within the IS-IS link-state database. 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 allow 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. Enter Privileged EXEC mode:

```
enable
```

2. Display the SPBM IP unicast FIB:

```
show isis spbm ip-unicast-fib [all] [id <1-16777215>] [spbm-nh-as-  
mac]
```

*** Note:**

To display the IPv6 unicast FIB, use the `show isis spbm ipv6-unicast-fib` command.

3. Display the SPBM multicast FIB:

```
show isis spbm multicast-fib [vlan <1-4059>] [i-sid <1-16777215>]  
[nick-name <x.xx.xx>] [summary]
```

4. Display the SPBM unicast FIB:

```
show isis spbm unicast-fib [b-mac <0x00:0x00:0x00:0x00:0x00:0x00>]  
[vlan <1-4059>] [summary]
```

5. Display the SPBM unicast tree:

```
show isis spbm unicast-tree <1-4059> [destination <xxxx.xxxx.xxxx>]
```

Example

```
Switch# show isis spbm ip-unicast-fib
```

SPBM and IS-IS infrastructure configuration

```

=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF  DEST
ISID  ISID  Destination      NH  BEB   VLAN  OUTGOING  SPBM  PREFIX  PREFIX  IP ROUTE
      ISID  Destination      NH  BEB   VLAN  INTERFACE COST  COST  TYPE   PREFERENCE
-----
GRT   -    -   10.133.136.0/24  4K3(*) 4058  1/3       10    1    Internal 7
GRT   -    -   10.133.136.0/24  4K3(*) 4059  1/3       10    1    Internal 7
GRT   -    -   10.133.136.0/24  4K4(*) 4058  to_4k4    10000 1    Internal 7
GRT   -    -   10.133.136.0/24  4K4(*) 4059  to_4k4    10000 1    Internal 7
-----
Total number of SPBM IP-UNICAST FIB entries 4
=====

```

```
Switch# show isis spbm ip-unicast-fib id 10002
```

```

=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF  DEST
ISID  ISID  Destination      NH  BEB   VLAN  OUTGOING  SPBM  PREFIX  PREFIX  IP ROUTE
      ISID  Destination      NH  BEB   VLAN  INTERFACE COST  COST  TYPE   PREFERENCE
-----
vrf2  -   10002 65.2.2.0/24     ESS2 1000  1/13      20    20   Internal 7
vrf2  -   10002 65.2.2.0/24     ESS2 1001  1/18      20    20   Internal 7
-----
Total number of SPBM IP-UNICAST FIB entries 2
=====

```

```
Switch# show isis spbm ip-unicast-fib all
```

```

=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF  DEST
ISID  ISID  Destination      NH  BEB   VLAN  OUTGOING  SPBM  PREFIX  PREFIX  IP ROUTE
      ISID  Destination      NH  BEB   VLAN  INTERFACE COST  COST  TYPE   PREFERENCE
-----
GRT   -    -   1.0.0.1/32      ESP0 1000  1/13      20    1    Internal 7
GRT   -    -   1.0.0.1/32      ESP0 1000  1/18      20    1    Internal 7
vrf2  -   10002 65.2.2.0/24     ESS2 1000  1/13      20    20   Internal 7
vrf2  -   10002 65.2.2.0/24     ESS2 1001  1/18      20    20   Internal 7
-----
Total number of SPBM IP-UNICAST FIB entries 4
=====

```

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

```

=====
                        SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA          ISID      BVLAN  SYSID      HOST-NAME  OUTGOING-INTERFACES  INCOMING
      DA          ISID      BVLAN  SYSID      HOST-NAME  OUTGOING-INTERFACES  INTERFACE
-----
03:00:07:e4:e2:02 15000066 1001   0077.0077.0077 Switch-25   1/33           MLT-2
03:00:08:e4:e2:02 15000066 1001   0088.0088.0088 Switch-33   1/50,1/33     40.40.40.40
03:00:41:00:04:4d 1101     4058   00bb.0000.4100 Switch-1(*) 1/3,1/49,0.0.0.0
Tunnel to HQ
03:00:41:00:04:4f 1103     4058   00bb.0000.4100 Switch-1(*) 1/3,1/49,0.0.0.0 cpp
-----
Total number of SPBM MULTICAST FIB entries 4
=====

```

```
Switch# show isis spbm unicast-fib
```

```

=====
                        SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION          BVLAN  SYSID      HOST-NAME  OUTGOING  COST
ADDRESS              BVLAN  SYSID      HOST-NAME  INTERFACE
-----

```


```

-----
00:16:ca:23:73:df 1000 0016.ca23.73df SPBM-1 1/21 10
00:16:ca:23:73:df 2000 0016.ca23.73df SPBM-1 1/21 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
-----

```

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 <code>show isis spbm ip-unicast-fib</code> the device displays only GRT entries. The command shows IP routes from remote Backbone Edge Bridges (BEBs).
id <1-16777215>	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
vlan <1-4059>	Displays the FIB for the specified SPBM VLAN.
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.

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>	Displays the FIB for the specified B-MAC.
vlan <1-4059>	Displays the FIB for the specified SPBM VLAN.
summary	Displays a summary of the FIB.

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

Variable	Value
<1-4059>	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	Description
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 of the IP unicast FIB entry.
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.
PREFIX TYPE	Specifies the prefix type of the IP unicast FIB entry.
IP ROUTE PREFERENCE	Specifies the IP route preference of the IP unicast FIB entry.

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	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.

Table continues...

Parameter	Description
OUTGOING INTERFACE	Indicates the outgoing port of the multicast FIB entry.
INCOMING INTERFACE	Indicates the outgoing port 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. Display the IS-IS LSDB:

```
show isis lsdb [level {l1|l2|l12}] [sysid <xxxx.xxxx.xxxx>] [lspid
<xxxx.xxxx.xxxx.xx-xx>] [tlv <1-184>] [detail]
```

2. Display IS-IS adjacencies:

```
show isis adjacencies
```

3. Clear IS-IS LSDB:

```
clear isis lsdb
```

Example

```
Switch: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    Switch-Lab2
0018.b0bb.b3df.00-00      1         708       0xb9      0xcb1a    Switch-Lab1
=====
Level-1 : 3 out of 3 Total Num of LSP Entries
```

SPBM and IS-IS infrastructure configuration

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

```
Switch:1# show isis adjacencies
```

```
=====
                        ISIS Adjacencies
=====
INTERFACE                L STATE  UPTIME   PRI    HOLDDTIME  SYSID  HOST-NAME
-----
Mlt2                      1 UP     1d      03:57:25 127     20     0018.b0bb.b3df Switch-Lab1
Port1/21                  1 UP     1d      03:57:16 127     27     0016.ca23.73df Switch-Lab2
103.103.103.103          1 UP     00:01:04 127     22     a425.1b51.4884 DUT-B
Tunnel to HQ             1 UP     00:01:06 127     23     a425.1b51.4885 DUT-C
Port-1/25                 1 UP     00:01:06 127     18     a420.1b51.4886 DUT-D
=====
5 out of 5 Total Num of Adjacencies
=====
```

```
Switch: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 (Switch0)      Metric:10
                SPBM Sub TLV:
                port id: 640 num_port 1
                Metric: 10
                0000.beb1.00b1.01 (Switch1)      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 (Switch2)      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 (Switch1)  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)

```

Variable definitions

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

Variable	Value
detail	Displays detailed information.
level {/1 /2 /12}	Displays the LSDB for the specified level: l1, l2, or l12.
local	Displays IS-IS local LSDB information.
sysid <xxxx.xxxx.xxxx>	Displays the LSDB for the specified system ID.

Table continues...

Variable	Value
lspid <xxxx.xxxx.xxxx.xx-xx>	Displays the LSDB for the specified LSP ID.
tlv <1-184>	Displays the LSDB by TLV type.

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: I1, I2, or I12.
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.
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, MLT, or logical interface on which IS-IS exists.
L	Indicates the level of the adjacent router.

Table continues...

Parameter	Description
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. Display IS-IS system statistics:

```
show isis statistics
```

2. Display IS-IS interface counters:

```
show isis int-counters
```

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

```
show isis int-l1-ctrl-pkts
```

* Note:

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

4. Clear IS-IS statistics:

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

Example

```
Switch: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
Switch:1# show isis int-counters
=====
ISIS Interface Counters
=====
IFIDX    LEVEL  AUTH  ADJ    INIT    REJ    ID LEN  MAX AREA LAN  DIS
          FAILS  CHANGES  FAILS  ADJ
-----
Mlt2     Level 1-2  0      1      0      0      0      0      0      0
Port1/21 Level 1-2  0      1      0      0      0      0      0      0

Switch:1# show isis int-l1-ctrl-pkts
=====
ISIS L1 Control Packet counters
=====
IFIDX    DIRECTION  HELLO    LSP      CSNP     PSNP
-----
Mlt2     Transmitted 13346    231      2        229
Mlt2     Received   13329    230      1        230
Port1/21 Transmitted 13340    227      2        226
Port1/21 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.
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.
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 switch 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.
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-ctrl-pkts

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

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

Running the vms endura script

Use the following procedure to execute the run vms endura script.

This procedure applies only to VSP 4000.

*** Note:**

The switch should be in a factory default state before executing the `run vms endura switch` CLI command. When the CLI command is executed, you will be prompted about this.

About this task

Use switch number 5 on a VSP 4450GSX-PWR+, which is suited for the network core where the Pelco Endura video management systems exist.

Use switch number 6 on a VSP 4850GTS or VSP 4850GTS-PWR+, which is suited for the network edge where the IP cameras will connect. Up to 48 IP cameras can connect to a VSP 4850GTS-PWR+ switch within an IP subnet zone.

For each additional area and switch, increment the switch number by one, for example, switch number 7 for the second edge switch. The configuration is customized based on that number for the IP subnet, loopback addresses, and SPB information.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Execute the run vms endura script:

```
run vms endura switch <switch #>
```

Where, <switch #> is a numerical value from 5 to 99 used to seed unique values in the configuration script.

Note:

If the script causes a configuration conflict or cannot execute a command, an error message displays and the script stops.

Example

The following example shows the configuration of a switch in the VMS core, and shows the configuration file created by the script.

```
Switch:1>enable
Switch:1#config terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config)#run vms endura switch 5
Do you want to execute the run endura script? Device needs to be in factory default
state. (y/n) ? y

CP1 [01/01/15 06:28:09.928:UTC] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.

**Previous configurations stored in pre_endura_install.cfg**
**New VMS Endura configurations stored in new primary config file spb-switch-5.cfg**

*** VMS Endura script execution complete ***

CP1 [01/01/15 06:28:13.075:UTC] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.
Switch:1(config)#
Switch:1(config)#exit
```

The following example shows the configuration of a switch at the edge, and shows the configuration file created by the script.

```
Switch:1>enable
Switch:1#config terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config)#run vms endura switch 6
```

```

Do you want to execute the run endura script? Device needs to be in factory default
state. (y/n) ? y

CP1 [01/01/15 06:28:09.928:UTC] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.

**Previous configurations stored in pre_endura_install.cfg**
**New VMS Endura configurations stored in new primary config file spb-switch-6.cfg**


*** VMS Endura script execution complete ***

CP1 [01/01/15 06:28:13.075:UTC] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.
Switch:1(config)#
Switch:1(config)#exit

```

Variable definitions

Use the data in the following table to use the `run vms endura switch` command.

Variable	Value
<5-99>	<p>The numeric switch value used as a common element to configure switch parameters such as name, VLAN ID, SPB and IP parameters.</p> <p> Note:</p> <p>Avaya recommends using a value of 5 with a VSP4450GSX-PWR+ for the network core where the VMS servers are connected. Avaya also recommends using a value of 6 onwards for all VSP4850GTS / VSP4850GTS-PWR+ switches used for connecting IP Cameras at the network edge/access layer.</p>

Fabric Extend configuration using the CLI

The following sections provide procedural information you can use to configure Fabric Extend (FE) using the Command Line Interface (CLI).

Configuring Fabric Extend

Use the following procedure to configure Fabric Extend (FE) between a VSP 8000 in a Main office to a VSP 4000 in a Branch office. This is a typical deployment. However, if your deployment creates tunnels between two switches that support Fabric Extend natively (VSP 7200 Series and the VSP 8000 Series) then repeat those steps and ignore the steps for switches that require an ONA.

 **Note:**

VRF is an optional parameter. If a VRF is not configured, then FE uses the GRT.

Before you begin

The tunnel source IP address can be either a brouter port interface or a CLIP IP.

If using the tunnel originating address on the **GRT**, Fabric Extend has the following requirements:

- The tunnel source IP address must be on the GRT, not on a VRF.

*** Note:**

A Best Practice is to use separate IP addresses for the SPBM IP Shortcuts `ip-source-address` command and the Fabric Extend `ip-tunnel-source-address` command. However, if you want these IP addresses to be the same, you **MUST** exclude the `ip-source-address` address with an IS-IS accept policy. You cannot use the redistribute command with a route map exclusion.

Specify a CLIP interface to use as the source address for SPBM IP shortcuts.

- If IP Shortcuts is enabled, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

If you are using the tunnel originating address on a **VRF**, Fabric Extend has the following requirements:

- Configure a CLIP and tunnel source IP address on the VRF.
- Remote management of the VSP 4000 is only possible after establishing IP Shortcut over IS-IS. (Alternatively, you can enable GRT-VRF redistribution locally.)

About this task

Configuring Fabric Extend consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Note that the VSP 4000 source address command is different than the VSP 7200/VSP 8000 Series command. Also note that the logical interface commands are different between Layer 2 and Layer 3 networks.

Procedure

The following steps are for platforms that support FE natively.

1. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

2. Configure the IP tunnel source address:

```
ip-tunnel-source-address <A.B.C.D> [vrf WORD<1-16>]
```

3. Enter Global Configuration mode:

```
exit
```

4. Use one of the following commands to create a logical IS-IS interface:

- In a network with a Layer 3 Core, enter `logical-intf isis <1-255> dest-ip <A.B.C.D> [name WORD<1-16>]`

- In a network with a Layer 2 Core, enter `logical-intf isis <1-255> vid <list of vids> primary-vid <2-4059> port <slot/port> mlt <mltId> [name WORD<1-16>]`

*** Note:**

The primary VLAN ID (`primary-vid` must be one of the VLANs in the `vid <list of vids>`).

The following steps are for platforms that require an ONA to support FE.

*** Note:**

The interface VLAN connecting to the ONA network port is always in the GRT and the member port that the VLAN is part of is always an access port.

5. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

6. Configure the IP tunnel source address on the port that connects to the Device side of the ONA:

```
ip-tunnel-source-address <A.B.C.D> port <slot/port> [mtu <750-1950>]
[vrf WORD<1-16>]
```

7. Exit back into Global Configuration mode:

```
exit
```

8. Use one of the following commands to create a logical IS-IS interface:

- In a network with a Layer 3 Core, enter:

```
logical-intf isis <1-255> dest-ip <A.B.C.D> [name WORD<1-16>]
```

- In a network with a Layer 2 Core, enter:

```
logical-intf isis <1-255> vid <list of vids> primary-vid <2-4059>
port <slot/port> mlt <mltId> [name WORD<1-16>]
```

*** Note:**

The primary VLAN ID (`primary-vid`) must be one of the VLANs in the `vid <list of vids>`.

Variable definitions

Use the data in the following tables to use the `ip-tunnel-source-address` command.

To delete an IS-IS IP tunnel source address, use the `no ip-tunnel-source-address` option.

*** Note:**

The `port` parameter is for the VSP 4000 only.

Variable	Value
<A.B.C.D>	Specifies the IS-IS IPv4 tunnel source address, which can be either a brouter interface IP or a CLIP IP.
port <slot/port>	Specifies the port that is connected to the ONA's Device port.
vrf WORD<1–16>	Specifies the VRF name associated with the IP tunnel.
mtu <750–1950>	Specifies the size of the maximum transmission unit (MTU). The default is 1950. This parameter only applies to an ONA configuration.

Use the data in one of the following tables to use the `logical-intf isis` command, depending on whether you have a Layer 2 or Layer 3 core.

To delete a logical IS-IS interface, use the `no logical-intf isis` option.

Table 5: Layer 2 core

Variable	Value
<1–255>	Specifies the index number that uniquely identifies this logical interface.
port {slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Specifies the physical port that the logical interface is connected to in a Layer 2 network.
vid <list of vids>	Specifies the list of VLANs that are associated with this logical interface.
primary-vid <2–4059>	Specifies the primary tunnel VLAN ID associated with this Layer 2 IS-IS logical interface.
mlt <mltld>	Specifies the MLT ID that the logical interface is connected to in a Layer 2 network.
name WORD<1–16>	Specifies the administratively-assigned name of this logical interface, which can be up to 16 characters.

Table 6: Layer 3 core

Variable	Value
<1–255>	Specifies the index number that uniquely identifies this logical interface.
dest-ip <A.B.C.D>	Specifies the tunnel destination IP address of the remote BEB.
name WORD<1–16>	Specifies the administratively-assigned name of this logical interface, which can be up to 16 characters.

Displaying IS-IS logical interfaces

Use the following procedure to display the IS-IS logical interfaces configured on the switch.

Procedure

Display the IS-IS logical interfaces:

```
show isis logical-interface [name]
```

Examples

Example of a Layer 2 Core:

```
Switch:1# show isis logical-interface
=====
ISIS Logical Interfaces
=====
IFIDX  NAME  ENCAP      L2_INFO      TUNNEL      L3_TUNNEL  NEXT_HOP_INFO
      TYPE      PORT/MLT  VIDS (PRIMARY)  DEST-IP      PORT/MLT  VLAN   VRF
-----
1      --    L2-P2P-VID  Port2/40    101,201 (101)  --          --      --   --
2      --    L2-P2P-VID  Port1/3     102,202 (102)  --          --      --   --
-----
2 out of 2 Total Num of Logical ISIS interfaces
=====
```

Example of a Layer 3 Core:

```
Switch:1# show isis logical-interface
=====
ISIS Logical Interfaces
=====
IFIDX  NAME  ENCAP      L2_INFO      TUNNEL      L3_TUNNEL  NEXT_HOP_INFO
      TYPE      PORT/MLT  VIDS (PRIMARY)  DEST-IP      PORT/MLT  VLAN   VRF
-----
1      SPBoIP_T1  IP        --           --           41.41.41.41  MLT10    2    vrf24
2      SPBoIP_T2  IP        --           --           42.42.42.42  MLT10    2    vrf24
3      SPBoIP_4K5  IP        --           --           187.187.187.187  MLT10    2    vrf24
-----
3 out of 3 Total Num of Logical ISIS interfaces
=====
```

Display the IS-IS logical interface names.

```
Switch:1# show isis logical-interface name
=====
ISIS Logical Interface name
=====
ID     NAME
-----
1      SPBoIP_T1
2      SPBoIP_T2
3      SPBoIP_4K5
-----
3 out of 3 Total Num of Logical ISIS interfaces
=====
```

The command **show isis logical-interface** truncates the IS-IS logical interface name to the first 16 characters. To view the entire name (up to a maximum of 64 characters), use the command **show isis logical-interface name**.

For example:

```
Switch:1# show isis logical-interface name
=====
ISIS Logical Interface name
=====
ID      NAME
-----
6       This_Is_A_50_Character_ISIS_Logical_Interface_Name
-----
1 out of 1 Total Num of Logical ISIS interfaces
=====
```

Job Aid

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

Parameter	Description
IFIDX	Specifies an index value for this logical interface.
NAME	Specifies the administratively-assigned name of this logical interface, which can be up to 16 characters.
ENCAP TYPE	Specifies whether the encapsulation type for the logical interface is Layer 2 (L2-P2P-VID) or Layer 3 (IP).
L2_INFO PORT/MLT	Specifies the port or MLT that the logical interface is connected to in an L2 network.
L2_INFO VLAN	Specifies the list of VLANs that are associated with this L2 logical interface.
TUNNEL DEST-IP	Specifies the destination IP address for the logical interface.
L3_TUNNEL_NEXT_HOP_INFO PORT/MLT	Specifies the outgoing interface (port or MLT) for VXLAN traffic.
L3_TUNNEL_NEXT_HOP_INFO VLAN	Specifies the outgoing VLAN interface for VXLAN traffic.
L3_TUNNEL_NEXT_HOP_INFO VRF	Specifies the name of the VRF that this L3 logical interface is configured on.

Displaying KHI Fabric Extend ONA status

About this task

* Note:

This feature only applies to platforms that have an Open Networking Adapter (ONA) connected to it.

Use the following command to display the current status of the Fabric Extend ONA, which includes release information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the ONA status:

```
show khi fe-ona status
```

Example

The following output displays the `show khi fe-ona status` when the ONA is operating normally.

```
Switch:1#show khi fe-ona status
```

```
=====
                        ONA STATUS
=====
ONA Device Status : UP
Running Release Name : v1.0.0.0int006-3-g9749735-dirty
Last Image Upgrade Status : UPGRADE_SUCCESS
Last Image File Used For Upgrde: gdb-secure_ona.tgz
=====
```

The following examples display the output when communication from the switch to the ONA is disrupted. Note that the `ONA Down reason` lists the cause of the failure. The reason changes depending on the context of the failure.

The following output displays when the configuration push from the switch to the ONA fails:

```
Switch:1#show khi fe-ona status
```

```
=====
                        ONA STATUS
=====
ONA Device Status : DOWN
ONA DOWN reason : ONA_CONFIG_DOWNLOAD_FAILED
Running Release Name :
Image Upgrade Status : UNKNOWN
=====
```

The following output displays when the port connecting to the ONA device port is DOWN:

```
Switch:1#show khi fe-ona status
```

```
=====
                        ONA STATUS
=====
ONA Device Status : DOWN
```

```
ONA DOWN reason : ONA_DEVICE_PORT_DOWN
Running Release Name :
Image Upgrade Status : UNKNOWN
Image File Is Being Used For Upgrade :
-----
```

The following output displays when the switch is not receiving LLDP packets from the ONA:

```
Switch:1#show khi fe-ona status

=====
                        ONA STATUS
=====
ONA Device Status : DOWN
ONA DOWN reason : ONA_LLDP_TIMEOUT
Running Release Name :
Image Upgrade Status : UNKNOWN
-----
```

*** Note:**

On the switch console, the following log message precedes all three of the above cases:

```
CP1 [03/22/71 09:30:15.336:UTC] 0x00378601 00000000 GlobalRouter ONA
WARNING ONA device status detected down
```

Displaying KHI Fabric Extend ONA global information

About this task

*** Note:**

This feature only applies to platforms that have an Open Networking Adapter (ONA) connected to it.

Use the following command to display Fabric Extend ONA global information such as port numbers, IP addresses, and MTU.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the ONA global information:

```
show khi fe-ona detail
```

Example

```
Switch:1#show khi fe-ona detail

=====
                        ONA RUNTIME INFORMATION
=====
ONA Port Number : 1/15
ONA Management Address : 100.1.1.11
Tunnel Source IP Address : 11.11.12.11
ONA LLDP Port Status : Enabled
ONA Device Port Status : UP
ONA Device Status : UP
MTU : 1000
ONA Network Port Number : 1/35
```

```

ONA Mac(ARP) Address : 10:cd:ae:69:b6:50
ONA Source VlanId : 1050
ONA Source VlanIP : 10.0.70.1
ONA Gateway IP : 10.0.70.1
ONA Management IP Mask : 255.255.255.0
ONA Bootmode : 1
ONA Uptime : 0 day(s), 00:00:00
pbit-to-dscp-map p0=16 p1=20 p2=24 p3=30 p4=36 p5=40 p6=48 p7=46
-----

```

*** Note:**

In the above example, the switch receives LLDP packets with the Management IP address of the ONA over the ONA Port (1/15). The switch extracts the ONA Management IP from the LLDP packet and resolves the ARP of the ONA over the network port (1/35). After the switch resolves the ARP of the ONA IP, the `show khi fe-ona detail` updates the following details:

- ONA Network Port Number
- ONA Mac(ARP) Address
- ONA Source VlanId

Note the following in regard to the `show khi fe-ona detail` output shown above:

- ONA Source VlanIP : 10.0.70.1—This is the IP address of the switch VLAN that switches traffic to the ONA network port. In the above output, this is VLAN 1050.
- ONA Gateway IP : 10.0.70.1—This is the ONA gateway IP address that the switch gets by querying the ONA. The ONA receives this gateway IP from the DHCP server.

! Important:

The ONA Source VlanIP, and ONA Gateway IP addresses must be the same for the tunnels to come up and the traffic to switch.

Fabric Attach configuration using the CLI

The following sections provide procedural information you can use to configure Fabric Attach (FA) and Logical Link Discovery Protocol (LLDP) using the Command Line Interface (CLI).

Configuring Fabric Attach globally

For proper operation, FA must be enabled at both the global level and at the interface level on the FA Server. By default, FA is globally enabled. However, FA is disabled by default at the interface level and must be explicitly enabled on each interface.

Use this procedure to enable Fabric Attach globally on a switch.

Procedure

1. Enter Global Configuration mode:

```

enable
configure terminal

```


2. Enable FA:

```
fa enable
```

3. (Optional) Disable FA:

```
no fa enable
```

 **Caution:**

Disabling FA flushes all FA element discovery and mappings.

4. View the FA configuration status. Use one of the following commands:

- show fa
- show fa agent

Example

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#fa enable
Switch:1(config)#show fa

=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 240
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm

Switch:1(config)#show fa agent

=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 240
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm
```

Configuring Fabric Attach discovery timeout

Use this procedure to configure the Fabric Attach discovery time-out.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the FA discovery time-out in seconds:

```
fa discovery-timeout <45-480>
```

*** Note:**

The discovery time-out must be greater than or equal to the assignment time-out.

3. (Optional) Configure the default FA discovery time-out:

```
default fa discovery-timeout
```

Example

Configure the FA discovery time-out.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#fa discovery-timeout 50
```

Verify the configuration.

```
Switch:1(config)#show fa

=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 45
                          FA Discovery Timeout : 50
                          FA Provision Mode : spbm
```

Variable definitions

Use the data in the following table to use the **fa discovery-timeout** command.

Variable	Value
<45-480>	Specifies the Fabric Attach discovery time-out in seconds. The default value is 240 seconds.

Configuring Fabric Attach assignment timeout

Use this procedure to configure the Fabric Attach assignment time-out.

Procedure**1. Enter Global Configuration mode:**

```
enable
configure terminal
```

2. Configure the FA assignment time-out in seconds:

```
fa assignment-timeout <45-480>
```

*** Note:**

The assignment time-out must be less than or equal to the discovery time-out.

3. (Optional) Configure the default FA assignment time-out value:

```
default fa assignment-timeout
```

Example

Configure the FA assignment time-out:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#fa assignment-timeout 50
```

Verify the configuration:

```
Switch:1(config)#show fa

=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 50
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm
```

Variable definitions

Use the data in the following table to use the **fa assignment-timeout** command.

Variable	Value
<45-480>	Specifies the Fabric Attach assignment time-out in seconds. The default value is 240 seconds.

Enabling Fabric Attach on an interface

Use this procedure to enable Fabric Attach on an interface (port, static MLT or LACP MLT). Enabling FA on an MLT enables FA on all ports of the MLT. If your platform supports channelization, FA can also be enabled on channelized ports.

Before you begin

Verify that FA is enabled globally on the switch.

About this task

Enabling FA on a port or MLT is necessary for element discovery.

On the FA Server, FA is enabled globally by default. However, you must explicitly enable FA on the desired port or MLT interface. FA is successfully enabled on an MLT only if all ports of the MLT have FA successfully enabled. Enabling FA automatically configures LLDP on all ports. Tagging is configured and spanning tree is disabled.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
```

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

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable FA on the interface:

```
fa enable
```

3. **(Optional)** Disable FA on the interface:

```
no fa enable
```

⚠ Caution:

Disabling FA flushes all FA element discovery and I-SID-to-VLAN mappings associated with the interface.

4. View the FA configuration status:

```
show fa interface [disabled-auth] [enabled-auth] [mlt <1-512>] [port
<{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}>]
```

Example

Enable FA on a port:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/2
Switch:1(config-if)#fa enable
Switch:1(config-if)#exit
Switch:1(config)#
```

Enable FA on an MLT:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#exit
Switch:1(config)#
```

Verify that FA is enabled on the interfaces.

*** Note:**

When FA is enabled, message authentication is enabled by default. The authentication key is set to the default value and appears encrypted on the output.

```
Switch:1(config)#show fa interface
```

```
=====
Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH  MSG AUTH
STATUS         ISID    CVID    STATUS  KEY
```

SPBM and IS-IS infrastructure configuration

```
-----  
Port1/1      enabled  0      0      enabled  ****  
Port1/2      enabled  0      0      enabled  ****  
Mlt1         enabled  0      0      enabled  ****  
Mlt10        enabled  0      0      enabled  ****  
-----  
4 out of 4 Total Num of fabric attach interfaces displayed  
-----
```

For example, disable FA on port 1/1 and Mlt1.

```
Switch:1(config)#interface gigabitethernet 1/1  
Switch:1(config-if)#no fa enable  
Switch:1(config-if)#exit  
Switch:1(config)#interface mlt 1  
Switch:1(config-mlt)#no fa enable  
Switch:1(config-mlt)#exit
```

Verify that FA is disabled on port 1/1 and Mlt1.

```
Switch:1(config)#show fa interface  
  
===== Fabric Attach Interfaces =====  
-----  
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH  MSG AUTH  
                STATUS  ISID    CVID    STATUS    KEY  
-----  
Port1/1        disabled 0      0      enabled   ****  
Port1/2        enabled  0      0      enabled   ****  
Mlt1           disabled 0      0      enabled   ****  
Mlt10          enabled  0      0      enabled   ****  
-----  
4 out of 4 Total Num of fabric attach interfaces displayed  
-----
```

View the FA interfaces that have authentication enabled:

```
Switch:1(config)#show fa interface enabled-auth  
  
===== Fabric Attach Interfaces =====  
-----  
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH  MSG AUTH  
                STATUS  ISID    CVID    STATUS    KEY  
-----  
Port1/2        enabled 0      0      enabled   ****  
Mlt10          enabled 0      0      enabled   ****  
-----  
2 out of 2 Total Num of fabric attach interfaces displayed  
-----
```

Optionally, disable FA message authentication on 1/1 and Mlt1.

```
Switch:1(config)#interface gigabitethernet 1/1  
Switch:1(config-if)#no fa message-authentication  
Switch:1(config-if)#exit  
Switch:1(config)#interface mlt 1  
Switch:1(config-mlt)#no fa message-authentication  
Switch:1(config-mlt)#exit
```

Verify that both FA and FA message authentication are disabled on 1/1 and Mlt1, as indicated by the SERVER STATUS and MSG AUTH STATUS fields respectively.

```
Switch:1(config)#show fa interface
```

```

=====
                          Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT   MGMT   MSG AUTH MSG AUTH
              STATUS ISID   CVID   STATUS   KEY
-----
Port1/1        disabled 0      0      disabled ****
Port1/2        enabled  0      0      enabled  ****
Mlt1           disabled 0      0      disabled ****
Mlt10          enabled  0      0      enabled  ****
-----
4 out of 4 Total Num of fabric attach interfaces displayed
=====

```

View the FA interfaces that have authentication disabled:

```
Switch:1(config)#show fa interface disabled-auth
```

```

=====
                          Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT   MGMT   MSG AUTH MSG AUTH
              STATUS ISID   CVID   STATUS   KEY
-----
Port1/1        disabled 0      0      disabled ****
Mlt1           disabled 0      0      disabled ****
-----
2 out of 2 Total Num of fabric attach interfaces displayed
=====

```

Variable definitions

Use the data in the following table to use the **show fa interface** command.

Variable	Value
disabled-auth	Displays the FA interfaces (port or MLT) that have authentication disabled.
enabled-auth	Displays the FA interfaces (port or MLT) that have authentication enabled.
<1-512>	The valid range for MLT ID. Displays FA configuration on the specified MLT interface.
port {slot/port[/sub-port] [-slot/port[/sub-port]] [...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. Displays FA configuration on the specified port.

Configuring FA message authentication on an interface

Use this procedure to configure FA message authentication on an interface (port or MLT).

Procedure

1. Enter Interface Configuration mode:

```
enable

configure terminal

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

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure FA message authentication on a port or MLT:

```
[default] [no] fa message-authentication
```

*** Note:**

When FA is enabled, message authentication is enabled by default. The authentication key is set to the default value and appears encrypted on the output.

Example

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#
```

Enable message authentication on a port.

```
Switch:1(config)#interface gigabitEthernet 1/2

Switch:1(config-if)#fa message-authentication
Switch:1(config-if)#show fa interface port 1/2

=====
Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH MSG AUTH
STATUS         ISID   CVID    STATUS  KEY
-----
Port1/2        enabled 0       0       enabled ****

-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
Switch:1(config-if)#exit
Switch:1(config)#
```

Enable message authentication on an MLT.

```
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#fa message-authentication
Switch:1(config-mlt)#show fa interface mlt 10

=====
Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH MSG AUTH
STATUS         ISID   CVID    STATUS  KEY
-----
Mlt10          enabled 0       0       enabled ****
```



```
-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
```

```
Switch:1(config-mlt)#exit
Switch:1(config)#
```

The following example demonstrates disabling message authentication on a port or MLT.

```
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#no fa message-authentication
Switch:1(config-if)#exit
Switch:1(config)
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#no fa message-authentication
```

```
Switch:1(config-mlt)#show fa interface
```

```
=====
Fabric Attach Interfaces
=====
```

INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY
Port1/2	enabled	0	0	disabled	****
Mlt10	enabled	0	0	disabled	****

```
-----
2 out of 2 Total Num of fabric attach interfaces displayed
-----
```

Configuring the FA authentication key on an interface

On the FA Server, you can configure an authentication key on an interface (port, static MLT or LACP MLT), to authenticate a client or proxy device on that interface. The authentication key is stored in encrypted form when you save configuration on the FA Server.

Before you begin

Ensure that:

- On the FA Server, FA is enabled globally and also on the interface.
- FA message authentication is enabled on the interface.

Note:

By default, enabling FA enables message authentication. The authentication key is set to the default value and appears encrypted on the output.

About this task

Use this procedure to configure an FA authentication key on a specified port or on all ports of an MLT, on the switch. If you do not configure an authentication key, the default value is used. If you specify a key, the default value is overridden and is stored in encrypted format in a separate file other than the configuration file, when you execute the `save config` command.

Caution:

For an FA Client or an FA Proxy device to successfully authenticate and attach to the FA Server, the authentication key *must* match on both the client and the server. If the authentication key is changed on the FA Server switch, it must correspondingly be changed on the FA Client or Proxy attached to it, for FA to operate properly.

Procedure

1. Enter Interface Configuration mode:

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

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the FA authentication key:

```
fa authentication-key WORD<0-32>
```

3. (Optional) Configure the default FA authentication key:

```
default fa authentication-key
```

Example

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
```

Enable FA and message authentication on a port. Configure the authentication key `phone-network` on the port.

```
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#fa enable
Switch:1(config-if)#fa message-authentication
Switch:1(config-mlt)#fa authentication-key phone-network
Switch:1(config-if)#exit
Switch:1(config)#
```

Enable FA and message authentication on an MLT. Configure the authentication key `client-network` on the MLT.

```
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#fa message-authentication
Switch:1(config-mlt)#fa authentication-key client-network
```

Verify configuration of the FA authentication key. The authentication key appears encrypted on the output.

```
Switch:1(config-if)#show fa interface
```

```
=====
                          Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT   MGMT   MSG AUTH  MSG AUTH
                STATUS ISID   CVID   STATUS    KEY
-----
Port1/2        enabled  0      0      enabled  ****
MLT10          enabled  0      0      enabled  ****
=====
```

2 out of 2 Total Num of fabric attach interfaces displayed

Variable Definitions

Use the data in the following table to use the `fa authentication-key` command.

Variable	Value
WORD<0–32>	Specifies the authentication key on the port or MLT.

Configuring FA management on a port or MLT

Use this procedure to configure a management I-SID on an FA enabled port or MLT on the switch.

Before you begin

Ensure that the port or MLT is enabled for Fabric Attach.

About this task

This command applies to all traffic sent or received on a port or MLT, carrying the VLAN ID specified using the `c-vid` parameter. This parameter is optional.

Depending on whether the `c-vid` parameter is specified or not, the behavior is as follows:

- If the `c-vid` parameter is specified, the FA Server transmits this VLAN ID as the management VLAN in the FA Element TLV. A client or proxy receiving this TLV uses this VLAN-ID for management traffic on the FA Server uplink.
- If the `c-vid` parameter is *not* specified, the FA Server transmits a management VLAN with a VLAN ID value of 4095 in the FA Element TLV. A client or proxy receiving this TLV uses **untagged** traffic for network management on the FA Server uplink.

Procedure

1. Enter Interface Configuration mode:

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

Note:

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the FA management I-SID:

```
fa management i-sid <1-16777215> [c-vid <1-4094>]
```

Important:

If you do not specify a C-VID value, the port or MLT is **untagged**.

3. Delete FA management I-SID on a port or MLT using one of the following commands:

- default fa management i-sid
- no fa management i-sid

4. Verify configuration of FA management on the port or MLT, using the following commands:

- show i-sid [<1-16777215>]
- show interfaces gigabitEthernet i-sid [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]]
- show mlt i-sid [<1-512>]

Example

The following example demonstrates configuring FA management on the port 1/2.

Configure FA management on port 1/2:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#fa management i-sid 101 c-vid 101

Switch:1(config-if)#show i-sid 101

=====
                                Isid Info
=====
ISID      ISID      VLANID    PORT      MLT      ORIGIN
ID        TYPE      ID        INTERFACES INTERFACES
-----
101      ELAN      N/A      c101:1/2  -        MANAGEMENT

Switch:1(config-if)#show interfaces gigabitEthernet i-sid

=====
                                PORT Isid Info
=====
PORTNUM  IFINDEX  ISID      VLANID  C-VID    ISID      ORIGIN      BPDU
          ID
-----
1/2      193     101      N/A     101     ELAN      MANAGEMENT

1 out of 1 Total Num of i-sid endpoints displayed
```

The following example demonstrates configuring FA management on an MLT.

Configure FA management on MLT 10.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#fa management i-sid 100

Verify configuration of FA management on the MLT. Since the C-VID is not specified, the MLT is displayed as untagged.

Switch:1(config-mlt)#show i-sid 100

=====
```

```

=====
                          Isid Info
=====
ISID      ISID      VLANID    PORT      MLT      ORIGIN
ID        TYPE                    INTERFACES INTERFACES
-----
100      ELAN      N/A       -         u:10    MANAGEMENT
=====

```

An FA management I-SID can have a platform VLAN associated with it. For Layer 3 support on the management I-SID, you must create a platform VLAN by port and associate the platform VLAN with the management I-SID. The C-VID can be of the same value or of a different value than that of the platform VLAN.

If the management I-SID matches one of the FA Switched UNI (ELAN) I-SIDs (as displayed by the command `show i-sid elan`), then the platform VLAN is automatically associated with the FA enabled interface (port or MLT).

In the following example, for Layer 3 support, create a platform VLAN 999 and associate it with the management I-SID 101.

```

Switch:1(config-if)#vlan create 999 type port-mstprstp 0
Switch:1(config-if)#vlan i-sid 999 101

```

```

Switch:1(config-if)#show i-sid

```

```

=====
                          Isid Info
=====
ISID      ISID      VLANID    PORT      MLT      ORIGIN
ID        TYPE                    INTERFACES INTERFACES
-----
101      ELAN      999      c101:1/2  -         MANAGEMENT
=====
c: customer vid      u: untagged-traffic
All 1 out of 1 Total Num of i-sids displayed

```

```

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

```

```

=====
                          Vlan I-SID
=====
VLAN_ID   I-SID
-----
999       101
1 out of 1 Total Num of Vlans displayed

```

Since the management I-SID matches one of the FA Switched UNI (ELAN) I-SIDs, the platform VLAN is automatically associated with the FA enabled port 1/2.

```

Switch:1(config-if)#show interfaces gigabitEthernet i-sid

```


```

=====
                          PORT Isid Info
=====
PORTNUM   IFINDEX  ISID      VLANID  C-VID  ISID      ORIGIN      BPDU
          ID      TYPE
-----
1/2       193     101      999     101    ELAN      MANAGEMENT
-----
1 out of 1 Total Num of i-sid endpoints displayed

```

Variable Definitions

Use the data in the following table to use the **fa management** command.

Variable	Value
i-sid <1-16777215>	Specifies the management I-SID
c-vid <1-4094>	Specifies the C-VID value of the port or MLT on the FA Server.  Important: If you do not specify a C-VID value, the port or MLT is untagged .

Viewing Fabric Attach global configuration status

Use this procedure to display the Fabric Attach global configuration status on a switch.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Display the FA configuration status using one of the following commands:
 - `show fa`
 - `show fa agent`

Example

Sample output for the **show fa** command:

```
Switch:1#show fa
=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 240
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm
```

Sample output for the **show fa agent** command:

```
Switch:1#show fa agent
=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 240
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm
```

Viewing Fabric Attach interface configuration

Use this procedure to view FA interface configuration.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View all FA interfaces (ports and MLTs):

```
show fa interface
```

3. To view FA interface configuration on ports, use one of the following commands:

- View FA configuration on all ports:

```
show fa interface port
```

- View FA configuration on a specific port, enter:

```
show fa interface port [{slot/port[/sub-port] [-slot/port[/sub-
port]] [,...]]
```

4. To view FA interface configuration on MLTs, use one of the following commands:

- View FA configuration on all MLTs:

```
show fa interface mlt
```

- View FA configuration on a specific MLT:

```
show fa interface mlt [<1-512>]
```

Example

The following example displays sample outputs for the **show fa interface** command.

```
Switch:1>en
Switch:1#show fa interface

=====
                          Fabric Attach Interfaces
=====
INTERFACE    SERVER   MGMT    MGMT    MSG AUTH  MSG AUTH
              STATUS  ISID    CVID    STATUS    KEY
-----
Port2/10     enabled  0       0       enabled   ****
Port4/6      enabled  0       0       enabled   ****
Port4/11     enabled  0       0       enabled   ****
Mlt2         enabled  0       0       enabled   ****

-----
4 out of 4 Total Num of fabric attach interfaces displayed
-----
```

The following is a sample output for the **show fa interface** command for the port 2/10.

```
Switch:1#show fa interface port 2/10

=====
                          Fabric Attach Interfaces
=====
INTERFACE    SERVER   MGMT    MGMT    MSG AUTH  MSG AUTH
              STATUS  ISID    CVID    STATUS    KEY
-----
Port2/10     enabled  0       0       enabled   ****
```



```
-----
1 out of 4 Total Num of fabric attach interfaces displayed
-----
```

The following is a sample output for the **show fa interface** command for the MLT 2.

```
Switch:1#show fa interface mlt 2

=====
                          Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT   MGMT   MSG AUTH  MSG AUTH
                STATUS ISID   CVID    STATUS    KEY
-----
Mlt2           enabled 0      0      enabled   ****

-----
1 out of 4 Total Num of fabric attach interfaces displayed
-----
```

Variable definitions

Use the data in the following table to use the `show fa interface port` command.

Variable	Value
{slot/port[/sub-port][/-slot/port[/sub-port]] [...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Use the data in the following table to use the `show fa interface mlt` command.

Variable	Value
<1-512>	The valid range for MLT ID.

Viewing Fabric Attach discovered elements

Use this procedure to view Fabric Attach discovered elements.

About this task

When FA is enabled on an FA Server switch, LLDP PDUs are exchanged between the FA Server and FA Clients or FA Proxies. Standard LLDPs allow neighbors to be learned. With the help of organizational-specific element discovery TLVs, the client or proxy recognizes that it has attached to the FA Server. Only after the discovery handshake is complete, an FA Client or FA Proxy can transmit I-SID-to-VLAN assignments to join the SPB Fabric network through the FA Server.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Display FA discovered elements:


```
show fa elements
```

3. Display FA discovered elements on a specific port:

```
show fa elements [{slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]]
```

Example

The following example displays the sample output for the `show fa elements` command.

```
Switch:1#show fa elements
=====
Fabric Attach Discovery Elements
=====
PORT      TYPE      MGMT      ELEM ASGN
          VLAN STATE  SYSTEM ID AUTH AUTH
-----
1/5      proxy     710 T / S    50:61:84:ee:8c:00:20:00:00:01  AP  AP
1/6      proxy     710 T / S    50:61:84:ee:8c:00:20:00:00:01  AP  AP
=====
Fabric Attach Authentication Detail
=====
PORT      ELEM OPER      ASGN OPER
          AUTH STATUS   AUTH STATUS
-----
1/5      successAuth    successAuth
1/6      successAuth    successAuth

State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid

Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,  N= None

-----
2 out of 2 Total Num of fabric attach discovery elements displayed
```

Variable definitions

Use the data in the following table to use the `show fa elements` command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Viewing Fabric Attach I-SID-to-VLAN assignments

Use this procedure to display the I-SID-to-VLAN assignments advertised by an FA Client or an FA Proxy, to be supported on the FA Server. These assignments can be accepted or rejected by the FA Server. An assignment that is successfully accepted by the FA Server results in the creation of a Switched UNI I-SID on the interface.

Before you begin

Verify that IS-IS and SPBM are properly configured on the FA Server switch.

- Verify SPBM configuration using the command `show running-config module spbm`.
- Verify IS-IS configuration using one of the following commands:
 - `show isis`
 - `show isis interface`
 - `show isis adjacency`
 - `show isis lsdb`

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display FA I-SID-to-VLAN assignments:

```
show fa assignment
```

3. Display FA I-SID-to-VLAN assignments on specific ports:

```
show fa assignment [{slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]]
```

Example

The following example displays a sample output for the `show fa assignment` command.

* Note:

The state of I-SID-to-VLAN assignments on a client or proxy device is pending until it is changed by the FA Server to active or reject.

```
Switch:>en
Switch:1#show fa assignment
=====
Fabric Attach Assignment Map
=====
Interface  I-SID      Vlan      State      Origin
-----
1/1        2          2         active     proxy
1/2        3          3         active     proxy
1/2        4          4         active     proxy
1/3        5          5         reject     proxy
-----
4 out of 4 Total Num of fabric attach assignment mappings displayed
=====
```

Variable definitions

Use the data in the following table to use the `show fa assignment` command.

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]] [,...]	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Viewing Fabric Attach statistics

If FA discovery fails, use this procedure to display FA statistics to determine if FA discovery TLVs were processed. You can also view the FA assignment statistics to determine the number of FA assignments that were accepted or rejected by the FA Server.

You can view the statistics at either the global level or at the port (interface) level.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View global level FA statistics:

```
show fa statistics [summary]
```

3. View FA statistics at the slot/port level:

```
show fa statistics [{slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...]]
```

Note:

If a slot is removed from the switch chassis, the FA statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

4. (Optional) Clear FA statistics:

```
clear fa statistics [summary] [{slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...]]
```

Examples

Viewing FA discovery and assignment statistics:

```
Switch:1>en  
Switch:1#show fa statistics
```

```
=====
                          Fabric Attach STATISTICS
=====
Port      DiscElem  DiscElem  DiscElem  DiscAuth
         Received  Expired   Deleted   Failed
-----
1/1       3057      0          1          0
1/2       2000      0          1          0
=====
                          Fabric Attach ASSIGNMENTS STATISTICS
```

SPBM and IS-IS infrastructure configuration

```
=====
```

Port	Asgn Received	Asgn Accepted	Asgn Rejected	Asgn Expired	Asgn Deleted	AsgnAuth Failed
1/1	3149	3	1	3	0	0
1/2	1500	0	1	2	0	0

```
=====
```

View a summary of the FA discovery and assignment statistics:

```
Switch:1#show fa statistics summary
```

```
=====
```

Fabric Attach STATISTICS SUMMARY				
Port	DiscElem Received	DiscElem Expired	DiscElem Deleted	DiscAuth Failed
1/1	3057	0	1	0
1/2	2000	0	1	0

```
=====
```

Fabric Attach ASSIGNMENTS STATISTICS SUMMARY						
Port	Asgn Received	Asgn Accepted	Asgn Rejected	Asgn Expired	Asgn Deleted	AsgnAuth Failed
1/1	3149	3	1	3	0	0
1/2	1500	0	1	2	0	0

```
=====
```

Viewing FA statistics on a specific port (port 1/1):

```
Switch:1>en  
Switch:1#show fa statistics 1/1
```

```
=====
```

Fabric Attach STATISTICS				
Port	DiscElem Received	DiscElem Expired	DiscElem Deleted	DiscAuth Failed
1/1	3057	0	1	0

```
=====
```

Fabric Attach ASSIGNMENTS STATISTICS						
Port	Asgn Received	Asgn Accepted	Asgn Rejected	Asgn Expired	Asgn Deleted	AsgnAuth Failed
1/1	3149	3	1	3	0	0

```
=====
```

Optionally, clear FA statistics and verify that the statistics are cleared.

```
Switch:1#clear fa statistics  
Switch:1#show fa statistics
```

```
=====
```

Fabric Attach STATISTICS				
Port	DiscElem Received	DiscElem Expired	DiscElem Deleted	DiscAuth Failed
1/1	0	0	0	0
1/2	0	0	0	0

```
=====
```

```

=====
Fabric Attach ASSIGNMENTS STATISTICS
=====
Port      Asgn      Asgn      Asgn      Asgn      Asgn      AsgnAuth
  Received Accepted Rejected Expired   Deleted   Failed
-----
1/1       0         0         0         0         0         0
1/2       0         0         0         0         0         0

```

Variable Definitions

Use the data in the following table to use the `show fa statistics` command.

Variable	Value
summary	Displays a summary of Fabric Attach element discovery and assignment statistics at the global level.
{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configuring global LLDP transmission parameters

Before you begin

- In the GigabitEthernet Interface Configuration mode, specify the LLDP port status as transmit only or transmit and receive.

About this task

Use this procedure to configure global LLDP transmission parameters on the switch. If required, you can also restore these parameters to their default values.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```

enable

configure terminal

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

```

* Note:

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. To configure the LLDP transmission parameters, enter:

```
lldp [tx-interval|tx-hold-multiplier]
```

3. **(Optional)** To restore specific LLDP transmission parameters to their default values, enter:

```
default lldp [tx-interval|tx-hold-multiplier]
```

4. **(Optional)** To restore all LLDP transmission parameters to their default values, enter:

```
default lldp
```

Example

Configure the LLDP transmission interval. The LLDP port status is set to transmit and receive prior to the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 4/4
Switch:1(config-if)#lldp status txAndRx
Switch:1(config-if)#exit
Switch:1(config)#lldp tx-interval 31
```

Optionally, restore the LLDP transmission interval to its default value:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#default lldp tx-interval
```

Variable definitions

Use the information in the following table to help you understand the **lldp** command.

Variable	Value
tx-interval<5–32768>	Specifies the global LLDP transmit interval in seconds, that is, the interval in which LLDP frames are transmitted. The default is 30 seconds.
tx-hold-multiplier <2–10>	Configures the multiplier for the transmit interval used to compute the Time To Live (TTL) value in LLDP frames. The default is 4 seconds.

Viewing global LLDP statistics

Use this procedure to view and verify global LLDP statistics.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. To view LLDP statistics, enter:
`show lldp stats`
3. To view LLDP reception statistics, enter:
`show lldp rx-stats`
4. To view LLDP transmission statistics, enter:
`show lldp tx-stats`

5. (Optional) Clear global LLDP statistics:

```
clear lldp stats summary
```

Example**View LLDP statistics:**

```
Switch:1>enable
Switch:1#show lldp stats

=====
LLDP Stats
=====
Inserts    Deletes    Drops    Ageouts
-----
0          0          0        0
-----
```

View LLDP transmission statistics:

```
Switch:1#show lldp tx-stats

=====
LLDP Tx-Stats
=====
PORT NUM          FRAMES
-----
1/2                100
```

View LLDP reception statistics:

```
Switch:1#show lldp rx-stats

=====
LLDP Rx-Stats
=====
Port  Frames    Frames    Frames    TLVs    TLVs    AgeOuts
Num   Discarded  Errors   Total    Discarded  Unrecognized
-----
1/2   0          0         46       0         0         0
```

Viewing port-based LLDP statistics

Use this procedure to verify port-based LLDP statistics.

About this task

LLDP operates at the interface level. Enabling FA on a port automatically enables LLDP transmission and reception on the port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on *all* ports in the MLT.

*** Note:**

When FA is enabled on ports in an MLT or LACP MLT, tagging is enabled and spanning tree is disabled on those ports.

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. To verify successful LLDP transmission on a port, enter:

```
show lldp tx-stats port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

3. To verify that a port receives LLDP PDUs successfully, enter:

```
show lldp rx-stats port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

4. **(Optional)** To clear LLDP statistics on a port, or ports, enter:

```
clear lldp stats {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

Example

Verify LLDP transmission statistics on a port:

```
Switch:1>en
Switch:1#show lldp tx-stats port 1/2
=====
LLDP Tx-Stats
=====
PORT NUM          FRAMES
-----
1/2                100
```

Verify that the port is receiving LLDP PDUs:

```
Switch:1#show lldp rx-stats port 1/2
=====
LLDP Rx-Stats
=====
Port Num          Frames Discarded  Frames Errors  Frames Total  TLVs Discarded (Non FA)  TLVs Unsupported (Non FA)  AgeOuts
-----
1/2                0              0           0           46           0              0              0              0
```

Variable definitions

Use the data in the following table to use the `show lldp tx-stats` and the `show lldp rx-stats` commands.

Variable	Value
{slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Displaying learned LLDP neighbors

Use this procedure to verify details of the LLDP neighbors learned.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Verify details of LLDP neighbors learned:

```
show lldp neighbor
```

3. Verify details of LLDP neighbors learned on a specific port:

```
show lldp neighbor port {slot/port[/sub-port] [-slot/port[/sub-  
port]] [,...]}
```

Example

The following example shows how two switches—an FA Server and an FA Proxy discover each other as LLDP neighbors. Switch A, which is the FA Server is an Avaya VSP 7200 Series switch (model 7254XSQ) and switch B which is the proxy device is an Avaya ERS 4826GTS switch.

The following examples shows neighbor discovery on non-channelized and channelized ports (if your platform supports channelization).

On the non-channelized port 1/1 on the FA Server, verify neighbor discovery of the proxy switch.

```
SwitchA:1>en
SwitchA:1#show lldp neighbor
=====
                        LLDP Neighbor
=====
Port: 1/1      Index      : 1                Time: 1 day(s), 04:03:52
                ChassisId: MAC Address    70:30:18:5a:05:00
                PortId   : MAC Address    70:30:18:5a:05:07
                SysName  :
                SysCap   : Br / Br
                PortDescr: Port 7
                SysDescr : Ethernet Routing Switch 4826GTS HW:10 FW:5.8.0.1 SW:v6.9.2.027
=====
Total Neighbors : 1
=====
Capabilities Legend: (Supported/Enabled)
B= Bridge,      D= DOCSIS,      O= Other,      R= Repeater,
```

SPBM and IS-IS infrastructure configuration

S= Station, T= Telephone, W= WLAN, r= Router
Switch:1(config)#

On the proxy switch, verify discovery of the FA Server switch.

```
SwitchB:1>en
SwitchB:1#show lldp neighbor
-----
LLDP neighbor
-----
Port: 7      Index: 71      Time: 12 days, 21:40:30
ChassisId:  MAC address      a4:25:1b:52:70:00
PortId:     MAC address      a4:25:1b:52:70:04
SysName:    BEB1-7254XSQ
SysCap:     rB / rB          (Supported/Enabled)
PortDesc:   Avaya Virtual Services Platform 7254XSQ - Gbic1000BaseT Port
1/1
SysDescr:   VSP-7254XSQ (6.0.0.0_GA)
-----
Sys capability: O-Other; R-Repeater; B-Bridge; W-WLAN accesspoint; r-Router;
T-Telephone; D-DOCSIS cable device; S-Station only.
Total neighbors: 1
```

On the channelized port 1/1/1 on the FA Server switch, verify discovery of the proxy switch.

```
SwitchA:1>en
SwitchA:1#show lldp neighbor
=====
LLDP Neighbor
=====
Port: 1/1/1  Index   : 1      Time: 1 day(s), 04:03:52
ChassisId:  MAC Address      70:30:18:5a:05:00
PortId     : MAC Address      70:30:18:5a:05:07
SysName    :
SysCap     : Br / Br
PortDescr  : Port 7
SysDescr   : FA Proxy 4826GTS HW:10 FW:5.8.0.1 SW:v5.9.2.027
-----
Total Neighbors : 1
-----
Capabilities Legend: (Supported/Enabled)
B= Bridge, D= DOCSIS, O= Other, R= Repeater,
S= Station, T= Telephone, W= WLAN, r= Router
Switch:1(config)#
```

Verify neighbor discovery on the proxy switch.

```
SwitchB:1>en
SwitchB:1#show lldp neighbor
-----
LLDP neighbor
-----
Port: 7      Index: 71      Time: 12 days, 21:40:30
ChassisId:  MAC address      a4:25:1b:52:70:00
PortId:     MAC address      a4:25:1b:52:70:04
SysName:    BEB1-7254XSQ
SysCap:     rB / rB          (Supported/Enabled)
PortDesc:   Avaya Virtual Services Platform 7254XSQ - 40GbCR4-Channel Port
1/1/1
SysDescr:   VSP-7254XSQ (6.0.0.0_GA)
-----
Sys capability: O-Other; R-Repeater; B-Bridge; W-WLAN accesspoint; r-Router;
```

T-Telephone; D-DOCSIS cable device; S-Station only.
Total neighbors: 1

Variable definitions

Use the data in the following table to use the `show lldp neighbor` command.

Variable	Value
port {slot/port[/sub-port][-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. Displays LLDP neighbor information on the specified port.

Displaying Switched UNI (ELAN) I-SID information

Use this procedure to display information on FA-created Switched UNI (ELAN) I-SIDs.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display all Switched UNI (ELAN) I-SIDs:

```
show i-sid elan
```

3. Display ELAN I-SID information on an MLT:

```
show mlt i-sid [<1-512>]
```

* Note:

Viewing ELAN I-SID information on an MLT is useful to understand the origin of the I-SID when multiple client or proxy devices connecting to the FA Server using SMLT MLT advertise the *same* I-SID-to-VLAN mappings. In the event of a link failure on an MLT, the origin of the I-SID helps determine on which MLT, and thereby from which proxy or client device, the mappings were successfully learnt.

4. Display ELAN I-SID information on ports:

```
show interfaces gigabitEthernet i-sid [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]]
```

Example

Display information on all Switched UNI (ELAN) I-SIDs.

The following sample output displays, for example, the I-SID information on one of the peer switches of the FA Server, in a dual-homed SMLT configuration.

```
Switch:1>en
Switch:1#show isid elan

=====
Isid Info
=====
ISID      ISID      VLANID    PORT      MLT      ORIGIN
ID        TYPE                               INTERFACES INTERFACES
-----
2002      ELAN      N/A       c2002:1/10 -        DISC_LOCAL
4000      ELAN      N/A       -         c4000:1  DISC_BOTH
4001      ELAN      N/A       -         c4001:1  DISC_LOCAL
4030      ELAN      N/A       -         c4030:1  DISC_BOTH
4051      ELAN      N/A       -         c4051:1  DISC_BOTH
10200     ELAN      N/A       -         c200:1   DISC_REMOTE

c: customer vid      u: untagged-traffic

All 6 out of 6 Total Num of Elan i-sids displayed
```

*** Note:**

The I-SID TYPE field displays once for each I-SID. The I-SID TYPE of an I-SID that is either learned through FA mapping assignments or configured as an FA management I-SID, is always ELAN. If a platform VLAN has the same I-SID value as that of the I-SID in an FA mapping assignment or in an FA management I-SID configuration, then the platform VLAN is associated with the I-SID endpoint and appears in the VLANID column.

*** Note:**

- The ORIGIN field displays once for each I-SID. It indicates the origin of the I-SID and *not* the origin of the I-SID endpoint. To view the origin of the I-SID endpoints, execute either the `show mlt i-sid` or the `show interfaces gigabitEthernet i-sid` command.
 - The origin of I-SID 4000 displays as DISC_BOTH, because it is discovered on both v1ST peers.
 - The origin of I-SID 4001 displays as DISC_LOCAL because it is first discovered on the local FA Server switch.
 - The origin of I-SID 10200 displays as DISC_REMOTE because it is first discovered on the peer switch and then synchronized with the local switch.
- If the origin of an I-SID is DISC_LOCAL, DISC_REMOTE, DISC_BOTH or MANAGEMENT, it changes to CONFIG, after you manually configure an endpoint on the I-SID.

Display MLT I-SID information for MLT 1.

In this sample output, the ORIGIN field indicates the origin of the I-SID endpoint.

```
Switch:1#show mlt i-sid 1

=====
MLT Isid Info
=====
```

```

-----
MLTID  IFINDEX  ISID      VLANID  C-VID  ISID      ORIGIN      BPDU
      ID
-----
1      6144    4000     N/A     4000   ELAN      DISC_BOTH
1      6144    4001     N/A     4001   ELAN      DISC_LOCAL
1      6144    4030     N/A     4030   ELAN      DISC_BOTH
1      6144    4051     N/A     4051   ELAN      DISC_BOTH
1      6144    10200    N/A     200    ELAN      DISC_REMOTE
-----

```

5 out of 6 Total Num of i-sid endpoints displayed

Display I-SID information on the port 1/10:

In this sample output, the ORIGIN field indicates the origin of the I-SID endpoint.

```
Switch:1#show interfaces gigabitEthernet i-sid 1/10
```

```

=====
                                PORT Isid Info
=====
PORTNUM  IFINDEX  ISID      VLANID  C-VID  ISID      ORIGIN      BPDU
      ID
-----
1/10     201     2002     N/A     601    ELAN      DISC_LOCAL
-----

```

1 out of 6 Total Num of i-sid endpoints displayed

Variable definitions

Use the data in the following table to use the `show i-sid` command.

Variable	Value
elan	Displays all ELAN I-SIDs.

Use the data in the following table to use the `show mlt i-sid` command.

Variable	Value
<1-512>	The valid range for MLT ID.

Use the data in the following table to use the `show interfaces gigabitEthernet i-sid` command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Enabling or disabling FA Zero Touch Client Attachment

Use this procedure to enable or disable the global FA Zero Touch Client Attachment feature on an FA Proxy or Server. By default, FA Zero Touch Client Attachment support is enabled.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable an FA Zero Touch client:

```
fa zero-touch-client standard <camera|ona-sdn|ona-spb-over-ip|phone|
router|security-device|srvr-endpt|switch|video|virtual-switch|wap-
type1|wap-type2> i-sid <1-16000000>
```

3. Disable an FA Zero Touch client:

```
no fa zero-touch-client standard <camera|ona-sdn|ona-spb-over-ip|
phone|router|security-device|srvr-endpt|switch|video|virtual-switch|
wap-type1|wap-type2>
```

Example

```
Switch:1(config)# fa zero-touch-client standard camera i-sid 1003
```

```
Switch:1(config)# no fa zero-touch-client standard camera
```

Variable definitions

Use the data in the following table to use the **fa zero-touch-client standard** command.

Variable	Value
camera	Specify element type to match camera.
ona-sdn	Specify element type to match ona-sdn.
ona-spb-over-ip	Specify element type to match ona-spb-over-ip.
phone	Specify element type to match phone.
router	Specify element type to match router.
security-device	Specify element type to match security-device.
srvr-endpt	Specify element type to match srvr-endpt.
switch	Specify element type to match switch.
video	Specify element type to match video.
virtual-switch	Specify element type to match virtual-switch.
wap-type1	Specify element type to match wap-type1.
wap-type2	Specify element type to match wap-type2.

Displaying FA Zero Touch Client Attachment

Use this procedure to display the Zero Touch Client Attachment data you have configured on an FA Server.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display Zero Touch Client Attachment data:

```
show fa zero-touch-client
```

Example

The following example displays sample output for the **show fa zero-touch-client** command.

```
Switch:1#show fa zero-touch-client
```

```
=====
                        Fabric Attach Zero Touch Client
=====
Type      Description      I-SID      VLAN
-----
6         wap-type1         11111      123
11        camera           2000      200
17        ona-spb-over-ip  40001     4001
-----
3 out of 3 Total Num of Fabric Attach Zero Touch Client entries displayed
=====
```

IS-IS external metric configuration using the CLI

This section provides procedures for IS-IS external metric configuration.

Matching metric type for IS-IS routes

About this task

Use this procedure to match the external metric-type by using a route-map for any of the following cases:

- accepting a remote IS-IS route with the help of IS-IS accept policies.
- redistributing IS-IS routes into other protocols.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must log on to the route-map configuration mode in the CLI.

Procedure

1. Enter Route-Map Configuration mode:

```
enable
```

```
configure terminal
```

```
route-map WORD<1-64> <1-65535>
```

2. Match IS-IS metric type:

```
match metric-type-isis {any|internal|external}
```


3. Permit the route policy action:

```
permit
```

4. Enable the route policy

```
enable
```

Example

Match metric type for IS-IS routes:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
```

Match metric type for IS-IS routes in accept policies:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router isis
Switch:1(config-isis)# accept route-map rol
Switch:1(config-isis)# exit
Switch:1(config)# isis apply accept
```

Match metric type to redistribute IS-IS routes into some other protocol (OSPF,RIP,BGP)

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router bgp
Switch:1(router-bgp)# redistribute isis route-map rol
Switch:1(router-bgp)# exit
Switch:1(config)# ip bgp apply redistribute
```

Variable definitions

Use the data in the following table to use the match metric-type-isis command.

Variable	Value
metric-type-isis {any internal external}	<p>Specifies the IS-IS metric type.</p> <ul style="list-style-type: none"> • internal – permits or denies routes that are internal to the IS-IS domain. • external – permits or denies routes that originate from an external routing protocol domain. • any – permits or denies both internal routes as well as external routes.

Setting metric type for IS-IS routes

About this task

Use this procedure to set the IS-IS external metric-type by using a route-map for any of the following cases:

- accepting a remote IS-IS route with the help of IS-IS accept policies.
- redistributing routes from other protocols into IS-IS.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must log on to the route-map configuration mode in the CLI.

Procedure

1. Enter Route-Map Configuration mode:

```
enable
configure terminal
route-map WORD<1-64> <1-65535>
```

2. Set IS-IS metric type:

```
set metric-type-isis {any|internal|external}
```

3. Permit the route policy action:

```
permit
```

4. Enable the route policy

```
enable
```

Example

Set metric type for IS-IS routes:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# set metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
```

Set metric type for IS-IS routes in accept policies:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# set metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router isis
Switch:1(config-isis)# accept route-map rol
Switch:1(config-isis)# exit
Switch:1(config)# isis apply accept
```

Set metric type to redistribute routes from other protocols into IS-IS:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router isis
Switch:1(config-isis)# redistribute bgp route-map rol
Switch:1(config-isis)# exit
Switch:1(config)# isis apply redistribute
```

Variable definitions

Use the data in the following table to use the set metric-type-isis command.

Variable	Value
metric-type-isis {any internal external}	<p>Specifies the IS-IS metric type.</p> <ul style="list-style-type: none"> • internal – permits or denies routes that are internal to the IS-IS domain. • external – permits or denies routes that originate from an external routing protocol domain. • any – permits or denies both internal routes as well as external routes.

Setting metric type for IS-IS routes using global redistribute command

About this task

Use this procedure to set the IS-IS external metric-type using the global redistribute command for the following cases redistributing routes from other protocols into IS-IS.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must log on to the IS-IS router configuration mode in the CLI.

Procedure

1. Enter IS-IS Router Configuration mode:


```
enable
configure terminal
router isis
```
2. Set IS-IS metric type using global redistribute command:


```
redistribute direct metric-type {internal|external}
```
3. Enable the route policy


```
redistribute direct enable
```

Example

Set metric type for IS-IS routes using global redistribute command:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# router isis
Switch:1(config-isis)# redistribute direct metric-type internal
Switch:1(config-isis)# redistribute direct enable
```

Variable definitions

Use the data in the following table to use the `redistribute direct metric-type` command.

Variable	Value
metric-type {internal external}	<p>Specifies the IS-IS metric type.</p> <ul style="list-style-type: none"> • internal – permits or denies routes that are internal to the IS-IS domain. • external – permits or denies routes that originate from an external routing protocol domain.

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).

Important:

The **EnableSpbmConfigMode** boot flag must be enabled (default) before you can configure SPBM or IS-IS. To verify the setting, navigate to **Configuration > Edit > Chassis** and click on the **Boot Config** tab.

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.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
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 (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. In the navigation pane, expand the **Configuration > VLAN > VLANs** folders.
10. Click the **Basic** tab.
11. Click **Insert**.
12. In the **Type** field, click **spbm-bvlan**.
13. Click **Insert**.
14. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
15. Click the **SPBM** tab.
16. In the **Vlans** field, specify the IDs of the SPBM B-VLANs to add to the SPBM instance.
17. In the **PrimaryVlan** field, specify which of the SPBM B-VLANs specified in the previous step is the primary B-VLAN.
18. Click **Apply**.
19. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
20. Click the **Manual Area** tab.
21. In the Manual Area tab, click **Insert** to add a manual area (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, you should 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 **State** field, select **enable**.
33. Click **Insert** to enable the SPBM instance on the IS-IS circuit.
34. Under the IS-IS tab, click the **Interfaces** tab.
35. In the **AdminState** field for the IS-IS circuit entry, select **on** to enable the IS-IS circuit.
36. 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, Spanning Tree, and NLB*.

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. 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.
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 vIST peer. The default is primary.
SmltVirtualBmac	Specifies a virtual MAC address that can be used by both peers.

Table continues...

Name	Description
SmltPeerSysId	Specifies the system ID of the SPBM SMLT 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.
Ipv6Shortcut	Enables or disables SPBM IPv6 shortcut state. The default is disable.
McastSpbPimGwControllerEnable	Enables or disables ISIS PLSB Multicast SPB PIM Gateway controller. Disabled by default.
McastSpbPimGwGatewayEnable	Enables or disables ISIS PLSB Multicast SPB PIM Gateway. Disabled by default.


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

Name	Description
Type	Specifies the type of VLAN: <ul style="list-style-type: none"> • byPort • byProtocolId • 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>. 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. <p> Note:</p> <p>Although it is not strictly required for SPBM operation, you should 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.</p>

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
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. Wait up to 20 minutes for the LSPs with the original system ID to age out.

Note:

To check the age out time, use the `show isis lsdb sysid <original-sys-id>` command on any of the other SPB nodes in the network. When there is no output from this command, proceed to the next step. The time left (in seconds) for the LSPs to age out is shown under the column LIFETIME.

6. Disable IS-IS.
7. Change the nickname to the original nickname.
8. Enable IS-IS.

Configuring IP Multicast over Fabric Connect globally

Use this procedure to globally enable IP Multicast over Fabric Connect 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 Fabric Connect is disabled. There is no need to enable IP Multicast over Fabric Connect on the Backbone Core Bridges (BCBs).

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

! Important:

IP Multicast over Fabric Connect 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 Fabric Connect.

* Note:

You must enable IP multicast over Fabric Connect globally on all DvR enabled nodes (Controllers and Leaf nodes) in a DvR domain.

For more information on DvR, see *Configuring IPv4 Routing*.

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.

Procedure

1. Determine if any I-SIDs are within the default range reserved for multicast. In the navigation pane, 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. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
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 exist, click **Insert**.
7. In the **Mcast** box, select **enable** to enable IP Multicast over Fabric Connect globally.
8. Click **Insert**.
9. 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. 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 Fabric Connect is enabled. The default is disabled.
McastFwdCacheTimeout	Specifies the global forward cache timeout in seconds. The default is 210 seconds.

Modifying IP Multicast over Fabric Connect globally

Use this procedure to modify IP Multicast over Fabric Connect 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 Fabric Connect is disabled. There is no need to enable IP Multicast over Fabric Connect on the Backbone Core Bridges (BCBs).

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

Important:

IP Multicast over Fabric Connect 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 Fabric Connect.

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.

Procedure

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

Displaying IP Multicast over Fabric Connect routes

Use this procedure to display IP Multicast over Fabric Connect routes.

Procedure

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

IpMcastRoutes field descriptions

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

Name	Description
VsnIsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
Source	Specifies the IP address where the IP Multicast over Fabric Connect route originated.
NickName	Specifies the nick name used to filter criteria.
SourceBeb	Specifies the source BEB for the IP multicast route.
VlanId	Specifies the ID for the C-VLAN.
VrfName	Specifies the VRF name.
Datalsid	Specifies the data I-SID for the IP Multicast over Fabric Connect route. 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.
Type	Specifies the type for the IP Multicast over Fabric Connect route.
Bvlan	Specifies the B-VLAN for the IP Multicast over Fabric Connect route.
NniInterfaces	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 customer VLAN are user-to-network interface (UNI) ports.

Displaying the UNI ports for IP multicast routes

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

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
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 Fabric Connect route.
Source	Specifies the IP address where the IP Multicast over Fabric Connect route originated.
Vsnlsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.
Datalsid	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 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 **Configuration > IS-IS** folders.
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.
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	

Table continues...

Name	Description
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: <ul style="list-style-type: none"> • 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 pane, expand the **Configuration > IS-IS** folders.
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.
NickName	Indicates the nickname of the node where this I-SID was configured or discovered.

Table continues...

Name	Description
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.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
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:

- The assignment of primary and secondary roles to the vIST peers is automatic. The switch with the lower system ID (between the two vIST peers) is primary, and the switch with the higher system ID is secondary when default system-id values are being used.
- 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.

- When using the default hardware assigned system-id value, the SMLT Virtual B-MAC is automatically derived by comparing the system-id values of the two vIST peers. A value of 0x01 plus the lower of the two system-id values is used as the SMLT Virtual B-MAC.

When using a manually configured system-id value, the SMLT Virtual B-MAC must also be manually configured.

- An I-SID must be assigned to every VLAN that is a member of an L2 VSN. Also if an L2 VSN is created on one vIST Peer, it must also be created on the other vIST peer.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
2. Click the **SPBM** tab.
3. Use the **SmltSplitBEB** field to see whether the switch is the primary or secondary vIST 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 vIST 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 pane, expand the **Configuration > IS-IS** folders.
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 pane, expand the **Configuration > IS-IS** folders.
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. 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.
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 vIST peer. The default is primary.

Table continues...

Name	Description
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 IP multicast over SPBM is enabled. The default is disabled.
McastFwdCacheTimeout	Specifies the global forward cache timeout in seconds. The default is 210 seconds.
Ipv6Shortcut	Enables or disables SPBM IPv6 shortcut state. The default is disable.
McastSpbPimGwControllerEnable	Enables or disables ISIS PLSB Multicast SPB PIM Gateway controller. Disabled by default.
McastSpbPimGwGatewayEnable	Enables or disables ISIS PLSB Multicast SPB PIM Gateway. Disabled by default.

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 189.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
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 pane, expand the **Configuration > IS-IS** folders.
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.
SpbmId	Specifies the SPBM ID.
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. 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 pane, expand the **Configuration > IS-IS** folders.
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. 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)

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 pane, expand the **Configuration > IS-IS** folders.
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.
DestId	Specifies the destination ISID 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 of the IP unicast FIB entry
MetricType	Specifies the IP Metric Type of the IP unicast FIB entry.

Displaying the IPv6 unicast FIB

Use the following procedure to display the IPv6 unicast Forwarding Information Base (FIB). The tab shows IPv6 routes from remote Backbone Edge Bridges (BEBs)

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.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **SPBM**.
3. Click the **IPv6 Unicast FIB** tab.

IPv6 Unicast FIB field descriptions

Use the data in the following table to use the **IPv6 Unicast FIB** tab.

Name	Description
Vrfid	Specifies the VRF ID of the IPv6 unicast FIB entry, 0 indicates NRE.
DestinationIpAddrType	Specifies the address type of the destination IPv6 address.
DestinationIpAddr	Specifies the destination IPv6 Address of the IPv6 unicast FIB entry.
DestinationMask	Specifies the destination IPv6 mask of the IPv6 unicast FIB entry.
NextHopBmac	Specifies the nexthop B-MAC of the IPv6 unicast FIB entry.
DestIsid	Specifies the destination I-SID of the IPv6 unicast FIB entry.
Vlan	Specifies the VLAN of the IPv6 unicast FIB entry.
Isid	Specifies the I-SID of the IPv6 unicast FIB entry.
NextHopName	Specifies the nexthop hostname of the IPv6 unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the IPv6 unicast FIB entry.
SpbmCost	Specifies the B-MAC cost of the IPv6 unicast FIB entry.
PrefixCost	Specifies the prefix cost of the IPv6 unicast FIB entry.

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 pane, expand the **Configuration > IS-IS** folders.
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.

Table continues...

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 pane, expand the **Configuration > IS-IS** folders.
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.
HostName	Host name of the node where the multicast FIB entry originated.
OutgoingInterfaces	Specifies the switched UNI port outgoing interface of multicast FIB entry.
IncomingInterface	Specifies the incoming interface (port or MLT) of the multicast FIB entry.

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. In the navigation pane, expand the **Configuration > IS-IS** folders.
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 pane, expand the **Configuration > IS-IS** folders.
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
Interface	Specifies the IS-IS interface on which the adjacency is found.
Level	Indicates the level of the IS-IS interface (Level 1 [default] or Level 2).
State	Specifies the state of the adjacency: <ul style="list-style-type: none"> • 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 pane, expand the **Configuration > IS-IS** folders.
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: <ul style="list-style-type: none"> • level1: Level-1 router type • level1and2: Level-1/2 router type is not supported. 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>. <p>! Important:</p> 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, see Job Aid on page 90.
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. <p>The default value is 900 seconds.</p>
Csnplnt	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. <p>The default value is 10.</p>
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. <p>The default value is 5 seconds.</p>
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. <p>The default value is 2.</p>
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. <p>The default value is 100 milliseconds.</p>
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.

Table continues...

Name	Description
	By default, the system name comes from the host name configured at the system level.
IpSourceAddress	Specifies IP source address for SPBM IP shortcuts.
Ipv6SourceAddressType	Click ipv6 to use IPv6 addresses.
Ipv6SourceAddress	Specifies IPv6 source address for SPBM IP shortcuts.
IpTunnelSourceAddress	Specifies the IS-IS IP tunnel source address.
IpTunnelVrf	Specifies the VRF name associated with the IP tunnel.
MgmtIpAddr	Specifies the DvR management IP address for this node, in the DvR domain.
BackboneEnable	Select to enable this node to join the DvR backbone so that it can receive redistributed DvR host routes from all DvR Controllers in the network.

Configuring system-level IS-IS parameters

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

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
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: I1 or I2. Only I1 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.
SetOverload	Sets or clears the overload condition. The possible values are true or false.

Table continues...

Name	Description
	The default value is false.
SetOverloadUntil	<p>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.</p> <p>* Note:</p> <p>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.</p> <p>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.</p> <p>If SetOverload is false, the system clears the overload bit after SetOverloadUntil seconds elapse, if the system has not run out of memory.</p> <p>The default value is 20.</p>
MetricStyle	Specifies the IS-IS metric type. Available values are narrow, wide or both. Only wide is supported.

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, expand the **Configuration > IS-IS** folders.
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.

Table continues...

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 opened 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 pane, expand the **Configuration > IS-IS** folders.
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. 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: <ul style="list-style-type: none"> • 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. <p>The default is none.</p>
AuthKey	Specifies the authentication key.
KeyId	Specifies the authentication key ID.
LevelType	Specifies the router type globally: <ul style="list-style-type: none"> • level1: Level-1 router type • level 1and2: Level-1/2 router type. This type is not supported. <p>The default value is level1.</p>
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. In the navigation pane, expand the **Configuration > IS-IS** folders.

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: <ul style="list-style-type: none"> • I1: Level1 router type • I12: Level1/Level2 router type. This type is not supported. The default value is I1.
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. In the navigation pane, expand the **Configuration > IS-IS** folders.
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.
Level	Shows the type of circuit that discovered the interface counters. The point to point Hello PDU includes both L1 and L2, and IS from a single adjacency on point to point links, therefore combining counts on point to point links into one group.
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 **Configuration > IS-IS** folders.
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.
Direction	Indicates whether the switch is sending or receiving the PDUs.
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 **Configuration > IS-IS** folders.
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.

Table continues...

Name	Description
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.
Cumulative	Displays the total value since you opened 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 **Configuration > IS-IS** folders.
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 opened the Stats tab.

Table continues...

Name	Description
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. In the navigation pane, expand the **Configuration > IS-IS** folders.
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 opened 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.

Configuring an IS-IS Manual Area

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

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
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>. 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.

Fabric Extend configuration using EDM

The following sections provide procedural information you can use to configure Fabric Extend (FE) using Enterprise Device Manager (EDM).

Configuring Fabric Extend tunnels

Use the following procedure to configure Fabric Extend (FE) between a VSP 8000 in a Main office to a VSP 4000 in a Branch office. This is a typical deployment. However, if your deployment creates tunnels between two switches that support Fabric Extend natively (VSP 7200 Series and the VSP 8000 Series) then repeat the VSP 8000 steps and ignore the VSP 4000 steps.

Note:

VRF is an optional parameter. If a VRF is not configured, then FE uses the GRT.

Before you begin

The tunnel source IP address can be either a brouter port interface or a CLIP IP.

If using the tunnel originating address on the **GRT**, Fabric Extend has the following requirements:

- The tunnel source IP address must be on the GRT, not on a VRF.

*** Note:**

A Best Practice is to use separate IP addresses for the SPBM IP Shortcuts `ip-source-address` command and the Fabric Extend `ip-tunnel-source-address` command. However, if you want these IP addresses to be the same, you **MUST** exclude the `ip-source-address` address with an IS-IS accept policy. You cannot use the redistribute command with a route map exclusion.

Specify a CLIP interface to use as the source address for SPBM IP shortcuts.

- If IP Shortcuts is enabled, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

If you are using the tunnel originating address on a **VRF**, Fabric Extend has the following requirements:

- Configure a CLIP and tunnel source IP address on the VRF.
- Remote management of the VSP 4000 is only possible after establishing IP Shortcut over IS-IS. (Alternatively, you can enable GRT-VRF redistribution locally.)

About this task

! Important:

In this procedure, the Switch A steps are for platforms that support Fabric Extend natively. The Switch B steps are for the 1 Gbps platforms that require an ONA to support Fabric Extend.

Configuring Fabric Extend consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Note that the VSP 4000 source address command is different than the VSP 7200/VSP 8000 Series command. Also note that the logical interface commands are different between Layer 2 and Layer 3 networks.

Procedure

VSP 8000 steps

1. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
2. Click the Globals tab.
3. In the **IpTunnelSourceAddress** field, enter the IP tunnel source address.
4. If you are using a VRF, select its name from the drop down menu in the **IpTunnelVrf** field.
5. Click **Apply**.

VSP 4000 steps

*** Note:**

The interface VLAN connecting to the ONA network port is always in the GRT and the member port that the VLAN is part of is always an access port.

6. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
7. Click the Globals tab.

8. In the **IpTunnelSourceAddress** field, enter the IP tunnel source address.
9. In the **IpTunnelPort** field, select from the drop down menu the physical port that the logical interface is connected to in an L2 network.
10. If you are using a VRF, select its name from the drop down menu in the **IpTunnelVrf** field.
11. In the **IpTunnelMtu** field, enter a value between 750 and 1950 to specify the size of the maximum transmission unit (MTU). The default is 1950. This parameter only applies to an ONA configuration.
12. Click **Apply**.

Fabric Extend field descriptions

Use the data in one of the following tables to use the **Globals** command, depending on whether you are configuring a VSP 4000 or a VSP 8000 switch.

Use the data in the following table to use the tab to configure a Fabric Extend tunnel source address.

Table 7: VSP 4000

Name	Description
IpTunnelSourceAddress	Specifies the IS-IS IPv4 tunnel source address.
IpTunnelPort	Specifies the physical port that the logical interface is connected to in an L2 network. The parameter is for the VSP 4000 only.
IpTunnelVrf	Specifies the VRF name associated with the IP tunnel.
IpTunnelMtu	Specifies the size of the maximum transmission unit (MTU). The default is 1950. This parameter only applies to an ONA configuration.

Table 8: VSP 8000

Name	Description
IpTunnelSourceAddress	Specifies the IS-IS IPv4 tunnel source address.
IpTunnelVrf	Specifies the VRF name associated with the IP tunnel.

Configuring Fabric Extend logical interfaces

Use the following procedure to configure Fabric Extend (FE) between a VSP 8000 in a Main office to a VSP 4000 in a Branch office. This is a typical deployment. However, if your deployment creates tunnels between two switches that support Fabric Extend natively (VSP 7200 Series and the VSP 8000 Series) then repeat the VSP 8000 steps and ignore the VSP 4000 steps.

*** Note:**

VRF is an optional parameter. If a VRF is not configured, then FE uses the GRT.

About this task

! Important:

In this procedure, the Switch A steps are for platforms that support Fabric Extend natively. The Switch B steps are for the 1 Gbps platforms that require an ONA to support Fabric Extend.

Configuring Fabric Extend consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Note that the VSP 4000 source address command is different than the VSP 7200/VSP 8000 Series command. Also note that the logical interface commands are different between Layer 2 and Layer 3 networks.

Procedure

VSP 8000 steps

1. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
2. Click the Logical Interfaces tab.
3. In the **Id** field, enter the index number that uniquely identifies this logical interface.
4. In the **Name** field, enter the name of this logical interface.
5. In the **Type** field, select the type of core network that the tunnel will be traversing. If it's a Layer 2 Core, select **layer2**. If it's a Layer 3 Core, select **ip**.

*** Note:**

Different fields will be available depending on which type of core network you select.

6. For a Layer 2 Core, complete the following fields:
 - a. In the **DestIfIndex** field, click the ellipsis button (...) to select the physical port that the logical interface is connected to or enter the name of the MLT.
 - b. In the **Vids** field, enter the list of VLANs for this logical interface.
 - c. In the **PrimaryVid** field, enter the primary tunnel VLAN ID.

*** Note:**

The primary VLAN ID must be one of the VIDs listed in the **Vids** field.

7. For a Layer 3 Core, complete the following field:

In the **DestIPAddr** field, enter the destination IP address for the logical interface.

8. Click **Insert**.

VSP 4000 steps

*** Note:**

The interface VLAN connecting to the ONA network port is always in the GRT and the member port that the VLAN is part of is always an access port.

9. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
10. Click the Logical Interfaces tab.
11. In the **Id** field, enter the index number that uniquely identifies this logical interface.
12. In the **Name** field, enter the name of this logical interface.
13. In the **Type** field, select the type of core network that the tunnel will be traversing. If it's a Layer 2 Core, select **layer2**. If it's a Layer 3 Core, select **ip**.

*** Note:**

Different fields will be available depending on which type of core network you select.

14. For a Layer 2 Core, complete the following fields:
 - a. In the **DestIfIndex** field, click the ellipsis button (...) to select the physical port that the logical interface is connected to or enter the name of the MLT.
 - b. In the **Vids** field, enter the list of VLANs for this logical interface.
 - c. In the **PrimaryVid** field, enter the primary tunnel VLAN ID.

*** Note:**

The primary VLAN ID must be one of the VIDs listed in the **Vids** field.

15. For a Layer 3 Core, complete the following field:

In the **DestIPAddr** field, enter the destination IP address for the logical interface.
16. Click **Insert**.

Fabric Extend logical interfaces field descriptions

Use the data in one of the following tables to use the **Logical Interfaces** command. The available fields are different depending on what type of core you select: **layer 2** or **ip**.

Table 9: Layer 2 core

Name	Description
Id	Specifies the index number that uniquely identifies this logical interface. This value is a read-only value.
Name	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
Type	Specify layer 2 for a Layer 2 core network that the tunnel will be traversing.
DestIfIndex	Specifies the physical port or MLT that the logical interface is connected to.

Table continues...

Name	Description
Vids	Specifies the list of VLANs that are associated with this logical interface.
PrimaryVid	Specifies the primary tunnel VLAN ID associated with this L2 IS-IS logical interface.

Table 10: Layer 3 core

Name	Description
Id	Specifies the index number that uniquely identifies this logical interface. This value is a read-only value.
Name	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
Type	Specify ip for a Layer 3 core network that the tunnel will be traversing.
DestIPAddr	Specifies the destination IP address for the logical interface.

Displaying the logical interface next hop

Use the following procedure to display the next hop for the logical interface.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
2. Click the Logical Interfaces NextHop tab.

Logical interfaces next hop field descriptions

Use the data in one of the following tables to view the next hop for the logical Interface.

Name	Description
Id	Shows a unique value that identifies the logical interface tunnel.
Ip	Shows a unique value that identifies the next hop IP address of the logical interface tunnel.
DestIfIndex	Shows the next hop destination interface index to reach the next hop IP of the logical interface tunnel.
DestVid	Shows the next hop destination VLAN ID to reach the next hop IP of the logical interface tunnel.

Fabric Attach configuration using the EDM

The following sections provide procedural information you can use to configure Fabric Attach (FA) and Logical Link Discovery Protocol (LLDP) using Enterprise Device Manager (EDM).

Configuring Fabric Attach globally

Use this procedure to configure FA globally or view existing FA global configuration.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Fabric Attach** folders.
2. In the content pane, click the **Globals** tab.
3. To enable or disable the Fabric Attach service, click **enabled** or **disabled** in the **Service** field.

 **Caution:**

Disabling FA flushes all FA element discovery and mappings.

4. View the element type in the **ElementType** field.

 **Note:**

The only supported element type is **faServer** (FA Server).

5. To specify the assignment time-out, enter a time-out value in seconds in the **AsgnTimeout** field.
6. View the provision mode in the **ProvisionMode** field.

 **Note:**

The supported provision mode is **spbm**.

7. To specify the discovery time-out, enter a time-out value in seconds in the **DiscTimeout** field.
8. To clear the FA statistics, select the **Clear FA Statistics** checkbox.
9. To clear the error counters, select the check boxes **ClearErrorCounters** and/or **ClearGlobalErrorCounters**.
10. Click **Apply**.

Fabric Attach Globals field descriptions

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

Name	Description
Service	Enables or disables Fabric Attach service globally. The default is enable.
ElementType	Specifies the Fabric Attach element type. The supported element type is Fabric Attach Server.
AsgnTimeout	Specifies the Fabric Attach assignment time-out in seconds. The range is 45 to 480 seconds. The default is 240 seconds.

Table continues...

Name	Description
ProvisionMode	Specifies the Fabric Attach provision mode. The supported provision mode is SPB.
DiscTimeout	Specifies the Fabric Attach discovery time-out in seconds. The range is 45 to 480 seconds. The default is 240 seconds.
Clear FA Statistics	Clears Fabric Attach statistics.
ClearGlobalErrorCounters	Clears Fabric Attach global error counters. Disabled by default.

Configuring Fabric Attach I-SID-to-VLAN assignments

Use this procedure to view or configure FA I-SID-to-VLAN assignment information.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Fabric Attach** folders.
2. Click the **Assignment** tab.
3. If you make configuration changes, click **Apply** to save changes.

Fabric Attach I-SID-to-VLAN assignments field descriptions

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

Name	Description
IfIndex	Specifies the interface identifier of the I-SID-to-VLAN assignment.
Isid	Specifies the I-SID value of the I-SID-to-VLAN assignment.
Vlan	Specifies the VLAN ID component of the I-SID-to-VLAN assignment.
State	Specifies the current state of the I-SID-to-VLAN assignment. It can be one of the following values: <ul style="list-style-type: none"> • Other • Pending • Active • Rejected
Origin	Specifies the origin information of the I-SID-to-VLAN assignment.

Configuring Fabric Attach interface-level settings

Use this procedure to configure FA interface-level settings or view existing interface-level settings.

You can enable Fabric Attach on a port, static MLT or an LACP MLT. Enabling FA on a port not only enables tagging but also disables spanning tree on that port. Enabling FA on an MLT enables FA on all ports of the MLT. When FA is enabled on ports in an MLT or LACP MLT, tagging is enabled and spanning tree is disabled on all those ports.

Before you begin

Ensure that FA is enabled globally on the switch.

About this task

Enabling FA on a port or MLT is necessary for element discovery. On the FA Server, FA is enabled globally by default. However, you must explicitly enable FA on a desired port or MLT interface, following which the FA Server can begin transmitting LLDP PDUs that contain the element discovery TLVs. This information is received by FA Client and FA Proxy devices which in turn also transmit their FA capabilities and settings. After the element handshake completes, the FA Server receives I-SID-to-VLAN assignment mappings from the connected client or proxy devices, on that port or MLT.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Fabric Attach** folders.
2. Click the **Ports** tab.

The FA interface-level settings are displayed.

3. To modify existing settings, double-click on the fields on this window. After making the required changes, click **Apply** to save your changes.
4. To configure FA on a new port or MLT interface:

- a. Click **Insert**.

The **Insert Ports** dialog box appears.

- b. To configure FA on a port, enter a port number in the format slot/port[/sub-port], or click **Port** to select from a list of available ports.
- c. To configure FA on an MLT, enter an MLT ID or click **Mlt** to select from a list of configured MLTs.

 **Note:**

FA is successfully enabled on the MLT, only if all ports of the MLT have FA successfully enabled. Enabling FA enables LLDP on all ports. Tagging is enabled and spanning tree is disabled.

- d. Click **Insert** to save your changes.
5. To remove (delete) FA on a port or MLT:
 - a. In the content pane, select a port or MLT from the list.
 - b. Click **Delete**.

 **Caution:**

Removing FA on an interface flushes all FA element discovery and I-SID-to-VLAN mappings associated with that interface.

Fabric Attach Ports field descriptions

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

Name	Description
IfIndex	Specifies the interface (port or MLT) on which Fabric Attach is configured.
State	Specifies the current state of the Fabric Attach port. It is either enabled or disabled. This field indicates whether LLDP PDUs (that include FA TLVs) are generated on the port (enabled) or not (disabled).
MsgAuthStatus	Specifies the Fabric Attach message authentication status on the port. It is either enabled or disabled.
MsgAuthKey	Specifies the Fabric Attach message authentication key for the associated port. The maximum length of this key is 32 characters.
MgmtIsid	Specifies the Fabric Attach management I-SID for the associated port. The range is 0 to 16777215. A zero value indicates that the management I-SID is not specified for the interface.
MgmtCvid	Specifies the Fabric Attach management customer VLAN ID (C-VID) for the interface. A zero value indicates that no C-VID is specified for the interface. A value of 4096 indicates the port is untagged.

Viewing Fabric Attach discovered elements

Use this procedure to view discovered Fabric Attach elements.

About this task

When FA is enabled on an FA Server switch, LLDP PDUs are exchanged between the FA Server and FA Clients or Proxies. Standard LLDPs allow neighbors to be learned. In addition, organizational specific element discovery TLVs allow the Client or Proxy to recognize that it has attached to an FA Server. Only after the discovery handshake is complete, an FA Client or Proxy can transmit I-SID-to-VLAN assignments to join the SPB Fabric through the FA Server.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Fabric Attach** folders.
2. In the content pane, click the **Elements** tab.

Fabric Attach Elements field descriptions

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

Name	Description
IfIndex	Specifies the interface (port or MLT) at which the Fabric Attach element was discovered.
ElementType	Specifies the element type of the discovered Fabric Attach element, as advertised using LLDP.

Table continues...

Name	Description
	The supported element type is the Fabric Attach Server.
ElementVlan	Specifies the VLAN ID of the discovered Fabric Attach element, as advertised using LLDP.
ElementId	Specifies the system ID of the discovered Fabric Attach element, as advertised using LLDP.
ElementState	Specifies the state flag data associated with the discovered Fabric Attach element, as advertised using LLDP.
ElementOperAuthStatus	Specifies the authentication status of the discovered Fabric Attach element.
ElementAsgnsOperAuthStatus	Specifies the authentication status of remote assignments.
ElementAuth	Specifies the discovered element authentication status.
AsgnsAuth	Specifies the assignment authentication status.

Viewing FA statistics

Use this procedure to view FA statistics.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Fabric Attach** folders.
2. In the content pane, click the **Stats** tab.

Fabric Attach Statistics field descriptions

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

Name	Description
PortIndex	Specifies the port for which the FA statistics are displayed.
DiscElemReceived	Specifies the number of element discoveries received on the port.
AsgnReceived	Specifies the number of remote assignments received on the port.
AsgnAccepted	Specifies the number of remote assignments accepted on the port.
AsgnRejected	Specifies the number of remote assignments rejected on the port.
AsgnExpired	Specifies the number of remote assignments that have expired, on the port.
AuthFailed	Specifies the number of authentications that have failed on the port.
DiscElemExpired	Specifies the number of discovery elements that have expired on the port.
DiscElemDeleted	Specifies the number of discovery elements that are deleted on the port.
AsgnDeleted	Specifies the number of remote assignments deleted on the port.

Configuring LLDP global information

Use this procedure to configure or view LLDP global information.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > 802_1ab.LLDP** folders.
2. In the content pane, click the **Globals** tab.
3. After you make the required configuration changes, click **Apply** to save changes.

Globals field descriptions

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

Field	Description
IldpMessageTxInterval	Specifies the interval at which LLDP messages are transmitted. The default is 30 seconds.
IldpMessageTxHoldMultiplier	Specifies the multiplier used to calculate the time-to-live (TTL) value of an LLDP message. The default value is 4 seconds.
IldpReinitDelay	Specifies the delay in seconds between the time a port is disabled and the time it is re-initialized. The default is 1 second.
IldpTxDelay	Specifies the delay in seconds between successive LLDP transmissions. The default is 1 second. The recommended value is as follows: $1 < \text{IldpTxDelay} < (0.25 \times \text{IldpMessageTxInterval})$
IldpNotificationInterval	Specifies the time interval between successive LLDP notifications. It controls the transmission of notifications. The default is 5 seconds.
Stats	
RemTablesLastChangeTime	Specifies the timestamp of LLDP missed notification events on a port, for example, due to transmission loss.
RemTablesInserts	Specifies the number of times the information advertised by a MAC Service Access Point (MSAP) is inserted into the respective tables.
RemTablesDeletes	Specifies the number of times the information advertised by an MSAP is deleted from the respective tables.
RemTablesDrops	Specifies the number of times the information advertised by an MSAP was not entered into the respective tables.
RemTablesAgeouts	Specifies the number of times the information advertised by an MSAP was deleted from the respective tables.

Viewing the LLDP port information

Use this procedure to view the LLDP port information.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > 802_1ab.LLDP** folders.
2. In the content pane, click the **Port** tab.
3. View the administrative status of the port in the **AdminStatus** field. To modify, double-click on a cell and select a value from the drop-down list.
4. View whether the port is enabled for notifications in the **NotificationEnable** field. To modify, double-click on a cell and select a value from the drop-down list.
5. View the set of TLVs whose transmission using LLDP is always allowed by network management in the **TLVsTxEnable** field.
6. **(Optional)** Modify the TLVs as follows:
 - a. To enable a TLV, select the appropriate check box, and click **Ok**. You can select more than one check box.
 - b. To enable all TLVs, click **Select All**, and click **Ok**.
 - c. To disable all TLVs, click **Disable All**, and click **Ok**.
7. View the CDP administrative status in the **CdpAdminState** field. To modify, double-click on a cell and select a value from the drop-down list.
8. Click **Apply** to save any configuration changes.
9. Click **Refresh** to verify the configuration.

Port field descriptions

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

Name	Description
PortNum	Specifies the port number. This is a read-only cell.
AdminStatus	<p>Specifies the administrative status of the port. The options are:</p> <ul style="list-style-type: none"> • txOnly: LLDP frames are only transmitted on this port. • rxOnly: LLDP frames are only received on this port. • txAndRx: LLDP frames are transmitted and received on this port. • disabled: LLDP frames are neither transmitted or received on this port. Any information received on this port from remote systems before this is disabled, ages out. <p>The default is disabled.</p>
NotificationEnable	<p>Specifies whether the port is enabled or disabled for notifications.</p> <ul style="list-style-type: none"> • true: indicates that the notifications are enabled. • false: indicates that the notifications are disabled. <p>The default is false.</p>

Table continues...

Name	Description
TLVsTxEnable	<p>Specifies the set of TLVs whose transmission using LLDP is always allowed by network management.</p> <p>The following list describes the TLV types:</p> <ul style="list-style-type: none"> • portDesc — indicates that the Port Description TLV is transmitted. • sysName — indicates that the System Name TLV. is transmitted. • sysDesc — indicates that the System Description TLV. is transmitted. • sysCap — indicates that the System Capabilities TLV. is transmitted. <p>The default is an empty set of TLVs.</p>
CdpAdminState	<p>Specifies the CDP administrative status of the port.</p> <ul style="list-style-type: none"> • true: indicates CDP is enabled. • false: indicates CDP is disabled. <p>The default is false.</p> <p>If CDP is enabled, the interface accepts only CDP packets. Similarly, if CDP is disabled but LLDP is enabled, the interface accepts only LLDP packets. To switch a port from CDP mode to LLDP mode, the LLDP status on that port must be txAndRx.</p>

Viewing LLDP transmission statistics

Use this procedure to view the LLDP transmission statistics. You can also view the statistics graphically.

About this task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on *all* ports in that MLT.

Note:

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Note:

If a slot is removed from the switch chassis, the statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > 802_1ab.LLDP** folders.

2. In the content pane, click the **TX Stats** tab.

The transmission statistics are displayed.

3. To view the transmission statistics graphically for a port:

- a. In the content pane (on the right-hand-side), select a row and click the **Graph** button.

The **TX Stats-Graph, <port-number>** tab displays.

You can view a graphical representation of the LLDP frames transmitted (**FramesTotal**), for the following parameters:

- AbsoluteValue
 - Cumulative
 - Average/sec
 - Minimum/sec
 - Maximum/sec
 - LastVal/sec
- b. To view the graph, select one of the above parameters and click the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
 - c. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.
 - d. Click **Export**, to export the statistical data to a file.
 - e. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

TX Stats field descriptions

Use the data in the following table to view the LLDP transmission statistics.

Field descriptions for the **TX Stats** tab.

Name	Description
PortNum	Specifies the port number.
FramesTotal	Specifies the total number of LLDP frames transmitted.

Field descriptions for the **TX Stats-Graph, <port-number>** tab.

Name	Description
AbsoluteValue	Specifies the absolute number of LLDP frames at a given point in time.
Cumulative	Specifies the cumulative rate of change of LLDP frames transmitted.

Table continues...

Name	Description
Average/sec	Specifies the average rate of change of LLDP frames transmitted.
Minimum/sec	Specifies the minimum rate of change of LLDP frames transmitted.
Maximum/sec	Specifies the maximum rate of change of LLDP frames transmitted.
LastVal/sec	Specifies the rate of change of LLDP frames transmitted in the last second.

Viewing LLDP reception statistics

Use this procedure to view the LLDP reception statistics. You can also view these statistics graphically.

About this task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on *all* ports in that MLT.

* Note:

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

* Note:

If a slot is removed from the switch chassis, the statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > 802_1ab.LLDP** folders.
2. In the content pane, click the **RX Stats** tab.
3. To view the reception statistics graphically for a port:
 - a. Select a row and click **Graph**.

The **RX Stats-Graph,<port-number>** tab displays.

You can view a graphical representation of the following data:

- **FramesDiscardedTotal** — Total number of LLDP received frames that were discarded.
- **FramesErrors** — Total number of erroneous LLDP frames received.
- **FramesTotal** — Total number of frames received.

- **TLVsDiscardedTotal** — Total number of received TLVs that were discarded.
 - **TLVsUnrecognizedTotal** — Total number of unrecognized TLVs received.
- b. Select one of the above parameters and click the appropriate icon on the top left-hand-side corner of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
 - c. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.
 - d. Click **Export**, to export the statistical data to a file.
 - e. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

RX Stats field descriptions

Use the data in the following table to view the LLDP reception statistics.

Name	Description
PortNum	Specifies the port number.
FramesDiscardedTotal	Specifies the number of LLDP frames received on the port, but discarded, for any reason. This counter provides an indication of possible LLDP header formatting problems in the sending system, or LLDP PDU validation problems in the receiving system.
FramesErrors	Specifies the number of invalid LLDP frames received on the port.
FramesTotal	Specifies the total number of LLDP frames received on the port.
TLVsDiscardedTotal	Specifies the number of LLDP TLVs discarded on the port, for any reason.
TLVsUnrecognizedTotal	Specifies the number of LLDP TLVs on the port, that are unrecognized on that port. An unrecognized TLV is referred to as the TLV whose type value is in the range of reserved TLV types (000 1001–111 1110). An unrecognized TLV could be, for example, a basic management TLV from a later LLDP version.
AgeoutsTotal	Specifies the number of LLDP age-outs that occur on a specific port. An age-out is the number of times the complete set of information advertised by a particular MSAP is deleted, because the information timeliness interval has expired.

Field descriptions for the **RX Stats-Graph, <port-number>** tab.

Name	Description
AbsoluteValue	Specifies the absolute number of LLDP frames at a given point in time.
Cumulative	Specifies the cumulative rate of change of LLDP frames received.
Average/sec	Specifies the average rate of change of LLDP frames received.
Minimum/sec	Specifies the minimum rate of change of LLDP frames received.
Maximum/sec	Specifies the maximum rate of change of LLDP frames received.
LastVal/sec	Specifies the rate of change of LLDP frames received in the last second.

Viewing LLDP local system information

Use this procedure to view the LLDP local system information.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > 802_1ab.LLDP** folders.
2. In the content pane, click the **Local System** tab.

Local System field descriptions

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

Name	Description
ChassisIdSubType	Indicates the encoding used to identify the local system chassis. <ul style="list-style-type: none"> • chassisComponent • interfaceAlias • portComponent • macAddress • networkAddress • interfaceName • local
ChassisId	Indicates the chassis ID of the local system.
SysName	Indicates local system name.
SysDesc	Indicates local system description.
SysCapSupported	Indicates the system capabilities supported on the local system.
SysCapEnabled	Indicates the system capabilities that are enabled on the local system.

Viewing LLDP local port information

Use this procedure to view the LLDP local port information.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > 802_1ab.LLDP** folders.
2. In the content pane, click the **Local Port** tab.

Local port field descriptions

Use the data in the following table to use the **Local Port** tab.

Name	Description
PortNum	Indicates the port number.

Table continues...

Name	Description
PortIdSubType	Indicates the type of port identifier. <ul style="list-style-type: none"> • interfaceAlias • portComponent • macAddress • networkAddress • interfaceName • agentCircuitId • local
PortId	Indicates the identifier associated with the port, on the local system.
PortDesc	Indicates the description of the port, on the local system.

Viewing LLDP neighbor information

Use this procedure to view the LLDP neighbor information.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > 802_1ab.LLDP** folders.
2. In the content pane, click the **Neighbor** tab.

Neighbor field descriptions

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

Name	Description
TimeMark	Indicates the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Identifies the port on which the remote system information is received.
Index	Indicates a particular connection instance that is unique to the remote system.
ChassisIdSubtype	Indicates the type of encoding used to identify the remote system chassis. <ul style="list-style-type: none"> • chassisComponent • interfaceAlias • portComponent • macAddress • networkAddress • interfaceName • local

Table continues...

Name	Description
ChassisId	Indicates the chassis ID of the remote system.
SysCapSupported	Identifies the system capabilities supported on the remote system.
SysCapEnabled	Identifies the system capabilities enabled on the remote system.
SysName	Indicates the name of the remote system.
SysDesc	Indicates the description of the remote system.
PortIdSubType	Indicates the type of encoding used to identify the remote port.
PortId	Indicates the remote port ID.
PortDesc	Indicates the remote port description.
ProtocolType	Indicates whether the entry protocol is CDP or LLDP.
IpAddress	Indicates the neighbor's IP address.

Viewing global FA statistics graphically

Use this procedure to view the global FA statistics graphically.

Procedure

1. In the navigation pane, expand the **Graph > Chassis** folders.
2. Click the **Fabric Attach** tab.
The global FA statistics are displayed.
3. To view a graphical representation of the statistics, select a row and click the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
4. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.
5. Click **Export**, to export the statistical data to a file.
6. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

Fabric Attach field descriptions

Use the data in the following table to use the **Fabric Attach** tab.

Name	Description
DiscElemReceived	Specifies the number of discovery elements received globally.
AsgnReceived	Specifies the number of remote I-SID-to-VLAN assignments received globally.
AsgnAccepted	Specifies the number of remote I-SID-to-VLAN assignments accepted globally.
AsgnRejected	Specifies the number of remote I-SID-to-VLAN assignments rejected globally.

Table continues...

Name	Description
AsgnExpired	Specifies the number of remote I-SID-to-VLAN assignments that expired globally.
DiscAuthFailed	Specifies the number of discovery authentications that failed globally.
DiscElemExpired	Specifies the number of discovery elements that expired globally.
DiscElemDeleted	Specifies the number of discovery elements that were deleted globally.
AsgnDeleted	Specifies the number of remote assignments that were deleted globally.
AsgnAuthFailed	Specifies the number of assignment authentications that failed globally.

Viewing FA port statistics graphically

Use this procedure to view the FA port statistics graphically.

Before you begin

Ensure that a switch port is selected in the **Device Physical View** tab.

Procedure

1. In the navigation pane, expand the **Graph > Port** folders.
2. Click the **Fabric Attach** tab.
The FA port statistics are displayed.
3. To view a graphical representation of the port statistics, select a row and click the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
4. Click **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.
5. Click **Export**, to export the statistical data to a file.
6. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

Fabric Attach field descriptions

Use the data in the following table to use the **Fabric Attach** tab.

Name	Description
DiscElemReceived	Specifies the number of discovery elements received on a given port.
AsgnReceived	Specifies the number of remote I-SID-to-VLAN assignments received on a given port.
AsgnAccepted	Specifies the number of remote I-SID-to-VLAN assignments accepted on a given port.
AsgnRejected	Specifies the number of remote I-SID-to-VLAN assignments rejected on a given port.
AsgnExpired	Specifies the number of remote I-SID-to-VLAN assignments that expired on a given port.

Table continues...

Name	Description
DiscAuthFailed	Specifies the number of authentications that failed on a given port.
DiscElemExpired	Specifies the number of discovery elements that expired on a given port.
DiscElemDeleted	Specifies the number of discovery elements that were deleted on a given port.
AsgnDeleted	Specifies the number of remote assignments that were deleted on a given port.
AsgnAuthFailed	Specifies the number of assignment authentications that failed on a given port.

Inserting a Zero Touch Client

Use this procedure to insert a FA Zero Touch Client.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **Fabric Attach**.
3. Click the **Zero Touch Client Auto Attach** tab.
4. Click **Insert**.

The **Insert Zero Touch Client** dialog box appears.

5. In the **Type** field click the ellipsis and select a client. Click **Ok** to select the client or **Refresh** to update the list.
6. In the ISID field enter the ISID value.
The ISID value is between 0 and 16777214.
7. Click **Insert**.

Configuring FA Zero Touch Client auto attach

Use this procedure to configure FA Zero Touch Client auto attach.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **Fabric Attach**
3. Click the **Zero Touch Client Auto Attach** tab.

From the Zero Touch Client Auto Attach tab you can configure a number of auto attach settings.

4. Click **Insert**.
5. In the **Type** field click the ellipsis and select a client.
6. Click **Ok** to select the client or **Refresh** to update the list.

7. In the **ISID** field enter the ISID value.
8. Click **Insert**.
9. **(Optional)** To **Delete** a FA Zero Touch client select it from the auto attach table and click **Delete**.

Zero Touch Client Auto Attach Field Descriptions

Use the data in the following table to use the Zero Touch Client Auto Attach tab

Field	Description
Type	This column describes the type of client assigned to auto attach. Available FA client types are: <ul style="list-style-type: none"> • Wireless AP (Type 1) • Wireless AP (Type 2) • Switch • Router • IP Phone • IP Camera • IP Video • Security Device • Virtual Switch • Server Endpoint • ONA (SDN) • ONA (spb0IP)
ClientName	FA client name.
Isid	Specifies the I-SID value of the I-SID-to-VLAN assignment.

SPBM configuration examples


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

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Basic SPBM configuration example

The following figure shows a sample greenfield deployment for SPBM.

Figure 24: Greenfield SPBM deployment



*** Note:**

For migration purposes, SPBM can coexist with existing SMLT configurations.

Ethernet and MLT configuration

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

SwitchC

```
PORT CONFIGURATION - PHASE 1

interface GigabitEthernet 1/30
encapsulation dot1q
exit
```

SwitchG

```
PORT CONFIGURATION - PHASE 1

interface GigabitEthernet 1/5
encapsulation dot1q
exit

MLT CONFIGURATION

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

SwitchD

```
MLT CONFIGURATION

mlt 1 enable
```

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

IS-IS SPBM global configuration

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




Figure 25: IS-IS SPBM global

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

SwitchC

```
enable
configure terminal
prompt SwitchC

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.13
spbm 1 b-vid 20

ISIS CONFIGURATION

is-type 11
manual-area 30.0000
sys-name SwitchC
exit
router isis enable

VLAN CONFIGURATION

vlan create 20 name "B-VLAN" type spbm-bvlan
```

SwitchG

```
enable
configure terminal
prompt SwitchG
```

SPBM and IS-IS infrastructure configuration

```
BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.10
spbm 1 b-vid 20

ISIS CONFIGURATION

is-type 11
manual-area 30.0000
sys-name SwitchG
exit
router isis enable

VLAN CONFIGURATION

vlan create 20 name "B-VLAN" type spbm-bvlan
```

SwitchD

```
enable
configure terminal
prompt SwitchD

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.14
spbm 1 b-vid 20

ISIS CONFIGURATION

is-type 11
manual-area 30.0000
sys-name SwitchD
exit
router isis enable


VLAN CONFIGURATION

vlan create 20 name "B-VLAN" type spbm-bvlan
```

IS-IS SPBM Interface Configuration

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

Figure 26: IS-IS SPBM interface



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

SwitchC

PORT CONFIGURATION - PHASE II

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

SwitchG

PORT CONFIGURATION - PHASE II

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

MLT INTERFACE CONFIGURATION

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

SwitchD

MLT INTERFACE CONFIGURATION

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

IP multicast over Fabric Connect global configuration

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

SwitchC

```
enable
configure terminal
prompt SwitchC

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

SwitchG

```
enable
configure terminal
prompt SwitchG

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

SwitchD

```
enable
configure terminal
prompt SwitchD

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 — SwitchC

```
SwitchC:1# show isis interface
=====
ISIS Interfaces
=====
IFIDX      TYPE      LEVEL    OP-STATE  ADM-STATE  ADJ    UP-ADJ    SPBM-L1-METRIC
-----
Port1/30   pt-pt    Level 1  UP        UP         1      1         10

SwitchC:1# show isis adjacencies
=====
ISIS Adjacencies
=====
INTERFACE  L STATE   UPTIME    PRI  HOLDTIME  SYSID          HOST-NAME
-----
Port1/30   1 UP      1d 19:11:30 127  26        000e.6225.a3df SwitchG

1 out of 1 interfaces have formed an adjacency
=====

SwitchC: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 SwitchG    1/30
00:14:0d:a0:13:df   4000   0014.0da0.13df SwitchD    1/30
-----

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

```

```

SwitchC:1# show isis spbm unicast-tree 4000
Node:000e.6225.a3df.00 (SwitchG) -> ROOT
Node:0014.0da0.13df.00 (SwitchD) -> Node:000e.6225.a3df.00 (SwitchG) ->
ROOT

```

Checking operation — SwitchG

```
SwitchG:1# show isis interface
```

```

=====
                        ISIS Interfaces
=====
IFIDX      TYPE      LEVEL    OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
-----
Port1/5    pt-pt    Level 1  UP        UP         1      1       10
Mlt1       pt-pt    Level 1  UP        UP         1      1       10

```

```
SwitchG:1# show isis adjacencies
```

```

=====
                        ISIS Adjacencies
=====
INTERFACE L STATE    UPTIME          PRI HOLDTIME SYSID          HOST-NAME
-----
Port1/5   1 UP      1d 19:19:52    127 26         0015.e89f.e3df SwitchC
Mlt1      1 UP      04:57:34       127 20         0014.0da0.13df SwitchD

```

```
SwitchG: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 SwitchD    MLT-1
00:15:e8:9f:e3:df   4000   0015.e89f.e3df SwitchC    1/5

```

```

SwitchG:1# show isis spbm unicast-tree 4000
Node:0015.e89f.e3df.00 (SwitchC) -> ROOT
Node:0014.0da0.13df.00 (SwitchD) -> ROOT

```

Checking operation — SwitchD

```
SwitchD: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

```

```
SwitchD: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      SwitchG

SwitchD: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 SwitchG    MLT-1
00:15:e8:9f:e3:df    4000   0015.e89f.e3df SwitchC    MLT-1

SwitchD:1# show isis spbm unicast-tree 4000
Node:000e.6225.a3df.00 (SwitchG) -> ROOT
Node:0015.e89f.e3df.00 (SwitchC) -> Node:000e.6225.a3df.00 (SwitchG) ->
ROOT

```

Fabric Extend configuration examples

This section provides configuration examples to configure Fabric Extend in the following deployment scenarios.

- [Fabric Extend over IP using the GRT](#) on page 230
- [Fabric Extend over IP using a VRF](#) on page 233
- [Fabric Extend over VPLS](#) on page 236
- [Fabric Extend over Pseudowires](#) on page 239
- [Fabric Extend with ONAs in the core and branches](#) on page 241

For more configuration examples, see *Shortest Path Bridging (802.1aq) Technical Configuration Guide*, NN48500–617.

Fabric Extend over IP using the GRT

This example shows a typical Fabric Extend deployment with a 10/40 Gbps switch in the core and a 1 Gbps switch in one of the branch offices. The 10/40 Gbps switch supports Fabric Extend natively and is connected over an IP network to a 1 Gbps switch, which requires an ONA to encapsulate SPB traffic with a VXLAN header. The ONA sets up a bridge between the ONA device-side port and the ONA network-side port. Fabric Extend uses a VXLAN tunnel to send traffic to and from the 10/40 Gbps switch through the 1 Gbps switch to the ONA.

* Note:

- This deployment uses the **GRT** so the tunnel source IP address must be on the GRT, not on a VRF.
- If IP Shortcuts is enabled, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

- Add any IP address used for setting up the logical tunnel (such as local network and loopback IP addresses) to the IS-IS accept policy or exclude route-map to prevent these addresses from being advertised into IS-IS.

The following figure shows a sample Fabric Extend deployment over IP using the GRT.

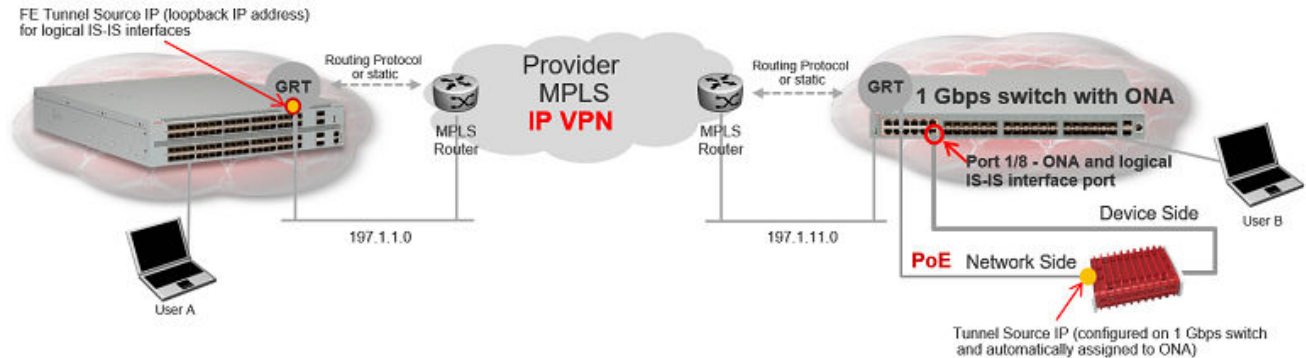


Figure 27: IP using GRT traffic flow




Figure 28: IP (GRT) traffic flow component view

For 10/40 Gbps Switches:

(The tunnel source IP address is configured in the GRT.)

```
Switch(config)# interface GigabitEthernet 2/14
Switch(config-if)# no shutdown
Switch(config-if)# default-vlan-id 0
Switch(config-if)# name "ospf-intf-SP-core"
Switch(config-if)# brouter port 2/14 vlan 3036 subnet 197.1.1.2/255.255.255.0
Switch(config-if)# no spanning-tree mstp force-port-state enable
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.1.2
Switch(config-isis)# exit

Switch(config)# logical-intf isis 255 dest-ip 197.1.6.2 name "Tunnel-to-Branch1"
Switch(config-isis-255-197.1.6.2)# isis
Switch(config-isis-255-197.1.6.2)# isis spbm 1
Switch(config-isis-255-197.1.6.2)# isis enable
Switch(config-isis-255-197.1.6.2)# exit

Switch(config)# ip prefix-list "isis-tunnel-addr" 197.1.0.0/16 ge 16 le 32
Switch(config)# route-map "deny-isis-tunnel-network" 1
Switch(route-map)# no permit
Switch(route-map)# enable
```

SPBM and IS-IS infrastructure configuration

```
Switch(route-map)# match network "isis-tunnel-addr"
Switch(route-map)# match protocol isis
Switch(route-map)# exit

Switch(config)# router isis
Switch(config-isis)# accept route-map "deny-isis-tunnel-network"
Switch(config-isis)# exit
Switch(config)# isis apply accept
```

For 1 Gbps Switches:

(The tunnel source address is a CLIP address on the GRT. This address is configured on the 1 Gbps switch and then automatically assigned to the ONA.)

```
Switch(config)# interface loopback 256
Switch(config-if)# ip address 256 197.1.6.2/255.255.255.0
Switch(config-if)# ip ospf 256
Switch(config-if)# exit

Switch(config)# interface GigabitEthernet 1/7-1/8; enables ports to ONA
Switch(config-if)# no shutdown
Switch(config-if)# exit

Switch(config)# vlan create 3037 name "ospf-intf-SP-core" type port-mstprstp 0
Switch(config)# vlan members 3037 1/49-1/50 portmember
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 1/49-1/50
Switch(config)# mlt 11 vlan 3037
Switch(config)# interface vlan 3037
Switch(config-if)# ip address 197.1.11.2 255.255.255.0 0
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3025 name "ONA-Mgmt-vlan" type port-mstprstp 0
Switch(config)# vlan members 3025 1/7 portmember
Switch(config)# interface vlan 3025
Switch(config-if)# ip address 197.2.1.1 255.255.255.0 3
Switch(config-if)# exit

Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 enable
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 mode bootp_dhcp

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.6.2 port 1/8 mtu 1950
Switch(config-isis)# exit
Switch(config)# logical-intf isis 255 dest-ip 197.1.1.2 name "Tunnel-to-HQ"
Switch(config-isis-255-197.1.1.2)# isis
Switch(config-isis-255-197.1.1.2)# isis spbm 1
Switch(config-isis-255-197.1.1.2)# isis enable
Switch(config-isis-255-197.1.1.2)# exit

Switch(config)# ip prefix-list "isis-tunnel-addr" 197.1.0.0/16 ge 16 le 32
Switch(config)# route-map "deny-isis-tunnel-network" 1
Switch(route-map)# no permit
Switch(route-map)# enable
Switch(route-map)# match network "isis-tunnel-addr"
Switch(route-map)# match protocol isis
Switch(route-map)# exit

Switch(config)# router isis
Switch(config-isis)# accept route-map "deny-isis-tunnel-network"
```

```
Switch(config-isis)# exit
Switch(config)# isis apply accept
```

Intermediate routers are typically configured by an Internet service provider (ISP). The following configurations are for reference only.

For Intermediate Router 1:

```
Switch(config)# interface GigabitEthernet 8/19
Switch(config-if)# default-vlan-id 0
Switch(config-if)# name "ospf-intf-from-Headoffice"
Switch(config-if)# no shutdown
Switch(config-if)# brouter port 8/19 vlan 3036 subnet 197.1.1.3/255.255.255.0 mac-offset 2
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.1 255.255.255.0 2
Switch(config)# ip ospf enable
Switch(config)# exit
```

For Intermediate Router 2:

```
Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.2 255.255.255.0 0
Switch(config)# ip ospf enable
Switch(config)# exit

Switch(config)# vlan create 3037 name "ospf-intf-from-branch1" type port-mstprstp 0
Switch(config)# vlan members 3037 8/21-8/22 portmember
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 8/21-8/22
Switch(config)# mlt 11 vlan 3037
Switch(config)# interface Vlan 3037
Switch(config)# ip address 197.1.11.1 255.255.255.0 5
Switch(config)# ip ospf enable
Switch(config)# exit
```

Fabric Extend over IP using a VRF

This example is the same as the previous IP example except this Fabric Extend deployment uses a VRF instead of the GRT. Because this deployment is using a **VRF**, Fabric Extend has the following requirements:

- Configure a CLIP and tunnel source IP address on the same VRF.
- Remote management of the 1 Gbps switch is only possible after establishing IP Shortcut over IS-IS. (Alternatively, you can enable GRT-VRF redistribution locally.)

The following figure shows a sample Fabric Extend deployment over IP using a VRF.

SPBM and IS-IS infrastructure configuration

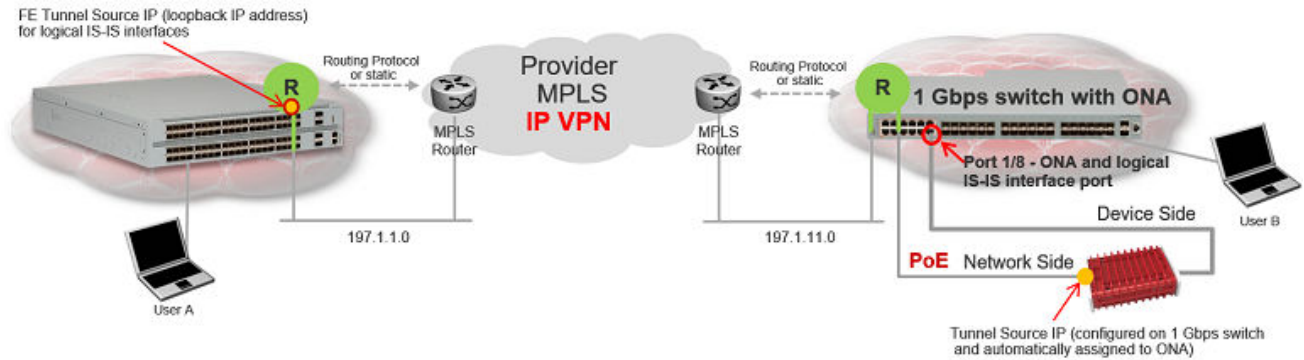


Figure 29: IP using VRF traffic flow




Figure 30: IP (VRF) traffic flow component view

For 10/40 Gbps Switches:

(The tunnel source IP address is configured as a brouter address in the VRF.)

```
Switch(config)# ip vrf vrf24
Switch(config)# router vrf vrf24
Switch(router-vrf)# ip ospf
Switch(router-vrf)# ip ospf admin-state
Switch(router-vrf)# exit

Switch(config)# interface GigabitEthernet 2/14
Switch(config-if)# no shutdown
Switch(config-if)# default-vlan-id 0
Switch(config-if)# name "ospf-intf-SP-core"
Switch(config-if)# vrf vrf24
Switch(config-if)# brouter port 2/14 vlan 3036 subnet 197.1.1.2/255.255.255.0 mac-offset 1
Switch(config-if)# no spanning-tree mstp force-port-state enable
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.1.2 vrf vrf24
Switch(config-isis)# exit

Switch(config)# logical-intf isis 255 dest-ip 197.1.6.2 name "Tunnel-to-Branch1"
Switch(config-isis-255-197.1.6.2)# isis
Switch(config-isis-255-197.1.6.2)# isis spbm 1
Switch(config-isis-255-197.1.6.2)# isis enable
Switch(config-isis-255-197.1.6.2)# exit
```

For 1 Gbps Switches:

(The tunnel source address is a CLIP address on the VRF. This address is configured on the 1 Gbps switch and then automatically assigned to the ONA.)

```
Switch(config)# ip vrf vrf24

Switch(config)# router vrf vrf24
Switch(router-vrf)# ip ospf
Switch(router-vrf)# ip ospf admin-state
Switch(router-vrf)# exit

Switch(config)# interface loopback 256
Switch(config-if)# ip address 197.1.6.2 255.255.255.255 vrf vrf24
Switch(config-if)# ip ospf vrf vrf24
Switch(config-if)# exit

Switch(config)# interface GigabitEthernet 1/7-1/8; enables ports to ONA
Switch(config-if)# no shutdown
Switch(config-if)# exit

Switch(config)# vlan create 3037 name "ospf-intf-SP-core" type port-mstprstp 0
Switch(config)# vlan members 3037 1/49-1/50 portmember
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 1/49-1/50
Switch(config)# mlt 11 vlan 3037
Switch(config)# interface vlan 3037
Switch(config-if)# vrf vrf24
Switch(config-if)# ip address 197.1.11.2 255.255.255.0 0
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3025 name "ONA-Mgmt-vlan" type port-mstprstp 0
Switch(config)# vlan members 3025 1/7 portmember
Switch(config)# interface vlan 3025
Switch(config-if)# ip address 197.2.1.1 255.255.255.0 3
Switch(config-if)# exit

Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 enable
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 mode bootp_dhcp

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.6.2 port 1/8 vrf vrf24 mtu 1950
Switch(config-isis)# exit

Switch(config)# logical-intf isis 255 dest-ip 197.1.1.2 name "Tunnel-to-HQ"
Switch(config-isis-255-197.1.1.2)# isis
Switch(config-isis-255-197.1.1.2)# isis spbm 1
Switch(config-isis-255-197.1.1.2)# isis enable
Switch(config-isis-255-197.1.1.2)# exit
```

Intermediate routers are typically configured by an Internet service provider (ISP). The following configurations are for reference only.

For Intermediate Router 1:

```
Switch(config)# interface GigabitEthernet 8/19
Switch(config-if)# default-vlan-id 0
Switch(config-if)# name "ospf-intf-from-Headoffice"
Switch(config-if)# no shutdown
Switch(config-if)# brouter port 8/19 vlan 3036 subnet 197.1.1.3/255.255.255.0 mac-offset 2
```

SPBM and IS-IS infrastructure configuration

```
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.1 255.255.255.0 2
Switch(config)# ip ospf enable
Switch(config)# exit
```

For Intermediate Router 2:

```
Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.2 255.255.255.0 0
Switch(config)# ip ospf enable
Switch(config)# exit

Switch(config)# vlan create 3037 name "ospf-intf-from-branch1" type port-mstprstp 0
Switch(config)# vlan members 3037 8/21-8/22 portmember
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 8/21-8/22
Switch(config)# mlt 11 vlan 3037
Switch(config)# interface Vlan 3037
Switch(config)# ip address 197.1.11.1 255.255.255.0 5
Switch(config)# ip ospf enable
Switch(config)# exit
```

Fabric Extend over VPLS

This example shows a Fabric Extend deployment over MPLS Virtual Private LAN Service (VPLS). In this scenario, VPLS emulates a LAN with full mesh connectivity. The SPB nodes connect with point-to-point Ethernet links and also use MPLS for normal forwarding.

The following figure shows a sample Fabric Extend deployment over VPLS.

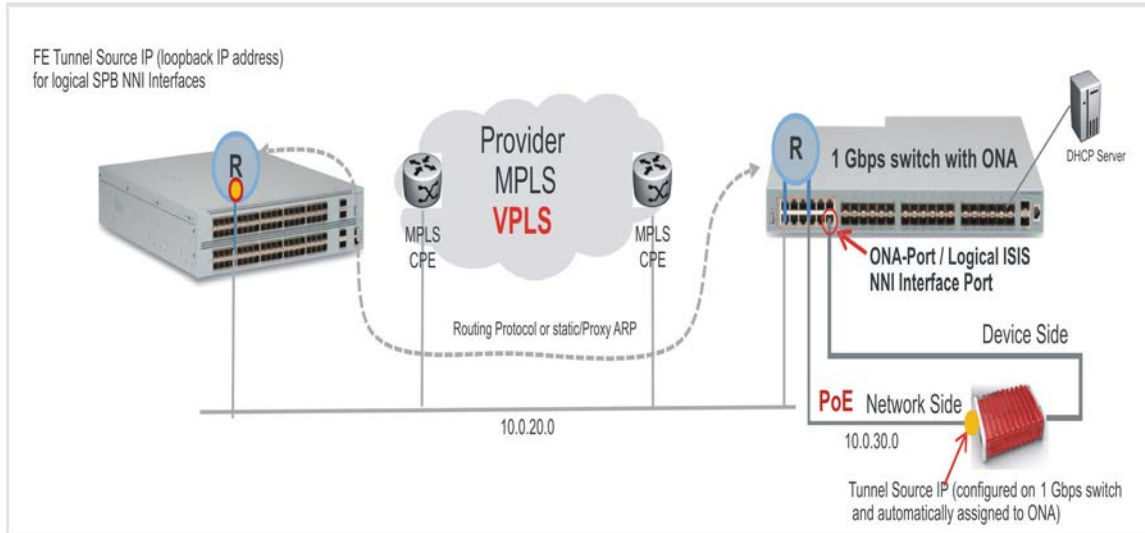


Figure 31: FE over VPLS traffic flow

For 10/40 Gbps Switches:

(The tunnel source IP address is configured as a brouter address in the VRF.)

```
Switch(config)# interface GigabitEthernet 2/14
Switch(config-if)# no shutdown
Switch(config-if)# default-vlan-id 0
Switch(config-if)# name "ospf-intf-SP-core"
Switch(config-if)# vrf vrf24
Switch(config-if)# brouter port 2/14 vlan 3036 subnet 197.1.1.2/255.255.255.0 mac-offset 1
Switch(config-if)# no spanning-tree mstp force-port-state enable
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.1.2 vrf vrf24
Switch(config-isis)# exit
Switch(config)# logical-intf isis 255 dest-ip 197.1.6.2 name "Tunnel-to-Branch1"
Switch(config-isis-255-197.1.6.2)# isis
Switch(config-isis-255-197.1.6.2)# isis spbm 1
Switch(config-isis-255-197.1.6.2)# isis enable
Switch(config-isis-255-197.1.6.2)# exit
```

For 1 Gbps Switches:

(The tunnel source address is a CLIP address on the VRF. This address is configured on the 1 Gbps switch and then automatically assigned to the ONA.)

```
Switch(config)# ip vrf vrf24

Switch(config)# router vrf vrf24
Switch(router-vrf)# ip ospf
Switch(router-vrf)# ip ospf admin-state
Switch(router-vrf)# exit

Switch(config)# interface loopback 256
Switch(config-if)# ip address 197.1.6.2 255.255.255.255 vrf vrf24
Switch(config-if)# ip ospf vrf vrf24
```


SPBM and IS-IS infrastructure configuration

```
Switch(config-if)# exit

Switch(config)# interface GigabitEthernet 1/7-1/8; enables ports to ONA
Switch(config-if)# no shutdown
Switch(config-if)# exit

Switch(config)# vlan create 3037 name "ospf-intf-SP-core" type port-mstprstp 0
Switch(config)# vlan members 3037 1/49-1/50 portmember
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 1/49-1/50
Switch(config)# mlt 11 vlan 3037
Switch(config)# interface vlan 3037
Switch(config-if)# vrf tunnel
Switch(config-if)# ip address 197.1.11.2 255.255.255.0 0
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3025 name "ONA-Mgmt-vlan" type port-mstprstp 0
Switch(config)# vlan members 3025 1/7 portmember
Switch(config)# interface vlan 3025
Switch(config-if)# ip address 197.2.1.1 255.255.255.0 3
Switch(config-if)# exit

Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 enable
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 mode bootp_dhcp

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.6.2 port 1/8 vrf vrf24 mtu 1950
Switch(config-isis)# exit
Switch(config)# logical-intf isis 255 dest-ip 197.1.1.2 name "Tunnel-to-HQ"
Switch(config-isis-255-197.1.1.2)# isis
Switch(config-isis-255-197.1.1.2)# isis spbm 1
Switch(config-isis-255-197.1.1.2)# isis enable
Switch(config-isis-255-197.1.1.2)# exit
```

Intermediate routers are typically configured by an Internet service provider (ISP). The following configurations are for reference only.

For Intermediate Router 1:

```
Switch(config)# interface GigabitEthernet 8/19
Switch(config-if)# default-vlan-id 0
Switch(config-if)# name "ospf-intf-from-Headoffice"
Switch(config-if)# no shutdown
Switch(config-if)# brouter port 8/19 vlan 3036 subnet 197.1.1.3/255.255.255.0 mac-offset 2
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.1 255.255.255.0 2
Switch(config)# ip ospf enable
Switch(config)# exit
```

For Intermediate Router 2:

```

Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.2 255.255.255.0 0
Switch(config)# ip ospf enable
Switch(config)# exit

Switch(config)# vlan create 3037 name "ospf-intf-from-branch1" type port-mstprstp 0
Switch(config)# vlan members 3037 8/21-8/22 portmember
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 8/21-8/22
Switch(config)# mlt 11 vlan 3037
Switch(config)# interface Vlan 3037
Switch(config)# ip address 197.1.11.1 255.255.255.0 5
Switch(config)# ip ospf enable
Switch(config)# exit

```

Fabric Extend over Layer 2 Pseudowire

This example shows a Fabric Extend deployment using service provider VLAN tunnels over MPLS Pseudowire. In this scenario, you map two dedicated VLAN IDs (VIDs) from the Hub to the Spoke sites. Then the logical IS-IS interfaces translate the BVIDs to map them to the per branch provider VIDs. Because the tunnels are point-to-point VLAN connections, not VXLAN, there is no need to encapsulate a VXLAN header to SPB packets. Therefore, the 1 Gbps switches in this type of deployment do not require ONAs.

! Important:

10/40 Gbps switch — — — — — Core — — — — — 1 Gbps switch

- You cannot have IS-IS in the Core.
- Do not create the two VLANs represented in the logical interface connection on the BEBs. If you do, you will not be able add any ports to be members of those VLANs. One links the port that is facing the core and those VLANs in the logical interface connection.

The following figure shows a sample Fabric Extend deployment over Pseudowire.

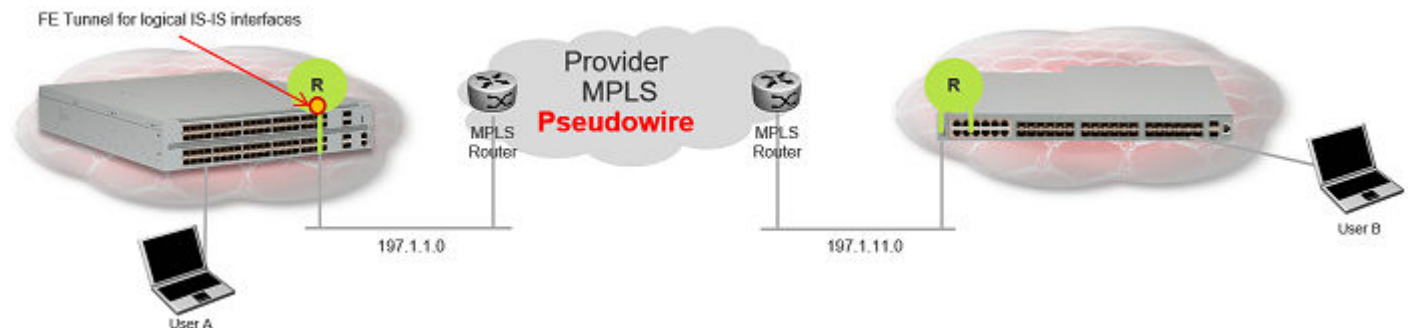


Figure 32: FE over Pseudowire traffic flow



Figure 33: FE over Pseudowire traffic flow component view

For 10/40 Gbps Switches:

*** Note:**

Logical interface VLANs cannot be the same as the SPBM B-VLANs and you cannot create these VLANs locally. Use these VLANs for configuring the logical interface only.

```
Switch(config)# logical-intf isis 255 vid 200,300 primary-vid 200 port 2/14 name
fe_to_Switch
Switch(config-isis-255)# isis
Switch(config-isis-255)# isis spbm 1
Switch(config-isis-255)# isis enable
Switch(config-isis-255)# exit
```

For 1 Gbps Switches:

*** Note:**

Logical interface VLANs cannot be the same as the SPBM B-VLANs and you cannot create these VLANs locally. Use these VLANs for configuring the logical interface only.

```
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 1/49-1/50
Switch(config)# router isis enable

Switch(config)# logical-intf isis 255 vid 200,300 primary-vid 200 mlt 11 name fe_to_Switch
Switch(config-isis-255)# isis
Switch(config-isis-255)# isis spbm 1
Switch(config-isis-255)# isis enable
Switch(config-isis-255)# exit
```

Intermediate routers are typically configured by an Internet service provider (ISP). The following configurations are for reference only.

For Intermediate Router 1:

```
Switch(config)# vlan create 200 type port-mstprstp 1
Switch(config)# vlan create 300 type port-mstprstp 1
Switch(config)# vlan member add 200 8/1,8/19
Switch(config)# vlan member add 300 8/1,8/19
```

For Intermediate Router 2:

```
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q
Switch(config)# mlt 11 mem 8/21-8/22
Switch(config)# vlan create 200 type port-mstprstp 1
Switch(config)# vlan create 300 type port-mstprstp 1
Switch(config)# vlan member add 200 8/1
Switch(config)# vlan mlt 200 11
```

```
Switch(config)# vlan member add 300 8/1
Switch(config)# vlan mlt 300 11
```

Fabric Extend with ONAs in the core and branches

This example shows a Fabric Extend deployment with 1 Gbps switches in the core of the network and in the branch sites. This type of deployment is not only a lower cost Fabric Extend solution, it also addresses situations where large MTU sizes (over 1594 bytes) are a problem for the Service Provider.

MTU sizes less than 1594 bytes require fragmentation and reassembly of packets and the 1 Gbps switch with ONA supports fragmentation and reassembly. However, you must have 1 Gbps switches with ONAs at BOTH ends of the IP WAN connection.

! Important:

There is no fragmentation/reassembly support in Layer 2 core solutions.

The following figure shows a sample Fabric Extend deployment using VRFs with both switches.

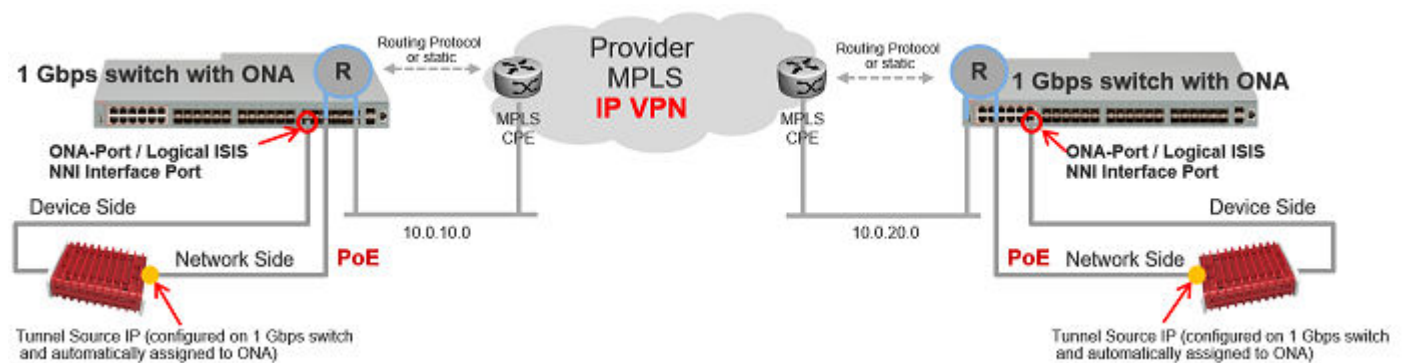


Figure 34: Fabric Extend traffic flow




Figure 35: Fabric Extend traffic flow component view

Core switch configuration:

(The tunnel source address is a CLIP address on the VRF. This address is configured on the switch and then automatically assigned to the ONA.)

```
Switch(config)# ip vrf vrf24
Switch(config)# router vrf vrf24
```

SPBM and IS-IS infrastructure configuration

```
Switch(router-vrf)# ip ospf
Switch(router-vrf)# ip ospf admin-state
Switch(router-vrf)# exit

Switch(config)# interface loopback 256
Switch(config-if)# ip address 197.1.5.2 255.255.255.255 vrf vrf24
Switch(config-if)# ip ospf vrf vrf24
Switch(config-if)# exit

Switch(config)# interface GigabitEthernet 1/7-1/8; enables ports to ONA
Switch(config-if)# no shutdown
Switch(config-if)# exit

Switch(config)# vlan create 3036 name "ospf-intf-SP-core" type port-mstprstp 0
Switch(config)# vlan members 3036 1/14 portmember
Switch(config)# interface vlan 3036
Switch(config)# vrf vrf24
Switch(config-if)# ip address 197.1.1.2 255.255.255.0 0
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3024 name "ONA-Mgmt-vlan" type port-mstprstp 0
Switch(config)# vlan members 3024 1/8 portmember
Switch(config)# interface vlan 3024
Switch(config-if)# ip address 197.3.1.1 255.255.255.0 3
Switch(config-if)# exit

Switch(config)# ip dhcp-relay fwd-path 197.3.1.1 197.10.1.11
Switch(config)# ip dhcp-relay fwd-path 197.3.1.1 197.10.1.11 enable
Switch(config)# ip dhcp-relay fwd-path 197.3.1.1 197.10.1.11 mode bootp_dhcp

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.5.2 port 1/7 vrf vrf24 mtu 1950
Switch(config-isis)# exit
Switch(config)# logical-intf isis 255 dest-ip 197.1.6.2 name "Tunnel-to-Branch1"
Switch(config-isis-255-197.1.6.2)# isis
Switch(config-isis-255-197.1.6.2)# isis spbm 1
Switch(config-isis-255-197.1.6.2)# isis enable
Switch(config-isis-255-197.2.1.1)# exit
```

Branch switch configuration:

(The tunnel source address is a CLIP address on the VRF. This address is configured on the switch and then automatically assigned to the ONA.)

```
Switch(config)# ip vrf vrf24

Switch(config)# router vrf vrf24
Switch(router-vrf)# ip ospf
Switch(router-vrf)# ip ospf admin-state
Switch(router-vrf)# exit

Switch(config)# interface loopback 256
Switch(config-if)# ip address 197.1.6.2 255.255.255.255 vrf vrf24
Switch(config-if)# ip ospf vrf vrf24
Switch(config-if)# exit

Switch(config)# interface GigabitEthernet 1/7-1/8; enables ports to ONA
Switch(config-if)# no shutdown
Switch(config-if)# exit
```

```

Switch(config)# vlan create 3037 name "ospf-intf-SP-core" type port-mstprstp 0
Switch(config)# vlan members 3037 1/49-1/50 portmember
Switch(config)# interface vlan 3037
Switch(config-if)# vrf vrf24
Switch(config-if)# ip address 197.1.11.2 255.255.255.0 0
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3025 name "ONA-Mgmt-vlan" type port-mstprstp 0
Switch(config)# vlan members 3025 1/8 portmember
Switch(config)# interface vlan 3025
Switch(config-if)# ip address 197.2.1.1 255.255.255.0 3
Switch(config-if)# exit

Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 enable
Switch(config)# ip dhcp-relay fwd-path 197.2.1.1 197.10.1.11 mode bootp_dhcp

Switch(config)# router isis
Switch(config-isis)# ip-tunnel-source-address 197.1.6.2 port 1/7 vrf vrf24 mtu 1950
Switch(config-isis)# exit
Switch(config)# logical-intf isis 255 dest-ip 197.1.5.2 name "Tunnel-to-HQ"
Switch(config-isis-255-197.1.1.2)# isis
Switch(config-isis-255-197.1.1.2)# isis spbm 1
Switch(config-isis-255-197.1.1.2)# isis enable
Switch(config-isis-255-197.1.1.2)# exit

```

Intermediate routers are typically configured by an Internet service provider (ISP). The following configurations are for reference only.

Intermediate Router 1 configuration:

```

Switch(config)# interface GigabitEthernet 8/19
Switch(config-if)# default-vlan-id 0
Switch(config-if)# name "ospf-intf-from-Headoffice"
Switch(config-if)# no shutdown
Switch(config-if)# brouter port 8/19 vlan 3036 subnet 197.1.1.3/255.255.255.0 mac-offset 2
Switch(config-if)# ip ospf enable
Switch(config-if)# exit

Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.1 255.255.255.0 2
Switch(config)# ip ospf enable
Switch(config)# exit

```

Intermediate Router 2 configuration:

```

Switch(config)# vlan create 3039 name "core-ospf-vlan" type port-mstprstp 0
Switch(config)# vlan members 3039 8/1 portmember
Switch(config)# interface Vlan 3039
Switch(config)# ip address 197.1.8.2 255.255.255.0 0
Switch(config)# ip ospf enable
Switch(config)# exit

Switch(config)# vlan create 3037 name "ospf-intf-from-branch1" type port-mstprstp 0
Switch(config)# vlan members 3037 8/21-8/22 portmember
Switch(config)# mlt 11
Switch(config)# mlt 11 encapsulation dot1q

```



```
Switch(config)# mlt 11 mem 8/21-8/22
Switch(config)# mlt 11 vlan 3037
Switch(config)# interface Vlan 3037
Switch(config)# ip address 197.1.11.1 255.255.255.0 5
Switch(config)# ip ospf enable
Switch(config)# exit
```

Fabric Attach configuration examples

This section provides configuration examples to configure Fabric Attach.


Configuring a Fabric Attach solution

The following section describes a simple configuration example to configure Fabric Attach (FA) at the edge of a Fabric Connect network. This is a typical deployment at its simplest level and is powerful because of its use in conjunction with a Fabric Connect core.

About this task

Configuring FA primarily consists of configuring the FA Server. The FA Server in turn *discovers* neighboring FA component devices (like the FA Proxies and FA Clients) using FA TLVs within the LLDP PDUs.

In the following deployment, the switch at the edge of the Fabric Connect cloud is configured as the FA Server. On this switch, FA is enabled globally and at the interface (port) level. Another switch, functioning as the FA Proxy connects to the FA enabled port (1/3) on the FA Server. User A is an end user device that needs to send and receive data traffic from User B (another end user device) across the network.



Before you begin

Configure SPBM and IS-IS on the edge and core switches. For more information, see [Configuring minimum SPBM and IS-IS parameters](#) on page 86.

Procedure

Configure the edge switch (BEB) as the FA Server:

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable FA globally:

```
fa enable
```

3. Enter port interface configuration mode:

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

4. Enable FA on the port:

```
fa enable
```

*** Note:**

Enabling FA automatically enables message authentication. Also, the authentication key is set to the default value and appears encrypted on the output.

*** Note:**

Enabling FA on a port not only enables tagging but also disables spanning tree on that port.

Verify global and interface level FA configuration:

5. Verify global configuration of FA using one of the following commands:

- `show fa`
- `show fa agent`

6. Verify interface level configuration of FA:

```
show fa interface
```

7. Verify the discovery of clients attaching to the FA Server:

```
show fa elements
```

8. Display the FA I-SID-to-VLAN assignments:

```
show fa assignment
```

To verify I-SID-to-VLAN assignments on a specific port, enter:


```
show fa assignment {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...]}
```

9. Verify creation of Switched UNI (ELAN) I-SIDs:

```
show i-sid elan
```

Example

SPBM and IS-IS configuration on the core and edge switches:

SPBM configuration:

```
Switch:1>en  
Switch:1#conf t  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#spbm  
Switch:1(config)#spbm ethertype 0x8100
```

IS-IS SPBM configuration:

```
Switch:1(config)#router isis  
Switch:1(config)#spbm 1  
Switch:1(config-isis)#spbm 1 nick-name 1.00.01  
Switch:1(config-isis)#spbm 1 b-vid 41-42 primary 41  
Switch:1(config-isis)#spbm 1 ip enable
```

IS-IS router configuration:

```
Switch:1(config-isis)#router isis  
Switch:1(config-isis)#sys-name BEB-Switch  
Switch:1(config-isis)#ip-source-address 3.3.3.3  
Switch:1(config-isis)#is-type 11  
Switch:1(config-isis)#system-id 0001.0001.0001  
Switch:1(config-isis)#manual-area c0.2000.000.00  
Switch:1(config-isis)#exit
```

Interface (port-level) configuration

```
Switch:1(config)#interface GigabitEthernet 1/2  
Switch:1(config-if)#no shutdown  
Switch:1(config-if)#isis  
Switch:1(config-if)#isis spbm 1  
Switch:1(config-if)#isis enable  
Switch:1(config-isis)#exit  
Switch(config)#vlan create 41 type spbm-vlan  
Switch(config)#vlan create 42 type spbm-vlan  
Switch(config)#router isis enable  
Switch(config)#show isis spbm
```

Configuration of the edge switch as the FA Server.

Enable FA globally.

```
Switch:1(config)#fa enable  
Switch:1(config)#show fa  
  
===== Fabric Attach Configuration =====  
  
FA Service : enabled  
FA Element Type : server  
FA Assignment Timeout : 240  
FA Discovery Timeout : 240  
FA Provision Mode : spbm
```

Enable FA on the port.

Enabling FA automatically enables message authentication. The authentication key is configured with the default value, which appears in encrypted format in the output.

```
Switch:1(config)#int gigabitEthernet 1/3
Switch:1(config-if)#fa enable
Switch:1(config-if)#show fa interface port 1/3
```

```
=====
                          Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT   MGMT   MSG AUTH  MSG AUTH
                STATUS ISID   CVID   STATUS    KEY
-----
Port1/3        enabled  0      0      enabled   ****

-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
```

```
Switch:1(config-if)#exit
Switch:1(config)#exit
```

Verify that the FA Proxy is discovered by the FA Server.

```
Switch:1(config)#show fa elements
```

```
=====
                          Fabric Attach Discovery Elements
=====
PORT   TYPE          MGMT          ELEM ASGN
      VLAN STATE  SYSTEM ID    AUTH AUTH
-----
1/3    proxy        2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP

-----
                          Fabric Attach Authentication Detail
=====
      ELEM OPER          ASGN OPER
PORT  AUTH STATUS          AUTH STATUS
-----
1/3   successAuth          successAuth

State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid

Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,  N= None

-----
2 out of 2 Total Num of fabric attach discovery elements displayed
```

Verify the FA I-SID-to-VLAN assignment. An `active` state indicates that the FA (ELAN) I-SID is successfully created with endpoint of type Switched UNI. By default, this I-SID is created for Layer 2.

```
Switch:1#show fa assignment
```

```
=====
                          Fabric Attach Assignment Map
=====
```

SPBM and IS-IS infrastructure configuration

```
=====
Interface  I-SID      Vlan      State      Origin
-----
1/3        44         2         active     proxy
-----
1 out of 1 Total Num of fabric attach assignment mappings displayed
=====
```

For Layer 3 support, you must configure a platform VLAN. The platform VLAN can have the same value as that of the C-VID or it can have a different value.

In this example, the platform VLAN has the same value as the C-VID.

```
Switch:1(config)#vlan create 2 type port-mstprstp 0
Switch:1(config)#vlan i-sid 2 44
```

```
Switch:1#show i-sid elan
```

```
=====
Isid Info
=====
ISID      ISID      VLANID    PORT      MLT      ORIGIN
ID        TYPE      ID        INTERFACES INTERFACES
-----
44        ELAN      2         c2:1/3    DISC_LOCAL
-----
c: customer vid    u: untagged-traffic
All 1 out of 1 Total Num of Elan i-sids displayed
```

Verify neighbor discovery on the FA Proxy switch:

Note that the edge switch (BEB) is discovered as the FA Server by the FA Proxy.

```
Switch:2(config)#show fa agent
```

```
Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Zero Touch Status: Enabled
Fabric Attach Auto Provision Setting: Proxy
Fabric Attach Provision Mode: SPBM
Fabric Attach Client Proxy Status: Enabled
Fabric Attach Standalone Proxy Status: Disabled
Fabric Attach Agent Timeout: 50 seconds
Fabric Attach Extended Logging Status: Enabled
Fabric Attach Primary Server Id: aa:bb:cc:dd:ee:11:30:01:00:01 (SPBM)
Fabric Attach Primary Server Descr: BEB-Switch (6.0.0.0_GA)
```

```
Switch:2(config)#show fa elements
```

```
Unit/   Element   Element   Element
Port    Type      Subtype   VLAN   Auth   System ID
-----
1/3     Server    Server (Auth)  0     AP     aa:bb:cc:dd:ee:11:30:01:00:01
```

```
Switch:2(config)#show fa i-sid
```

```
I-SID   VLAN   Source   Status
-----
44      2      Proxy   Active
```

Configuring Fabric Attach in an SMLT

The following example describes FA configuration and behavior in a dual-homed SMLT deployment.

The following figure shows a simple FA solution in a dual-homed SMLT deployment. In this deployment, a pair of BEB switches (BEB A and BEB B) operating as IST peers are configured as the FA Server. An access switch or a wiring closet switch configured as an FA Proxy connects to the FA Server. The FA Proxy advertises I-SID-to-VLAN assignment mappings to the FA Server. Both BEB switches receive the mapping information using LLDP PDUs containing assignment TLVs. The switch that learns the mapping first considers the I-SID to be discovered locally and creates the I-SID on its device. The mapping information is then shared with its IST peer switch. When the peer switch receives the mapping across IST in a new SMLT message, it too creates the I-SID on its device. This I-SID however, is considered to be discovered remotely because it is learnt from synchronization with the peer switch. The mappings can also be learned on the FA Server from *both* LLDP PDUs and from IST synchronization.




Figure 36: FA configuration in dual-homed SMLT

Before you begin

Ensure that the proxy device (for example, an access switch) is properly configured for FA. See the corresponding product documentation for information on how to configure FA on the switch.

Procedure

1. Configure SMLT and vIST on switches BEB A and BEB B.

Caution:

For the IST peer switches acting as the FA Server to transmit the same FA System ID (based on the virtual MAC), SMLT configuration on both the switches *must* be the same.

For detailed information on configuring SMLT and vIST, see *Configuring Link Aggregation, MLT, SMLT, and vIST*.

Configure BEB A and BEB B as the FA Server

Perform the following configuration on each switch.

2. Enter Global Configuration mode:

```
enable
configure terminal
```

3. Enable FA globally:

```
fa enable
```

4. Enter MLT interface configuration mode:

```
interface mlt <1-512>
```

5. Enable FA on the MLT:

```
fa enable
```

 **Note:**

Enabling FA automatically enables message authentication. Also, the authentication key is set to the default value and appears encrypted on the output.

6. **(Optional)** Configure an FA authentication key with a value different from that of the default value:

```
fa authentication-key [WORD<0-32>]
```

 **Caution:**

When you configure the FA authentication key, you *must* configure the same value on both BEB switches in the SMLT.

Verify global and MLT-level FA configuration on BEB A and BEB B:

7. Verify global configuration of FA using one of the following commands:

- `show fa`
- `show fa agent`

8. Verify MLT-level (interface-level) configuration of FA:

```
show fa interface
```

Verify FA discovery on BEB A and BEB B:

9. Verify discovery of the FA Proxy.

```
show fa elements
```

View FA I-SID-to-VLAN assignments on BEB A and BEB B:

10. View the FA I-SID-to-VLAN assignments:

```
show fa assignment
```

To view FA I-SID-to-VLAN assignments on specific ports, enter:

```
show fa assignment {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}
```

Verify creation of Switched UNI I-SIDs on BEB A and BEB B:

11. Verify creation of Switched UNI (ELAN) I-SIDs. Use the following commands:

- View ELAN I-SID information using `show i-sid elan`.
- View ELAN I-SID information on a specific MLT using `show mlt i-sid [<1-512>]`.

Note:

Viewing ELAN I-SID information on an MLT is very useful to understand the origin of the I-SID, when multiple client or proxy devices connecting to the FA Server using SMLT MLT advertise the *same* I-SID-to-VLAN mappings. In the event of a link failure on an MLT, the origin of the I-SID helps determine on which MLT, and thereby from which proxy or client device, the mappings were successfully learnt.

Example

SMLT configuration on BEB A and BEB B:

On BEB A:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface mlt 1
Switch:1(config)#smlt
```

On BEB B:

```
Switch:2>en
Switch:2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:2(config)#interface mlt 1
Switch:2(config)#smlt
```

vIST configuration on BEB A and BEB B:

On BEB A:

```
Switch:1(config)#vlan create 2261 type port-mstprstp 0
Switch:1(config)#vlan i-sid 2261 1502261
Switch:1(config)#interface vlan 2261
Switch:1(config)#ip address 10.22.61.1 255.255.255.0 2
```

Configure BEB B (IP address 10.22.61.2) as the IST peer.

```
Switch:1(config)#virtual-ist peer-ip 10.22.61.2 vlan 2261
Switch:1(config)#show virtual-ist
Switch:1(config)#exit
```

On BEB B:

```
Switch:2(config)#vlan create 2261 type port-mstprstp 0
Switch:2(config)#vlan i-sid 2261 1502261
Switch:2(config)#interface vlan 2261
Switch:2(config)#ip address 10.22.61.2 255.255.255.0 2
```

Configure BEB A (IP address 10.22.61.1) as the IST peer.

```
Switch:2(config)#virtual-ist peer-ip 10.22.61.1 vlan 2261
Switch:2(config)#show virtual-ist
Switch:2(config)#exit
```

FA configuration on BEB A:

Enable FA globally and on the MLT:

```
Switch:1(config)#fa enable
Switch:1(config)#show fa
=====
Fabric Attach Configuration
=====
FA Service : enabled
FA Element Type : server
FA Assignment Timeout : 240
FA Discovery Timeout : 240
FA Provision Mode : spbm
```

Optionally, configure an FA authentication key with the value `dual-homed-smlt`. Ensure that you configure the **same** value on both switches BEB A and BEB B.

```
Switch:1(config)#interface mlt 1
Switch:1(config-mlt)#fa authentication-key dual-homed-smlt
```

Enable FA on the MLT:

```
Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#exit
Switch:1(config)#show fa interface
=====
Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH  MSG AUTH
                STATUS ISID    CVID    STATUS    KEY
-----
Mlt1           enabled  0       0       enabled   ****
-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
```

Verify discovery of the FA Proxy:

```
Switch:1(config)#show fa elements
=====
Fabric Attach Discovery Elements
=====
PORT  TYPE      MGMT      ELEM ASGN
      TYPE      VLAN STATE  SYSTEM ID  AUTH AUTH
-----
3/21  proxy    2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
3/22  proxy    2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
3/23  proxy    2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
-----
Fabric Attach Authentication Detail
=====
PORT  ELEM OPER      ASGN OPER
      AUTH STATUS  AUTH STATUS
-----
3/21  successAuth    successAuth
```

```

3/22    successAuth                successAuth
3/23    successAuth                successAuth

State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid

Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,  N= None
-----

3 out of 3 Total Num of fabric attach discovery elements displayed

```

The FA Proxy advertises I-SID-to-VLAN assignment mappings to BEB A, on MLT ports 3/21 to 3/23. View the FA I-SID-to-VLAN assignments on BEB-A:

All ports in the MLT receive the FA assignment mappings, as shown in the following output.

```

Switch:1(config)#show fa assignment

=====
Fabric Attach Assignment Map
=====
Interface  I-SID    Vlan    State    Origin
-----
3/21       2         2       active   proxy
3/21       3         3       active   proxy
3/21       4         4       active   proxy
3/22       2         2       active   proxy
3/22       3         3       active   proxy
3/22       4         4       active   proxy
3/23       2         2       active   proxy
3/23       3         3       active   proxy
3/23       4         4       active   proxy

```

FA configuration on BEB B:

Enable FA globally and on the MLT:

```

Switch:2(config)#fa enable
Switch:2(config)#show fa

=====
Fabric Attach Configuration
=====

FA Service : enabled
FA Element Type : server
FA Assignment Timeout : 240
FA Discovery Timeout : 240
FA Provision Mode : spbm

```

Configure the FA authentication key `dual-homed-smlt`. Ensure that you configure the **same** value as on BEB A.

```

Switch:2(config)#interface mlt 1
Switch:2(config-mlt)#fa authentication-key dual-homed-smlt

```

Enable FA on the MLT:

```

Switch:2(config-mlt)#fa enable
Switch:2(config-mlt)#exit
Switch:2(config)#show fa interface

```


SPBM and IS-IS infrastructure configuration

```
=====
Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH MSG AUTH
                STATUS ISID    CVID    STATUS   KEY
-----
Mlt1           enabled 0       0       enabled  ****
-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
```

Verify discovery of FA Proxy:

```
Switch:2(config)#show fa elements
```

```
=====
Fabric Attach Discovery Elements
=====
PORT  TYPE          MGMT          ELEM ASGN
      VLAN STATE  SYSTEM ID    AUTH AUTH
-----
3/17  proxy        2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
3/18  proxy        2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
3/19  proxy        2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
3/20  proxy        2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
-----
```

```
=====
Fabric Attach Authentication Detail
=====
PORT  ELEM OPER          ASGN OPER
      AUTH STATUS    AUTH STATUS
-----
3/17  successAuth      successAuth
3/18  successAuth      successAuth
3/19  successAuth      successAuth
3/20  successAuth      successAuth
-----
```

State Legend: (Tagging/AutoConfig)

T= Tagged, U= Untagged, D= Disabled, S= Spbm, V= Vlan, I= Invalid

Auth Legend:

AP= Authentication Pass, AF= Authentication Fail,

NA= Not Authenticated, N= None

```
-----
4 out of 4 Total Num of fabric attach discovery elements displayed
-----
```

The FA Proxy device advertises I-SID-to-VLAN assignment mapping requests to BEB B on MLT ports 3/17 to 3/20.

View FA I-SID-to-VLAN assignments on BEB-B:

```
Switch:2(config)#show fa assignment 3/17
```

```
=====
Fabric Attach Assignment Map
=====
Interface  I-SID    Vlan    State    Origin
-----
3/17      2        2       active   proxy
-----
```

```
3/17      3      3      active  proxy
3/17      4      4      active  proxy
```

Verify creation of FA Switched UNI (ELAN) I-SIDs on BEB A and BEB B:

Verify the creation of FA Switched UNI (ELAN) I-SIDs on BEB A and BEB B. Note that the `ORIGIN` of the I-SIDs displays as `DISC_BOTH`

Since the I-SID-to-VLAN mappings are learnt from both LLDP PDUs (containing the Assignment TLVs) and from IST synchronization between the peers, the origin displays as `DISC_BOTH`.

On BEB A:

```
Switch:1(config)#show i-sid elan
=====
                        Isid Info
=====
ISID      ISID      VLANID    PORT      MLT      ORIGIN
ID        TYPE      VLANID    INTERFACES INTERFACES
-----
2         ELAN      N/A       -         c2:1     DISC_BOTH
3         ELAN      N/A       -         c3:1     DISC_BOTH
4         ELAN      N/A       -         c4:1     DISC_BOTH
```

View the I-SID information for MLT 1 on BEB A.

```
Switch:1(config)#show mlt i-sid 1
=====
                        MLT Isid Info
=====
MLTID     IFINDEX  ISID      VLANID    C-VID    ISID      ORIGIN      BPDU
          ID      ID        C-VID    TYPE     TYPE     ORIGIN
-----
1         6144    2         N/A       2        ELAN     DISC_BOTH
1         6144    3         N/A       3        ELAN     DISC_BOTH
1         6144    4         N/A       4        ELAN     DISC_BOTH
=====
3 out of 3 Total Num of i-sid endpoints displayed
```

On BEB B:

```
Switch:2(config)#show i-sid elan
=====
                        Isid Info
=====
ISID      ISID      VLANID    PORT      MLT      ORIGIN
ID        TYPE      VLANID    INTERFACES INTERFACES
-----
2         ELAN      N/A       -         c2:1     DISC_BOTH
3         ELAN      N/A       -         c3:1     DISC_BOTH
4         ELAN      N/A       -         c4:1     DISC_BOTH
```

View the I-SID information for MLT 1 on BEB B.

```
Switch:1(config)#show mlt i-sid 1
=====
                        MLT Isid Info
=====
MLTID     IFINDEX  ISID      VLANID    C-VID    ISID      ORIGIN      BPDU
          ID      ID        C-VID    TYPE     TYPE     ORIGIN
-----
```

```
-----
1      6144      2      N/A      2      ELAN      DISC_BOTH
1      6144      3      N/A      3      ELAN      DISC_BOTH
1      6144      4      N/A      4      ELAN      DISC_BOTH
-----
```

3 out of 3 Total Num of i-sid endpoints displayed

The following section describes the behavior if, for example, a link failure occurs between the FA Proxy and BEB B, as shown in the following figure.




Figure 37: FA behavior in dual-homed SMLT during a link failure

View the I-SID-to-VLAN assignments on BEB A:

```
Switch:1(config)#show fa assignment 3/21
-----
Fabric Attach Assignment Map
-----
Interface  I-SID      Vlan      State      Origin
-----
3/21       2          2         active     proxy
3/21       3          3         active     proxy
3/21       4          4         active     proxy
-----
```

View the Switched UNI (ELAN) I-SIDs created on BEB A.

Since BEB A first learns the mappings from the LLDP PDUs (containing the Assignment TLVs), the origin of the I-SIDs displays as DISC_LOCAL.

```
Switch:1(config)#show i-sid elan
-----
Isid Info
-----
ISID      ISID      VLANID    PORT      MLT        ORIGIN
ID        TYPE      VLANID    INTERFACES INTERFACES
-----
2         ELAN     N/A      -         c2:1      DISC_LOCAL
3         ELAN     N/A      -         c3:1      DISC_LOCAL
4         ELAN     N/A      -         c4:1      DISC_LOCAL
-----
```

View the I-SID information for MLT 1 on BEB A.

Switch:1(config)#show mlt i-sid 1

```
=====
                                MLT Isid Info
=====
```

MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	BPDU
1	6144	2	N/A	2	ELAN	DISC_LOCAL	
1	6144	3	N/A	3	ELAN	DISC_LOCAL	
1	6144	4	N/A	4	ELAN	DISC_LOCAL	

3 out of 3 Total Num of i-sid endpoints displayed

View the Switched UNI (ELAN) I-SIDs created on BEB B.

Since BEB B learns the mappings only through IST peer synchronization, the origin of the I-SIDs displays as DISC_REMOTE.

BEB-B:1(config-mlt)#show i-sid elan

```
=====
                                Isid Info
=====
```

ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN
2	ELAN	N/A	-	c2:1	DISC_REMOTE
3	ELAN	N/A	-	c3:1	DISC_REMOTE
4	ELAN	N/A	-	c4:1	DISC_REMOTE

View the I-SID information for MLT 1 on BEB B.

Switch:1(config)#show mlt i-sid 1

```
=====
                                MLT Isid Info
=====
```

MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	BPDU
1	6144	2	N/A	2	ELAN	DISC_REMOTE	
1	6144	3	N/A	3	ELAN	DISC_REMOTE	
1	6144	4	N/A	4	ELAN	DISC_REMOTE	

3 out of 3 Total Num of i-sid endpoints displayed

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 VSNs, and Inter-Virtual Services Networks (VSNs) routing.

Fabric Connect Service Types

The Fabric Connect technology delivers Layer 2 and Layer 3 virtualization. These virtualized Layer 2 (L2) and Layer 3 (L3) instances are referred to as Virtual Service Networks (VSNs). A Service Identifier (I-SID) is used to uniquely distinguish these service instances network-wide, and a User Network Interface (UNI) is the boundary or demarcation point between the “service layer” of traditional networks, that is VLANs and VRFs, and the Fabric Connect “service layer”, that is L2 & L3 VSNs.

- Layer 2 VSNs are virtual broadcast domains interconnecting UNI members that share the same L2 VSN I-SID. MAC learning/aging is applied to all L2 VSNs.
- Layer 3 VSNs are virtual routed L3 networks (L3 VPN) leveraging IS-IS as the routing protocol between VRFs that share the same L3 VSN I-SID.

Fabric Connect uses the User-Network-Interface (UNI) to denote the capabilities and attributes of the service interfaces. Fabric connect devices support the following UNI types:

- *VLAN UNI (C-VLAN)* — a device-specific VLAN-ID maps to a L2 VSN I-SID – all device physical ports that are associated with the VLAN are therefore associated with the UNI.
- *Flex UNI* — it has the following sub-types:
 - *Switched UNI* — a VLAN-ID and a given port (VID, port) maps to a L2 VSN I-SID. With this UNI type, VLAN-IDs can be reused on other ports and therefore mapped to different I-SIDs.
 - *Transparent Port UNI* — a physical port maps to a L2 VSN I-SID (all traffic through that port, 802.1Q tagged or untagged, ingress and egress is mapped to the I-SID). Note: All VLANs on a *Transparent Port UNI* interface now share the same single MAC learning table of the *Transparent Port UNI* I-SID.
- *E-Tree UNI* — it extends Private VLANs beyond one Switch to form a network-wide E-Tree service infrastructure. An E-Tree UNI is a L2 VSN where broadcast traffic flows from Hub sites to Spokes sites, and from Spokes to Hubs, but not between Spoke sites. E-Tree Hubs can be formed with any VLAN UNI, while E-Tree Spokes must be configured as Private VLAN UNIs.

- *L3 VSN UNI* — a device-specific VRF maps to an I-SID, and the control plane exchanges the L3 routes belonging to the same I-SID. All VRFs in a network sharing the same L3 I-SID effectively form an L3 VPN. L3 VSNs can be configured to simultaneously support both IP Unicast and IP Multicast.

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 B-MAC-DA, using the shortest path topology learned using IS-IS.

The following figure shows a sample campus SPBM Layer 2 VSN network.

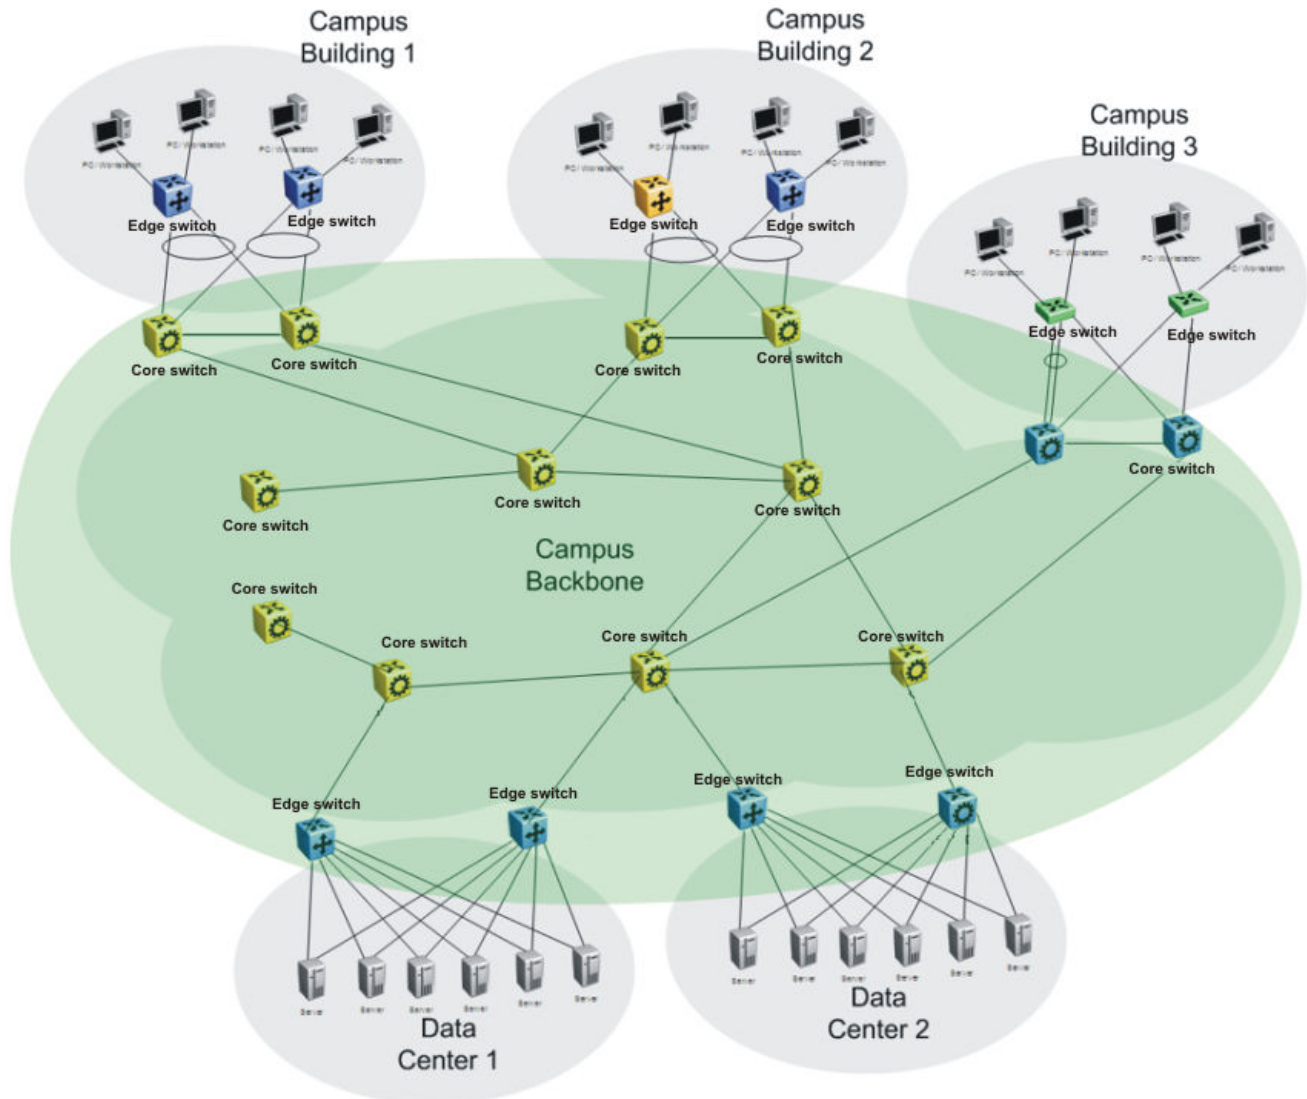


Figure 38: SPBM L2 VSN in a campus

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.

Configuration difference from ERS 8800

One major difference between these VSP switches and the ERS 8800 is how they connect to two SMLT devices.

The ERS 8800 uses an interswitch trunk (IST). The IST connects directly to two SMLT devices with a dedicated MLT and runs IS-IS over it. The dedicated MLT carries the IST control traffic and data traffic during an SMLT failover. This feature dramatically improves resiliency over other methods. However, if the dedicated MLT breaks, then there is no way to communicate between the IST peers, which causes traffic loss.

These VSP switches use a virtual IST (vIST) that eliminates this single point of failure. The vIST feature creates a virtualized IST channel in the SPBM cloud. With vIST, the IST tunnel is always up as long as there is SPBM connectivity between the vIST peers. vIST also interoperates between any two devices that support vIST, and the devices do not have to be the same type of device.

Before you can create a vIST, you must do the following:

- Enable SPBM and IS-IS globally.
- Configure SPBM and IS-IS.
- Create a VLAN (that is not used anywhere else) for each peer.
- Create an I-SID that is not used anywhere else.
- Configure an IP address for the vIST VLAN.
- Configure an L2 VSN by assigning an I-SID to the C-VLAN, which is used by the vIST.

Important:

An I-SID must be assigned to every VLAN that is a member of an L2 VSN. Also if an L2 VSN is created on one vIST Peer, it must also be created on the other vIST peer.

For information about vIST, see *Configuring Link Aggregation, MLT, SMLT, and vIST*.

Transparent Port UNI

The Transparent port user-network-interface (T-UNI) feature enables you to map an entire port or an MLT to an I-SID. CMAC learning is done against the I-SID. *Transparent Port UNI* configures a transparent port where all traffic is MAC switched on an internal virtual port using the assigned I-SID. No VLAN is involved in this process. Devices switch tagged and untagged traffic in the assigned I-SID regardless of the VLAN ID. The T-UNI port or MLT can be either static or LACP and is not a member of any VLAN or Spanning Tree Group (STG). The T-UNI port or MLT is always in the forwarding state.

You can map multiple ports to a T-UNI I-SID. Multiple ports on the same switch and on other BEBs can use the common I-SID to switch traffic.

Transparent Port UNI is a point to point service and all traffic that ingress the UNI egress from the remote UNI end-point.

Transparent

Transparent Port UNI is transparent because the MAC learning occurs within the I-SID, and packets that ingress from any CVLAN are processed in an identical manner. Devices switch tagged and untagged traffic in the assigned I-SID. Devices switch control protocols, such as BPDU, LACP, LLDP, and others, in the assigned I-SID, rather than forwarding to the CP.

The service classification of packets that are received on a T-UNI port, is independent of the VLAN ID values present in those packets. All data packets received on a T-UNI port are classified into the same service. When data packets enter and exit the T-UNI service, no VLAN tag modifications are performed on the data packets.

T-UNI based MAC learning

When a packet ingresses a port or MLT associated with a T-UNI I-SID, the system performs MAC lookup based on the I-SID. A packet that ingresses a T-UNI port on a BEB can transfer through the SPB network, or can egress out another T-UNI port configured to the same I-SID.

When a packet ingresses an NNI port, before egressing a T-UNI port, the system performs a MAC Destination Address (DA) lookup based on the I-SID. If the DA lookup fails, the packet floods to all T-UNI ports.

Considerations

Consider the following when you configure a *Transparent Port UNI*:

- You cannot configure a *Transparent Port UNI* on the same I-SID as a C-VLAN.
- A *Transparent Port UNI* port or MLT is not a member of any VLAN and does not belong to any STG.
- Ensure that you always associate a T-UNI LACP MLT with a VLAN (even if it is the default VLAN) before adding it to a T-UNI ISID. Otherwise, traffic is not forwarded on the T-UNI LACP MLT.
- No Layer 3 processing takes place on packets ingressing on a T-UNI port.
- Pause frames do not switch through the T-UNI I-SID.
- You can map more than one *Transparent Port UNI* port to same ISID.
- A T-UNI port can be a part of only one ISID.
- Static MAC is not supported for a T-UNI port.
- An ISID mapped to a T-UNI service must not be mapped to any other service, such as L2 VSN and L3 VSN, on any of the remote BEBs in the SPBM network.

Use *Transparent Port UNI* when either of the following apply:

- All tagged and untagged traffic on a port must be classified into the same broadcast domain.
- You want to offer a transparent provider solution.

An example of an application for *Transparent Port UNI* is a typical Ethernet provider deployment with port-based classification and transparent forwarding.

QoS re-marking on a Transparent Port UNI

A *Transparent Port UNI* port is normally configured as a Layer 2 trusted port. The T-UNI port honors incoming customer 802.1p bits and derives an internal QoS level. The 802.1p bit marking of the Backbone VLAN (BVLAN) is derived from the internal QoS level. If the T-UNI port is set as a Layer 2 untrusted port, a best-effort queue is assigned. Customer packet headers are not modified.

The T-UNI port QoS configurations are:

- DiffServ = disable
- Layer3Trusted = access

QoS considerations when a port is associated with a T-UNI I-SID

- You cannot configure `access-diffserv` and `enable diffserv` on a T-UNI port.
- When a port is associated with a T-UNI ISID, the T-UNI QoS configuration automatically takes effect.
- When the port is removed from the T-UNI ISID, the default port QoS configuration takes effect.

QoS considerations when an MLT is associated with a T-UNI I-SID

- When an MLT, static or LACP, is added to a T-UNI ISID, the T-UNI QoS configuration take effect on all the ports of the MLT.
- When an MLT, static or LACP, is removed from a T-UNI ISID, the port default QoS configuration is configured on all the member ports of the MLT.
- If a port is added dynamically to a T-UNI MLT, static or LACP, the port inherits the QoS properties of the T-UNI MLT ports.
- If a port is dynamically removed from a T-UNI MLT, static or LACP, the port retains the QoS configuration inherited from the MLT.

Transparent Port UNI over vIST

Virtual IST (vIST) provides the ability to dual-home hosts, servers and other network devices to a pair of Multi-Chassis Link Aggregation (MC-LAG) enabled devices. The MC-LAG nodes appear to the connected devices as one link-aggregated group. So, although the physical connection is spread between two individual network nodes, logically they appear as a single connection.

Transparent Port UNI (T-UNI) over vIST peers extends the capability of dual-home hosts on the SPB cloud to achieve higher network resiliency. The MACs learnt on the T-UNI interface of any one vIST peer is synchronized with the other peer through MAC synchronization.

In the following figure, the T-UNI access switch ACCESS-1 is dual-homed into vIST peer hosts VIST-PEER 1 and VIST-PEER 2. At ACCESS-1, a link aggregation is created to connect to the SPBM cluster. On the VIST peers, an SMLT is created towards ACCESS-1. Depending on the link aggregation hashing logic, traffic is hashed on to VIST-PEER 1 and VIST-PEER 2. The MACs learnt on the T-UNI interfaces of either host is synchronized with the other host.




Figure 39: Example of Transparent Port UNI over vIST

If one of the links between ACCESS-1 and the vIST cluster goes down, all traffic is serviced through the other link. The same applies when any of the vIST peers go down. Since MAC learning on both peers are synchronized, both peers can switch traffic with the same efficiency.

Single-homed T-UNI service on a vIST-enabled node

If you configure a T-UNI service as a single-homed service on a vIST-enabled node, you must configure the same ISID service without port/MLT being mapped to ISID, on the other vIST peer node. Failure to perform this configuration on the vIST peer node can result in the loss of traffic to the single-homed T-UNI service in various scenarios.

Switched UNI

Switched User Network Interface (S-UNI) allows the association of local endpoints to I-SIDs based on local port and VLAN together. With switched UNI, the same VLAN can be used on one port to create an endpoint to one I-SID, and on another port to create an endpoint to another I-SID.

Switched UNI summary:

- S-UNI is a VLAN and ports associated with I-SIDs.
- Local significance on the ports.
- You can re-use the same VLAN to associate different ports with different I-SIDs.
- You can use a different VLAN to the same ports, or you can assign different ports to the same I-SID.
- Supports VLAN mapping on the local switch.
- To accept untagged traffic, the port needs to be configured as untagged-traffic in the I-SID.

Use Switched UNI when either of the following apply:

- Vlan ID (VID) reuse is required. The same VID is used on different broadcast domains (multi-tenant applications).
- Multiple VLANs must be part of the same broadcast domain.
- VID translation is required.

An example of an application for Switched UNI is a typical host and provider deployment, with a port and VID-based classification.

S-UNI based MAC learning

MAC learning is done on I-SID MAC. When a packet ingresses on a port or MLT which is associated with S-UNI I-SID, the system performs MAC look up based on the I-SID. S-UNI operates on Any-To-Any (ELAN) mode, there can be one or more ports associated to a S-UNI I-SID. A packet that ingresses to a S-UNI port on a BEB can transfer through the SPBM cloud, or can egress out another S-UNI port configured to the same I-SID.

When a packet ingresses an NNI port, before egressing a S-UNI port, the system performs a MAC Destination Address (DA) lookup based on the I-SID. If the DA lookup fails, the packet floods to all S-UNI ports in the I-SID.

Considerations

Consider the following when you configure a Switched UNI:

- The VLAN tag is removed before the traffic egresses out on the untagged-traffic port or MLT.
- VLAN priority received on the packet is maintained across VLAN IDs.
- Spanning tree is disabled on all S-UNI ports, and the ports remain in forwarding state.

- The S-UNI I-SID is advertised to the SPBM cloud.
- The Broadcast and unknown Unicast packets are flooded to all ports in the I-SID.

Limitations

- You cannot change from one UNI type to another dynamically. The I-SID has to be deleted and created with new UNI type (Customer VLAN (C-VLAN), Transparent port user-network-interface (T-UNI), ELAN).
- I-SID cannot be used by IPVPN, MVPN, SPBM dynamic multicast range, or *Transparent Port UNI*.
- If the port is a member of MLT, the entire MLT has to be added to the VID.
- The port is always in the forwarding state.
- The same VID, port, or MLT cannot be member of more than one I-SID.
- Static MAC, Static ARP and static IGMP group are not supported on S-UNI enabled ports.

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 40: 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 41: 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 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 42: Shortest path tree for source node A




Figure 43: Shortest path tree for source node B




Figure 44: 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 Fabric Connect

IP Multicast over Fabric Connect 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 Fabric Connect 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 Fabric Connect services. No explicit configuration exists separate from that to enable Layer 2 VSN IP Multicast over Fabric Connect.

Multicast traffic remains in the same Layer 2 VSN across the SPBM cloud for Layer 2 VSN IP Multicast over Fabric Connect. IP Multicast over Fabric Connect 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 Fabric Connect 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 Fabric Connect, 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 the CLI

This section provides procedures to configure Layer 2 VSNs using the CLI.

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.

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.

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 <1-4059> <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.

The switch reserves I-SID 0x00ffffff. The switch 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

```
Switch> enable
Switch# configure terminal
Switch:1(config)# vlan i-sid 5 5
Switch:1(config)# show vlan i-sid
```

```
=====
                        Vlan I-SID
=====
VLAN_ID    I-SID
-----
1          2
```

```
5           5
10          20
```

Variable definitions

Use the data in the following table to use the `vlan i-sid` command.

Variable	Value
<code>vlan i-sid <1-4059> <0-16777215></code>	<p>Specifies the customer VLAN (CVLAN) to associate with the I-SID.</p> <p>Use the <code>no</code> or <code>default</code> options to remove the I-SID from the specified VLAN.</p> <p>* Note:</p> <p>The switch reserves I-SID 0x00ffffff. The switch 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.</p>

Displaying C-VLAN I-SID information

Use the following procedure to display C-VLAN I-SID information.

Procedure

1. Display the C-VLAN to I-SID associations:

```
show vlan i-sid <1-4059>
```

2. Display the IS-IS SPBM multicast-FIB calculation results by I-SID:

```
show isis spbm i-sid {all|config|discover} [vlan <1-4059>] [id <1-16777215>] [nick-name <x.xx.xx>]
```

3. Discover where entries are learned:

```
show vlan mac-address-entry [spbm-tunnel-as-mac]
```

4. Display the VLAN remote MAC table for a C-VLAN:

```
show vlan remote-mac-table <1-4059>
```

Example

```
Switch# show vlan i-sid
```

```
=====
                        Vlan I-SID
=====
VLAN_ID   I-SID
-----
1
2
5          5
10
20
```

```
Switch# 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         SwitchA
300    1.11.16        1000   0014.c7e1.33df  config         SwitchA
400    1.11.16        1000   0014.c7e1.33df  config         SwitchA
200    1.11.16        2000   0014.c7e1.33df  config         SwitchA
300    1.11.16        2000   0014.c7e1.33df  config         SwitchA
400    1.11.16        2000   0014.c7e1.33df  config         SwitchA
200    1.12.45        1000   0016.ca23.73df  discover       SwitchA
300    1.12.45        1000   0016.ca23.73df  discover       SwitchA
=====

Total number of SPBM ISID entries configed: 6
=====
Total number of SPBM ISID entries discovered: 2
=====
Total number of SPBM ISID entries: 8
=====

```

```

Switch:# 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-1/9           false SwitchB
1       learned          00:80:2d:22:ac:46  Port-1/15          false SwitchB
2       self             a4:25:1b:51:48:84  103.103.103.103    false -
2       self             02:01:03:ff:ff:ff  Tunnel_to_HQ       false -
5       learned          00:00:00:00:00:1a  access             false SwitchB
10      self             00:00:00:00:49:50  Port-1/9           false -
10      self             00:00:00:00:50:50  Port-1/9           false -
=====

```

```

Switch# 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   Switch-6005  MLT-2      false
100  learned b4:a9:5a:04:c8:83  b4:a9:5a:04:c8:65  3    Switch-174   103.103.103.103 true
100  learned b4:a9:5a:04:c8:84  b4:a9:5a:04:c8:66  3    Switch-175   Tunnel_to_HQ true
=====
3 of 3 matching entries out of total of 3 Remote Mac entries in all fdb(s) displayed.
=====

```

Variable definitions

Use the data in the following table to use the **show vlan** commands.

Variable	Value
i-sid <1-4059>	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

Table continues...

Variable	Value
	entry is not learned in the SPBM network, the TUNNEL column will be empty (-).
remote-mac-table <1-4059>	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}	<ul style="list-style-type: none"> all: displays all I-SID entries config: displays configured I-SID entries discover: displays discovered I-SID entries
vlan <1-4059>	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: <ul style="list-style-type: none"> • 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 vIST peer.
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: <ul style="list-style-type: none"> • 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.

Table continues...

Parameter	Description
PORTS	Either displays the value 0 or indicates the port in which a frame comes from.
SMLT REMOTE	Indicates the MAC address entry for the remote vIST peer.

Configuring Layer 2 VSN IP Multicast over Fabric Connect

Use this procedure to configure IP Multicast over Fabric Connect for Layer 2 VSN functionality. With Layer 2 VSN IP Multicast over Fabric Connect, 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 IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect services. No explicit configuration exists separate from that to enable Layer 2 VSN IP Multicast over Fabric Connect.

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.

The switch only supports IPv4 multicast traffic.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4059>
```

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 Fabric Connect 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 Fabric Connect route information:


```
show isis spbm ip-multicast-route [all]
```
3. Display detailed IP Multicast over Fabric Connect route information:

```
show isis spbm ip-multicast-route [detail]
```

4. Display IP multicast route information by VLAN:

```
show isis spbm ip-multicast-route [vlan <1-4059>]
```

5. Display IP Multicast over Fabric Connect route information by VSN I-SID:

```
show isis spbm ip-multicast-route [vsn-isid <1-16777215>]
```

6. Display IP Multicast over Fabric Connect route information by group address:

```
show isis spbm ip-multicast-route [group {A.B.C.D}]
```

7. Display IP Multicast over Fabric Connect route information by source address:

```
show isis spbm ip-multicast-route [source {A.B.C.D}]
```

! Important:

When you use the command `show isis spbm ip-multicast-route` without parameters or use the detail or group optional parameters without specifying a VLAN ID or VSN-ISID, the command output displays Layer 3 context only. No Layer 2 context is displayed.

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

```
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    e12
snoop  GRT      501  192.0.2.1  233.252.0.2  5010      16300002  20    e12
snoop  GRT      501  192.0.2.1  233.252.0.3  5010      16300003  10    e12
snoop  GRT      501  192.0.2.1  233.252.0.4  5010      16300004  20    e12
snoop  GRT      501  192.0.2.1  233.252.0.5  5010      16300005  10    e12
snoop  GRT      501  192.0.2.1  233.252.0.6  5010      16300006  20    e12
snoop  GRT      501  192.0.2.1  233.252.0.7  5010      16300007  10    e12
snoop  GRT      501  192.0.2.1  233.252.0.8  5010      16300008  20    e12
snoop  GRT      501  192.0.2.1  233.252.0.9  5010      16300009  10    e12
snoop  GRT      501  192.0.2.1  233.252.0.10 5010      16300010  20    e12
-----
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    e12
snoop  GRT      501  192.0.2.1  233.252.0.2  5010      16300002  20    e12
```

SPBM and IS-IS services configuration

```
snoop GRT 501 192.0.2.1 233.252.0.3 5010 16300003 10 e12
snoop GRT 501 192.0.2.1 233.252.0.4 5010 16300004 20 e12
snoop GRT 501 192.0.2.1 233.252.0.5 5010 16300005 10 e12
snoop GRT 501 192.0.2.1 233.252.0.6 5010 16300006 20 e12
snoop GRT 501 192.0.2.1 233.252.0.7 5010 16300007 10 e12
snoop GRT 501 192.0.2.1 233.252.0.8 5010 16300008 20 e12
snoop GRT 501 192.0.2.1 233.252.0.9 5010 16300009 10 e12
snoop GRT 501 192.0.2.1 233.252.0.10 5010 16300010 20 e12
```

```
-----
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
-----
```

```
Switch:1# show isis spbm ip-multicast-route vsn-isis 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	e12
192.0.2.1	233.252.0.2	16300002	20	e12
192.0.2.1	233.252.0.3	16300003	10	e12
192.0.2.1	233.252.0.4	16300004	20	e12
192.0.2.1	233.252.0.5	16300005	10	e12
192.0.2.1	233.252.0.6	16300006	20	e12
192.0.2.1	233.252.0.7	16300007	10	e12
192.0.2.1	233.252.0.8	16300008	20	e12
192.0.2.1	233.252.0.9	16300009	10	e12
192.0.2.1	233.252.0.10	16300010	20	e12

```
-----
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
-----
```

```
Switch:1# show isis spbm ip-multicast-route vsn-isis 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	1/3	V501:9/38	e12
192.0.2.1	233.252.0.2	16300002	20	1/2,1/3	V501:9/38	e12
192.0.2.1	233.252.0.3	16300003	10	1/3	V501:9/38	e12
192.0.2.1	233.252.0.4	16300004	20	1/2,1/3	V501:9/38	e12
192.0.2.1	233.252.0.5	16300005	10	1/3	V501:9/38	e12
192.0.2.1	233.252.0.6	16300006	20	1/2,1/3	V501:9/38	e12
192.0.2.1	233.252.0.7	16300007	10	1/3	V501:9/38	e12
192.0.2.1	233.252.0.8	16300008	20	1/2,1/3	V501:9/38	e12
192.0.2.1	233.252.0.9	16300009	10	1/3	V501:9/38	e12
192.0.2.1	233.252.0.10	16300010	20	1/2,1/3	V501:9/38	e12

```
-----
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 e12
5010	192.0.2.1	233.252.0.3	16300003	10	0x0 e12
5010	192.0.2.1	233.252.0.5	16300005	10	0x0 e12

```

5010    192.0.2.1    233.252.0.7    16300007  10    0x0  e12
5010    192.0.2.1    233.252.0.9    16300009  10    0x0  e12
5010    192.0.2.1    233.252.0.2    16300002  20    0x0  e12
5010    192.0.2.1    233.252.0.4    16300004  20    0x0  e12
5010    192.0.2.1    233.252.0.6    16300006  20    0x0  e12
5010    192.0.2.1    233.252.0.8    16300008  20    0x0  e12
5010    192.0.2.1    233.252.0.10   16300010  20    0x0  e12

```

```
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	1/4,MLT-35	V501:9/32-9/33	e12
192.0.2.1	233.252.0.3	16300002	20	-	V501:9/32-9/33	e12
192.0.2.1	233.252.0.5	16300003	10	1/4,MLT-35	V501:9/32-9/33	e12
192.0.2.1	233.252.0.7	16300004	20	-	V501:9/32-9/33	e12
192.0.2.1	233.252.0.9	16300005	10	1/4,MLT-35	V501:9/32-9/33	e12
192.0.2.1	233.252.0.2	16300006	20	-	V501:9/32-9/33	e12
192.0.2.1	233.252.0.4	16300007	10	1/4,MLT-35	V501:9/32-9/33	e12
192.0.2.1	233.252.0.6	16300008	20	-	V501:9/32-9/33	e12
192.0.2.1	233.252.0.8	16300009	10	1/4,MLT-35	V501:9/32-9/33	e12
192.0.2.1	233.252.0.10	16300010	20	-	V501:9/32-9/33	e12

```
-----
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 Fabric Connect route information.
detail	Displays detailed IP Multicast over Fabric Connect route information.
group {A.B.C.D} source {A.B.C.D}	Displays information on the group IP address for the IP Multicast over Fabric Connect route. If you select source it will also display the source IP address.
vlan <0-4084>	Displays IP Multicast over Fabric Connect route information by VLAN.
vrf WORD<1-16>	Displays IP Multicast over Fabric Connect route information by VRF.
vsn-isid <1-16777215>	Displays IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect summary for a given host-name.

Table continues...

Variable	Value
lspid <xxxx.xxxx.xxxx.xx-xx>	Displays the IP Multicast over Fabric Connect 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: <ul style="list-style-type: none"> • 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 Fabric Connect route.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
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 Fabric Connect because IP Shortcuts IP Multicast over Fabric Connect does not use a VSN I-SID.
Data ISID	Specifies the data I-SID for the IP Multicast over Fabric Connect 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 Fabric Connect route.
Source-BEB	Specifies the source BEB for the IP Multicast over Fabric Connect 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 Fabric Connect route.

Table continues...

Parameter	Description
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
Data ISID	Specifies the data I-SID for the IP Multicast over Fabric Connect route. After a BEB receives the IP Multicast over Fabric Connect 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 Fabric Connect route.
Source-BEB	Specifies the source BEB for the IP Multicast over Fabric Connect 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.
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.

Table continues...

Parameter	Description
SOURCE ADDRESS	Specifies the group IP address for the IP Multicast over Fabric Connect route.
GROUP ADDRESS	Specifies the group IP address for the IP Multicast over Fabric Connect route.
DATA I-SID	Specifies the data I-SID for the IP Multicast over Fabric Connect route. After a BEB receives the IP Multicast over Fabric Connect 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[/sub-port]}[-slot/
port[/sub-port]][,...]] [vlan <1-4059>[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

```
Switch:#enable
```

```
Switch:1#show ip igmp interface
```

```
=====
                        Igmp Interface - GlobalRouter
=====
```

IF	QUERY INTVL	STATUS	VERS.	OPER VERS	QUERIER	QUERY MAXRSPT	WRONG QUERY	JOINS	ROBUST	LASTMEM 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 ID	QUERY INTVL	QUERY MAX RESP	ROBUST	VERSION	LAST MEMB QUERY	PROXY SNOOP ENABLE	SNOOP ENABLE	SSM SNOOP ENABLE	FAST LEAVE ENABLE	FAST LEAVE PORTS
1	125	100	2	2	10	false	false	false	false	

VLAN ID	SNOOP QUERIER ENABLE	SNOOP QUERIER ADDRESS	DYNAMIC DOWNGRADE VERSION	COMPATIBILITY MODE	EXPLICIT HOST TRACKING
1	false	0.0.0.0	enable	disable	disable

```
Switch:1# show ip igmp sender
```

```
=====
                        IGMP Sender - GlobalRouter
=====
```

GRPADDR	IFINDEX	MEMBER	PORT/ 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
---------	--------	--------	------------	------

SPBM and IS-IS services configuration

```
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
=====
GRPADDR      IFINDEX      MEMBER      PORT/
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  1/7,1/9 NETWORK
233.252.0.2  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.3  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.4  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.5  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.6  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.7  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.8  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.9  192.2.0.1  501  spb  501  1/7,1/9 NETWORK
233.252.0.10 192.2.0.1  501  spb  501  1/7,1/9 NETWORK
```

Variable definitions

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

Variable	Value
gigabitethernet {slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vlan <1-4059>	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.

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.

Table continues...

Variable	Value
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.
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

Table continues...

Parameter	Description
	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 Fabric Connect 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.
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.

Table continues...

Parameter	Description
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.

Table continues...

Parameter	Description
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.

The following table describes the fields for the `show ip igmp snoop-trace` command.

Parameter	Description
GROUP ADDRESS	Indicates the IP multicast group address for which this entry contains information.
SOURCE ADDRESS	Indicates the source of the multicast traffic.
IN VLAN	Indicates the incoming VLAN ID.
IN PORT	Indicates the incoming port number.
OUT VLAN	Indicates the outgoing VLAN ID.
OUT PORT	Indicates the outgoing port number.
TYPE	Indicates where the stream is learned. ACCESS indicates the stream is learned on UNI ports. NETWORK indicates the stream is learned over the SPBM network.

Viewing TLV information for Layer 2 VSN IP Multicast over Fabric Connect

Use the following commands to check TLV information.

For Layer 2 VSN with IP multicast over Fabric Connect, 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 <1-236> [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 <1-236> [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 {1, 12, 112}	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 (112) function is disabled.
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.

Table continues...

Variable	Value
tlv <1–236>	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.

Configuring an SPBM Layer 2 Transparent Port UNI

Transparent Port UNI (T-UNI) maps a port or an MLT to an I-SID. T-UNI configures a transparent port where all traffic is MAC switched on an internal virtual port using the assigned I-SID. Multiple ports on the same unit and on other Backbone Edge Bridges (BEBs) are switched on a common I-SID. No VLAN is involved in this process. The T-UNI port is not a member of any VLAN or STG.

Consider the following design requirements when implementing this feature:

- Only E-LAN based T-UNI is supported. All T-UNI I-SID end points become members of the same shared E-LAN service. If an E-LINE type of service is required, provision T-UNI at the two end points comprising the point-to-point service.
- A port or MLT associated with a T-UNI I-SID cannot be part of any VLAN.
- Any Spanning Tree Protocol implementation is disabled on the port or MLT associated with the T-UNI I-SID. The port will always be in a Forwarding state.
- MACs are learned against the combination of the I-SID and port or MLT.
- Multiple ports or MLTs can be associated with same T-UNI I-SID.
- One port or MLT cannot be part of multiple T-UNI I-SIDs.

- No additional IS-IS TLVs are added to advertise or withdraw T-UNI I-SID services. Avaya makes use of the existing IS-IS TLV-144 and sub TLV-3 to carry I-SID information.
- The MAC address limit is supported on a per-I-SID basis. For example, the MAC addresses learned on the T-UNI I-SID can be limited.
- IP traffic and control packets are transparently bridged over T-UNI endpoints.
- Untagged traffic ingressing on the T-UNI port will use port QoS. B-TAG and I-TAG priorities are derived from the port QoS.
- The 802.1p bits of the incoming traffic are used to derive the B-TAG and I-TAG priorities for tagged traffic.
- LACP and VLACP PDUs are extracted to the CP and all other control packets are transparently bridged over the T-UNI port or MLT.

This feature handles control PDUs in the following manner:

- All the Layer 2 and Layer 3 control packets are transparently bridged over the T-UNI port or MLT with the exception of LACP and VLACP PDUs. LACP PDUs and VLACP PDUs are not transparently bridged over the T-UNI port or MLT if LACP or VLACP is enabled on the port or MLT.
 - If an LACP MLT is associated with a T-UNI I-SID, LACP PDUs are extracted to CP and processed locally.
 - If LACP is not enabled globally and LACP MLT is not associated with the T-UNI I-SID, LACP PDUs are transparently bridged across the T-UNI port or MLT.
 - If a VLACP enabled port is added to a T-UNI I-SID, VLACP PDUs are extracted to the CP for local processing. If a port that is not VLACP enabled is added to the T-UNI I-SID, VLACP PDUs are transparently bridged across T-UNI port.
- The following list of control packet types are transparently bridged across the T-UNI I-SID:
 - SLPP
 - VRRP
 - OSPF
 - RIP
 - BGP
 - ISIS
 - CFM
 - STP
 - SONMP

*** Note:**

If you are configuring a T-UNI to terminate on a port or MLT on a switch in a vIST switch cluster, you must also configure the T-UNI I-SID on the other switch of the vIST switch cluster. You

must configure the T-UNI I-SID on both switches of a vIST pair. It is not necessary to assign an actual port or MLT to the T-UNI on the second switch.

Use this procedure to configure a Transparent Port UNI or Elan-Transparent based service.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes creating the SPBM BVLANS.
- You must associate a T-UNI LACP MLT with a VLAN before mapping the LACP MLT to a T-UNI I-SID.

Caution:

In the case of T-UNI LACP SMLT, before you configure SMLT on switch peers, ensure that the T-UNI LACP MLT on each peer is always associated with a VLAN, even if it is the default VLAN, and that it is added to a T-UNI I-SID. Otherwise, traffic is not forwarded on the T-UNI LACP MLT.

About this task

You can configure Transparent Port UNI when either of the following apply:

- You want all tagged and untagged traffic on a port to be classified into the same broadcast domain.
- You want to offer a transparent provider solution.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure a Transparent Port UNI (Elan-Transparent based service). Enter:

```
i-sid <1-16777215> elan-transparent
```

This command automatically takes you to the Elan-Transparent I-SID Configuration mode.

3. Add ports to the Elan-Transparent based service. Enter:

```
port {slot/port[/sub-port] [-slot/port[/sub-port]][, ...]}
```

A warning message displays indicating that adding a port to a T-UNI I-SID removes the port from all VLANs. Click *y* when prompted, to continue.

4. Add an MLT to the Elan-Transparent based service. Enter:

```
mlt <1-512>
```

A warning message displays indicating that adding an MLT to a Transparent Port UNI I-SID removes the MLT from all VLANs. Click *y* when prompted, to continue.

5. To verify the Transparent Port UNI configuration, enter:

```
show i-sid <1-16777215>
```

6. To remove ports or MLT from the Elan-Transparent based service, enter one of the following commands:

```
no port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}
```

OR

```
no mlt <1-512>
```

7. To delete the Elan-Transparent based service, enter:

```
no i-sid <1-16777215>
```

Example

Configure a Transparent Port UNI I-SID (elan-transparent based service).

```
Switch:1(config)#i-sid 300 elan-transparent

Switch:1(elan-tp:300)#port 1/25
Adding Ports to Transparent UNI i-sid removes it from all VLANS.
Do you wish to continue (y/n) ? y
Switch:1(elan-tp:300)#

Switch:1(elan-tp:300)#mlt 1
Adding MLTs to Transparent UNI i-sid removes it from all VLANS.
Do you wish to continue (y/n) ? y
Switch:1(elan-tp:300)#
```

Verify Transparent Port UNI or Elan-Transparent based service configuration.

```
Switch:1(elan-tp:300)#show i-sid 300
=====
                                Isid Info
=====
ISID      ISID      VLANID    PORT      MLT
ID        TYPE                               INTERFACES INTERFACES
-----
200       CVLAN     200       1/25      1
300       ELAN_TR   N/A       1/25      1
```

Variable definitions

Use the data in the following table to use the `i-sid` command.

* Note:

When SPB is enabled, ISID IDs 16777216 and greater are reserved for dynamic data-isid used to carry Multicast traffic over SPB.

Variable	Value
<code>i-sid <1-16777215> elan-transparent</code>	Creates an Elan-Transparent based service. The service interface identifier (I-SID) range is 1 to 16777215.
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}</code>	Add ports to the Elan-Transparent based service.
<code>mlt<1-512></code>	Add MLTs to the Elan-Transparent based service. The MLT range is 1 to 512.

Viewing all configured I-SIDs

Perform this procedure to view all the configured I-SIDs including their types, ports, and MLTs.

About this task

View all configured I-SIDs (both CVLAN and T-UNI). View also the I-SID types and the ports or MLTs that are assigned to each I-SID.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View all configured I-SIDs. This command displays both CVLAN and T-UNI based I-SIDs.

```
show i-sid
```

3. View all T-UNI (Elan-Transparent) I-SIDs.

```
show i-sid [elan-transparent]
```

4. View information for a particular T-UNI I-SID.

```
show i-sid [<1-16777215>]
```

5. View all IS-IS SPBM I-SID information by I-SID ID:

```
show isis spbm i-sid {all|config|discover} [vlan <2-4059>] [id <1-16777215>] [nick-name <x.xx.xx>]
```

Example

View all configured I-SIDs.

```
Switch:1(config)#show i-sid
=====
Isid Info
=====
ISID      ISID      VLANID    PORT      MLT
ID        TYPE                               INTERFACES INTERFACES
-----
6         ELAN_TR   N/A       1/2       6
12        CVLAN     11        1/7
34        ELAN_TR   4450      1/17
100       ELAN_TR   7254      1/12

All 4 out of 4 Total Num of i-sids displayed
Switch(config)#
```

View T-UNI (ELAN Transparent) I-SIDs.

```
Switch:1(config)#show i-sid elan-transparent
=====
Isid Info
=====
ISID      ISID      VLANID    PORT      MLT
ID        TYPE                               INTERFACES INTERFACES
-----
100       ELAN_TR   N/A       1/12
6         ELAN_TR   N/A       1/2
```

```
All 2 out of 2 Total Num of elan-tp i-sids displayed
```

View MLT or port information for a particular T-UNI I-SID.

```
Switch:1(config)#show i-sid 6
```

```
=====
                                Isid Info
=====
ISID      ISID      VLANID    PORT      MLT
ID        TYPE                               INTERFACES INTERFACES
-----
6         ELAN_TR   N/A       1/2       6
```

View all IS-IS SPBM I-SID information:

```
Switch:1#show isis spbm i-sid all
```

```
=====
                                SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID                               TYPE      HOST_NAME
-----
100      1.11.16     20    0014.c7e1.33df                       config    Switch1
6        1.11.20     10    0014.c723.67df                       discover  Switch2
-----
Total number of SPBM ISID entries configured: 1
-----
Total number of SPBM ISID entries discovered: 1
-----
Total number of SPBM ISID entries: 2
-----
```

View all IS-IS SPBM I-SID information by I-SID ID:

```
Switch:1#show isis spbm i-sid all id 300
```

```
=====
                                SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID                               TYPE      HOST_NAME
-----
300      7.15.16     20    a425.1b51.9484                       config    Switch1
300      4.01.18     10    b4a9.5a2a.d065                       discover  Switch2
-----
Total number of SPBM ISID entries configured: 1
-----
Total number of SPBM ISID entries discovered: 1
-----
Total number of SPBM ISID entries: 2
-----
```

Variable definitions

Use the data in the following table to use the `show i-sid` command.

* Note:

When SPB is enabled, I-SID IDs 16777216 and greater are reserved for dynamic data I-SIDs, used to carry Multicast traffic over SPB.

Variable	Value
<1-16777215>	Specifies the service interface identifier (ISID).
elan-transparent	Displays only all the Elan-Transparent (T-UNI based) ISIDs.
spbm i-sid {all config discover}	<ul style="list-style-type: none"> • all: displays all I-SID entries • config: displays configured I-SID entries • discover: displays discovered I-SID entries
vlan <2-4059>	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 table describes the fields in the output for the `show i-sid` command.

Table 11: show i-sid

Field	Description
ISID ID	Specifies the service interface identifier (I-SID)
ISID TYPE	Specifies the type of I-SID
VLANID	Specifies the backbone VLAN
PORT INTERFACES	Specifies the port that is assigned to the I-SID
MLT INTERFACES	Specifies the mlt that is assigned to the I-SID

The following describes the fields in the output for the `show isis spbm i-sid` command.

Table 12: show isis spbm i-sid

Field	Description
ISID	Indicates the IS-IS SPBM I-SID identifier.
SOURCE NAME	<p>Indicates the nickname of the node where this I-SID was configured or discovered.</p> <p> Note: SOURCE NAME is equivalent to nickname.</p>
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.

Viewing C-MACs learned on T-UNI ports for an ISID

Perform this procedure to view the I-SID bridge forwarding database.

About this task

The `show i-sid mac-address-entry` command displays the C-MACs learned on T-UNI I-SIDs. It also displays the C-MACs learned on T-UNI I-SIDs for a specific I-SID, MAC address, port or port list or remote MAC address.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View C-MACs learned on the T-UNI I-SIDs:

```
show i-sid mac-address-entry [<1-16777215>] [mac
<0x00:0x00:0x00:0x00:0x00:0x00>] [port {slot/port[/sub-port]} [-slot/
port[/sub-port]] [,...]] [remote]
```

Example

View C-MACs learned on all T-UNI I-SIDs.

```
Switch:1#show i-sid mac-address-entry
```

```
=====
I-SID Fdb Table
=====
I-SID  STATUS  MAC-ADDRESS  INTERFACE  TYPE  DEST-MAC  BVLAN  DEST-SYSNAME
-----
100    learned  cc:f9:54:ae:28:81  Port-1/16  LOCAL  00:00:00:00:00:00  0
4      learned  cc:f9:54:ae:2c:18  mlt-6      LOCAL  00:00:00:00:00:00  0
252    learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0  128  DIST-1B
```

```
All 3 out of 3 Total Num of i-sid FDB Entries displayed
```

View C-MACs learned on a specific T-UNI I-SID.

```
Switch:1#show i-sid mac-address-entry 100
```

```
=====
I-SID Fdb Table
=====
I-SID  STATUS  MAC-ADDRESS  INTERFACE  TYPE  DEST-MAC  BVLAN  DEST-SYSNAME
-----
100    learned  cc:f9:54:ae:28:81  Port-1/16  LOCAL  00:00:00:00:00:00  0
```

```
All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

```
Switch:1#show i-sid mac-address-entry 252
```

```
=====
I-SID Fdb Table
=====
I-SID  STATUS  MAC-ADDRESS  INTERFACE  TYPE  DEST-MAC  BVLAN  DEST-SYSNAME
-----
252    learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0  128  DIST-1B
```

```
All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View C-MACs learned on a T-UNI I-SID for a specific MAC address.

```
Switch:1#show i-sid mac-address-entry mac cc:f9:54:ae:38:64
```



```

=====
I-SID Fdb Table
=====
I-SID STATUS  MAC-ADDRESS      INTERFACE  TYPE    DEST-MAC          BVLAN  DEST-SYSNAME
-----
252   learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0  128    DIST-1B

All 1 out of 1 Total Num of i-sid FDB Entries displayed
    
```

View C-MACs learned on aT-UNI I-SID for a specific port.

```

Switch:1#show i-sid mac-address-entry port 1/15

=====
I-SID Fdb Table
=====
I-SID STATUS  MAC-ADDRESS      INTERFACE  TYPE    DEST-MAC          BVLAN  DEST-SYSNAME
-----
252   learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0  128    DIST-1B

All 1 out of 1 Total Num of i-sid FDB Entries displayed
    
```

View C-MACs learned on a T-UNI I-SID as a remote MAC address.

```

Switch:1#show i-sid mac-address-entry remote

=====
I-SID Fdb Table
=====
I-SID STATUS  MAC-ADDRESS      INTERFACE  TYPE    DEST-MAC          BVLAN  DEST-SYSNAME
-----
252   learned  cc:f9:54:ae:38:64  Port-1/15  REMOTE  00:13:0a:0c:d3:e0  128    DIST-1B

All 1 out of 1 Total Num of i-sid FDB Entries displayed
    
```

Variable definitions

Use the data in the following table to use the **show i-sid mac-address-entry** command.

Variable	Value
<1-16777215>	Displays the MAC address learned on the service interface identifier (ISID).
mac <0x00:0x00:0x00:0x00:0x00:0x00>	Displays the I-SID FDB details for the specified MAC address.
port {slot/port[/sub-port] [-slot/port[/sub-port]] [...]}	Displays the MAC address learned on the specified port or port list.
remote	Displays the remote MAC address learned on the I-SID.

Job aid

The following table describes the fields in the output for the **show i-sid mac-address-entry** command.

Table 13: show i-sid

Field	Description
I-SID	Specifies the service interface identifier (I-SID).
STATUS	Specifies the learning status of the associated MAC.
MAC-ADDRESS	Specifies the MAC address of the port assigned to the specific I-SID or MAC learned on the specific I-SID.
INTERFACE	Specifies the port or MLT on which the MAC is learned for the specific I-SID.
TYPE	Specifies whether the MAC is a Local or IST PEER or a Remote MAC.
DEST-MAC	Specifies the virtual BMAC address or system ID, in MAC format, of the destination node.
BVLAN	Specifies the BVLAN on which the destination node is discovered for the I-SID.
DEST-SYSNAME	Specifies the destination system name.

Viewing I-SID maximum MAC-limit

Perform this procedure to view the maximum MAC learning limit information for an I-SID.

Important:

The command `show i-sid limit-fdb-learning` is supported only on the Avaya Virtual Services Platform 4000 Series.

About this task

The total MAC learning limit per switch is 32000. MAC learning on I-SID stops when the maximum limit is reached.

Procedure

View the maximum MAC learning limit configured for an I-SID:

```
show i-sid limit-fdb-learning <1-16777215>
```

Example

View maximum MAC learning limit for all I-SIDs.

```
Switch:1#show i-sid limit-fdb-learning
```

```
=====
                Isid MAC-Limit Info
=====
ISID      MAC-LIMIT  MAXMAC
ID        STATUS    COUNT
-----
10        disabled   32000
11        disabled   32000
12        disabled   32000
15        disabled   32000
```

```
101          disabled    32000
All 5 out of 5 Total Num of i-sid Info displayed
```

View maximum MAC learning limit for a specific I-SID.

```
Switch:1#show i-sid limit-fdb-learning 10

=====
          Isid MAC-Limit Info
=====
ISID      MAC-LIMIT  MAXMAC
ID        STATUS    COUNT
-----
10        disabled   32000
All 1 out of 1 Total Num of i-sid Info displayed
```

Variable definitions

Use the data in the following table to use the `show i-sid limit-fdb-learning` command.

! Important:

The command `show i-sid limit-fdb-learning` is supported only on the Avaya Virtual Services Platform 4000 Series.

Variable	Value
limit-fdb-learning	Displays the I-SID-based maximum MAC limit information.
<1-6777215>	Displays the service interface identifier (ISID). The ISID range is 1 to 16777215.

Job aid

The following table describes the fields in the output of the `show i-sid limit-fdb-learning` command.

! Important:

The command `show i-sid limit-fdb-learning` is supported only on the Avaya Virtual Services Platform 4000 Series.

Table 14: show i-sid limit-fdb-learning

Field	Description
ISID ID	Specifies the service interface identifier (ISID)
MAC-LIMIT STATUS	Specifies whether the MAC learning limit is enabled or disabled
MAXMAC COUNT	Specifies the MAC learning limit

Configuring an SPBM Layer 2 Switched UNI on an MLT

Shortest Path Bridging MAC (SPBM) supports Layer 2 Virtual Service Network (VSN) functionality where Switched UNIs are bridged over the SPBM core infrastructure.

Switched User Network Interface (S-UNI) allows the association of local endpoints to I-SIDs based on local port and VLAN together. With switched UNI, the same VLAN can be used on one port to create an endpoint to one I-SID, and on another port to create an endpoint to another I-SID.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.

About this task

To configure a Switched UNI on an MLT, you must create a Switched UNI I-SID, and map an MLT to the Switched UNI I-SID.

Procedure

1. Enter MLT Interface Configuration mode:

```
enable
configure terminal
interface mlt <1-512>
```

2. Enable S-UNI on MLT:

```
flex-uni enable
```

*** Note:**

Spanning tree is disabled on all the Switched UNI ports.

*** Note:**

You cannot enable Switched UNI on EAPoL enabled interface.

3. Configure a Switched UNI Service Instance Identifier (I-SID):

```
i-sid <1-16777215> [elan]
```

This command automatically takes you to the Elan I-SID Configuration mode.

4. Add an MLT to a Switched UNI I-SID:

```
c-vid <1-4094> mlt <1-512>
```

C-VID is customer VLAN ID. This command maps a VLAN ID to an MLT.

*** Note:**

You can run this command again to map a Switched UNI MLT to multiple I-SIDs.

5. Display the Switched UNI information:

```
show mlt i-sid
```

Example

```
Switch> enable
Switch# configure terminal
```

SPBM and IS-IS services configuration

```
Switch(config)# mlt 10
Switch(config)# interface mlt 10
Switch(config-mlt)# flex-uni enable
Switch(config-mlt)# i-sid 100
Switch(elan:100)# c-vid 20 mlt 10
Switch(elan:100)# show mlt i-sid
```

```
=====
                        MLT Isid Info
=====
MLTID      IFINDEX  ISID      VLANID  C-VID  ISID      ORIGIN      BPDU
          ID          TYPE
-----
10         6153    100      N/A     20     ELAN     CONFIG
-----
1 out of 1 Total Num of i-sid endpoints displayed
```

Variable definitions

Use the data in the following table to use the `i-sid` command to configure a Switched UNI.

Variable	Value
<code>i-sid <1-16777215> elan</code>	Creates an Elan based service. The service interface identifier (I-SID) range is 1 to 16777215.
<code>c-vid <1-4094> mlt <mlt-id></code>	Add MLT to the Elan based service. C-VID is customer VLAN ID. This command maps a VLAN ID to an MLT.

Configuring an SPBM Layer 2 Switched UNI on a Port

Shortest Path Bridging MAC (SPBM) supports Layer 2 Virtual Service Network (VSN) functionality where Switched UNIs are bridged over the SPBM core infrastructure.

Switched User Network Interface (S-UNI) allows the association of local endpoints to I-SIDs based on local port and VLAN together. With switched UNI, the same VLAN can be used on one port to create an endpoint to one I-SID, and on another port to create an endpoint to another I-SID.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.

About this task

To configure a Switched UNI on a port, you must create a Switched UNI I-SID, and map the port to the Switched UNI I-SID.

Procedure

- Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
```

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

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable S-UNI on a port:

```
flex-uni enable
```

*** Note:**

Spanning tree is disabled on all the Switched UNI ports.

*** Note:**

You cannot enable Switched UNI on EAPoL enabled interface.

3. Configure a Switched UNI Service Instance Identifier (I-SID):

```
i-sid <1-16777215> [elan]
```

This command automatically takes you to the Elan I-SID Configuration mode.

4. Add ports to a Switched UNI I-SID:

```
c-vid <1-4094> port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}
```

5. Display the Switched UNI information:

```
show interface gigabitethernet i-sid {slot/port[/sub-port] [-slot/
port[/sub-port]] [,...]}
```

Example

```
Switch> enable
Switch# configure terminal
Switch(config)# interface gigabitethernet 1/1,1/2
Switch(config-if)# flex-uni enable
Switch(config-if)# i-sid 100
Switch(elan:100)# c-vid 10 port 1/1,1/2
Switch(elan:100)# show interface gigabitethernet i-sid
```

PORT Isid Info							
PORTNUM	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	BPDU
1/1	192	100	N/A	10	ELAN	CONFIG	
1/2	193	100	N/A	10	ELAN	CONFIG	

2 out of 3 Total Num of i-sid endpoints displayed

Variable definitions

Use the data in the following table to use the `i-sid` command to configure a Switched UNI.

Variable	Value
i-sid <1-16777215> elan	Creates an Elan based service. The service interface identifier (I-SID) range is 1 to 16777215.
c-vid <1-4094> port {slot/port[/sub-port] [-slot/port[/sub-port]][,...]}	Add ports to the Elan based service.

Viewing all configured Switched UNI I-SIDs

Perform this procedure to view all the configured S-UNI I-SIDs including their types, ports, and MLTs.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View all configured I-SIDs. This command displays CVLAN, T-UNI, and S-UNI based I-SIDs.
`show i-sid`
3. View all S-UNI I-SIDs.
`show i-sid [elan]`
4. View all associated MLT on the S-UNI I-SID.
`show mlt i-sid [MLT ID <1-512>]`
5. View all associated ports on the S-UNI I-SID.
`show interface gigabitethernet i-sid {slot/port[/sub-port] [-slot/port[/sub-port]][,...]}`
6. View all IS-IS SPBM multicast FIB entries.
`show isis spbm multicast-fib detail`

Example

View all configured I-SIDs.

Switch:1# show i-sid

```

=====
                                Isid Info
=====
ISID      ISID      VLANID   PORT      MLT      ORIGIN
ID        TYPE
-----
999      ELAN      99       -         c110:100  CONFIG
          99       1/21     -
=====
    
```

```
c: customer vid      u: untagged-traffic
```

```
All 1 out of 1 Total Num of i-sids displayed
```

View all S-UNI I-SIDs.

```
Switch:1# show i-sid elan
```

```
=====
```

Isid Info					
ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN
312	ELAN	N/A	-	c710:1	DISC_BOTH
513	ELAN	N/A	-	c513:1	DISC_BOTH
2000	ELAN	200	c20:1/3	-	CONFIG
2001	ELAN	201	-	c121:1	CONFIG
2002	ELAN	202	-	c222:2	CONFIG
2004	ELAN	204	c420:1/3	-	CONFIG
			-	c421:1	
			-	c422:2	
3000	ELAN	N/A	c30:1/3	-	CONFIG
			-	c31:1	
			-	c32:2	
3010	ELAN	N/A	c130:1/3	-	CONFIG
3020	ELAN	N/A	-	-	CONFIG
3030	ELAN	N/A	c330:1/3	-	CONFIG
3040	ELAN	N/A	-	c431:1	CONFIG
			-	u:2	
3050	ELAN	N/A	-	c531:1	DISC_BOTH
3070	ELAN	N/A	c730:1/3	-	CONFIG
			-	c731:1	
			-	c732:2	
4000	ELAN	400	c40:1/3	-	CONFIG
			-	c41:1	
			-	c42:2	
4011	ELAN	411	c140:1/3	-	CONFIG
4013	ELAN	413	c340:1/3	-	CONFIG
4014	ELAN	414	-	c441:1	CONFIG
			-	c442:2	
4015	ELAN	415	-	c541:1	DISC_BOTH
4017	ELAN	417	c740:1/3	-	CONFIG
			-	c741:1	
			-	c742:2	
5010	ELAN	500	c50:1/3	-	CONFIG
			-	c51:1	
			-	c52:2	
5011	ELAN	501	u:1/3	-	CONFIG
5013	ELAN	503	c350:1/3	-	CONFIG
5014	ELAN	504	-	c451:1	CONFIG
			-	c452:2	
5015	ELAN	505	-	c551:1	CONFIG
			-	c552:2	
5017	ELAN	507	c750:1/3	-	CONFIG
			-	c751:1	
			-	c752:2	
6000	ELAN	600	c60:1/3	-	CONFIG
6001	ELAN	601	-	c161:1	DISC_BOTH
6002	ELAN	602	-	c262:2	CONFIG
6004	ELAN	604	c460:1/3	-	CONFIG
			-	c461:1	
			-	c462:2	

```
=====
```


SPBM and IS-IS services configuration

7000	ELAN	700	c70:1/3	-		CONFIG
7001	ELAN	701	-	c171:1		DISC_BOTH
7002	ELAN	702	-	c272:2		CONFIG
7004	ELAN	704	c470:1/3	-		CONFIG
			-	c471:1		
			-	c472:2		
8000	ELAN	800	c80:1/3	-		CONFIG
8001	ELAN	801	-	c181:1		DISC_BOTH
8002	ELAN	802	-	c282:2		CONFIG
8004	ELAN	804	c480:1/3	-		CONFIG
			-	c481:1		
			-	c482:2		
9000	ELAN	900	c90:1/3	-		CONFIG
9001	ELAN	901	-	c191:1		CONFIG
9002	ELAN	902	-	292:2		CONFIG
9004	ELAN	904	c490:1/3	-		CONFIG
			-	c491:1		
			-	c492:2		
9005	ELAN	N/A	-	c400:1		DISC_BOTH

c: customer vid u: untagged-traffic

All 42 out of 42 Total Num of Elan i-sids displayed

View all associated MLT on the S-UNI I-SID.

Switch:1# show mlt i-sid

```
=====
MLT Isid Info
=====
```

MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	BPDU
10	6153	100	N/A	20	ELAN	CONFIG	

```
-----
```

1 out of 1 Total Num of i-sid endpoints displayed

View all associated ports on the S-UNI I-SID.

Switch:1# show interface gigabitethernet i-sid

```
=====
PORT Isid Info
=====
```

PORTNUM	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	BPDU
1/1	192	100	N/A	10	ELAN	CONFIG	
1/2	193	100	N/A	10	ELAN	CONFIG	

```
-----
```

2 out of 3 Total Num of i-sid endpoints displayed

View all IS-IS SPBM multicast FIB entries.

Switch:1# show isis spbm multicast-fib detail

```
=====
SPBM MULTICAST FIB ENTRY DETAIL INFO
=====
```

MCAST DA	ISID	BVLAN	SYSID	HOST-NAME	OUTGOING-INTERFACES	INCOMING INTERFACE	CVLAN
----------	------	-------	-------	-----------	---------------------	--------------------	-------

```
-----
```

```

03:77:77:00:0b:b8 3000 1001 0000.beb0.0007 BEB-07 MLT-1 1/2 0
c30:1/3
c31:MLT-1
c32:MLT-2
03:77:77:00:0f:a0 4000 1001 0000.beb0.0007 BEB-07 c40:1/3 1/2 400
c41:MLT-1
c42:MLT-2
03:77:77:00:13:92 5010 1001 0000.beb0.0007 BEB-07 c50:1/3 1/2 500
c51:MLT-1
c52:MLT-2
03:88:88:00:0b:b8 3000 1001 0000.beb0.0008 BEB-08 MLT-1 1/2 0
c30:1/3
c31:MLT-1
c32:MLT-2
03:88:88:00:0f:a0 4000 1001 0000.beb0.0008 BEB-08 c40:1/3 1/2 400
c41:MLT-1

```

Total number of SPBM MULTICAST FIB entries 157

Variable definitions

Use the data in the following table to use the `i-sid` command.

Variable	Value
elan	Displays only all the Elan (S-UNI based) I-SIDs.
MLT ID <1–512>	Specifies the MLT associated with the Switched UNI I-SID.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Specifies the ports associated with the Switched UNI I-SID.

Displaying C-VLAN and Switched UNI I-SID information

Use the following procedure to display C-VLAN I-SID information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the C-VLAN to I-SID associations:

```
show vlan i-sid <1-4094>
```

3. Display I-SID information and Switched UNI to I-SID associations:

```
show i-sid <1-16777215>
```

4. Display the IS-IS SPBM multicast-FIB calculation results by I-SID:

```
show isis spbm i-sid {all|config|discover} [vlan <1-4094>] [id <1-16777215>] [nick-name <x.xx.xx>]
```

5. Display all elan I-SID:

- show i-sid elan

6. Display I-SID configured on MLT:

- show mlt i-sid

7. Display I-SID configured on port:

- show interfaces gigabitethernet i-sid

Example

```
Switch:1#show vlan i-sid
```

```
=====
                                Vlan I-SID
=====
VLAN_ID    I-SID
-----
1
2
5           5
10
20
```

```
Switch: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      Switch1
300      1.11.16        1000    0014.c7e1.33df      config      Switch1
400      1.11.16        1000    0014.c7e1.33df      config      Switch1
200      1.11.16        2000    0014.c7e1.33df      config      Switch1
300      1.11.16        2000    0014.c7e1.33df      config      Switch1
400      1.11.16        2000    0014.c7e1.33df      config      Switch1
200      1.12.45        1000    0016.ca23.73df      discover    Switch2
300      1.12.45        1000    0016.ca23.73df      discover    Switch2

-----
Total number of SPBM ISID entries configed: 6
-----
Total number of SPBM ISID entries discovered: 2
-----
Total number of SPBM ISID entries: 8
-----
```

```
switch:1#show i-sid
```

```
=====
                                Isid Info
=====
ISID      ISID           VLANID    PORT                MLT              ORIGIN
ID        TYPE          INTERFACES INTERFACES          INTERFACES
-----
999      ELAN          99        -                   c110:100        CONFIG
          u: untagged-traffic
          99        1/21              -

All 1 out of 1 Total Num of i-sids displayed
```

```
switch:1#show mlt i-sid
```

```
=====
                                MLT Isid Info
=====
ISID      ISID
```

```

MLTID IFINDEX ID          VLANID C-VID  TYPE      ORIGIN      BPDU
-----
10     6153    100       N/A    20       ELAN        CONFIG
-----

1 out of 1 Total Num of i-sid endpoints displayed

switch:1#show interfaces gigabitEthernet i-sid
=====
                                PORT Isid Info
=====
                                ISID          ISID
PORTNUM IFINDEX ID          VLANID C-VID  TYPE      ORIGIN      BPDU
-----
1/1     192    100       N/A    10       ELAN        CONFIG
1/2     193    100       N/A    10       ELAN        CONFIG
-----

2 out of 3 Total Num of i-sid endpoints displayed

```

Variable definitions

Use the data in the following table to use the **show vlan i-sid** commands.

Variable	Value
<1-4059>	Displays I-SID information for the specified C-VLAN. You can specify the VLAN ID.

Use the data in the following table to use the **show i-sid** commands

Variable	Value
<1-16777215>	Displays I-SID information. You can specify the I-SID ID.

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

Variable	Value
spbm i-sid {all config discover}	<ul style="list-style-type: none"> all: displays all I-SID entries config: displays configured I-SID entries discover: displays discovered I-SID entries

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 i-sid

The following table describes the fields in the output for the `show i-sid` command.

Parameter	Description
I-SID	Indicates the I-SID IDs.
I-SID TYPE	Indicated the I-SID type. <ul style="list-style-type: none"> • T-UNI: Transparent UNI service. • ELAN: any to any service (switched service). • CVLAN: CVLAN based service.
VLANID	Indicates the VLAN IDs.
PORT INTERFACES	Indicated the port interface.
MLT INTERFACES	Indicates the MLT interface.
ORIGIN	Indicates if the I-SID is discovered by Fabric Attach or manually added.

show isis spbm i-sid

The following describes the fields in the output for the `show isis spbm i-sid` command.

Parameter	Description
ISID {all discover config}	Indicates the IS-IS SPBM I-SID identifier. <ul style="list-style-type: none"> • all: display all SPBM I-SID • discover: display discovered SPBM I-SID • config: display configured SPBM I-SID
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.

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.

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.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
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.
- The switch reserves I-SID 0x00ffffff. The switch 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 pane, expand the **Configuration > VLAN** folders.
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.
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 vIST peer.
Status	Indicates the status of this entry: <ul style="list-style-type: none"> • other • invalid • learned • self • mgmt

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 pane, expand the **Configuration > IP** folders.
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	<p>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).</p> <p> Important: You must configure this value lower than the QueryInterval.</p>
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.
LastMembQueryIntvl	Configures the maximum response time (in tenths of a second) that is inserted into group-specific queries

Table continues...


Name	Description
	<p>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 a second.</p> <p>Configure this parameter to values greater than 3. If a fast leave process is not required, use values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)</p>
OtherQuerierPresentTimeout	Specifies the timeout interval.
FlushAction	<p>Configures the flush action to one of the following:</p> <ul style="list-style-type: none"> • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	<p>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.</p> <p> Important:</p> <p>To maximize your network performance, set this parameter according to the version of IGMP currently in use.</p> <ul style="list-style-type: none"> • 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.
Mcast Mode	<p>Specifies the multicast mode:</p> <ul style="list-style-type: none"> • snoop

Table continues...

Name	Description
	<ul style="list-style-type: none"> • pim • snoopSpb • routedSpb • none <p>The default is none.</p>

Configuring IP Multicast over Fabric Connect on a Layer 2 VSN

Use this procedure to enable IP Multicast over Fabric Connect for a Layer 2 VSN. With Layer 2 VSN IP Multicast over Fabric Connect, 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 IP Multicast over Fabric Connect 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 enable IP Multicast over Fabric Connect 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.

This switch only supports IPv4 multicast traffic.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
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**.

Configuring UNI

Use the following procedure to configure a Transparent Port UNI or Switched UNI by mapping an I-SID to a port or MLT and VLAN together.

Consider the following design requirements when implementing Transparent Port UNI (T-UNI):

- Only E-LAN based T-UNI is supported. All T-UNI I-SID end points become members of the same shared E-LAN service. If an E-LINE type of service is required, provision T-UNI at the two end points comprising the point-to-point service.
- A port or MLT associated with a T-UNI I-SID cannot be part of any VLAN.
- Any Spanning Tree Protocol implementation is disabled on the port or MLT associated with the T-UNI I-SID. The port will always be in a Forwarding state.
- MACs are learned against the combination of the I-SID and port or MLT.
- Multiple ports or MLTs can be associated with same T-UNI I-SID.
- One port or MLT cannot be part of multiple T-UNI I-SIDs.
- No additional IS-IS TLVs are added to advertise or withdraw T-UNI I-SID services. Avaya makes use of the existing IS-IS TLV-144 and sub TLV-3 to carry I-SID information.
- The MAC address limit is supported on a per-I-SID basis. For example, the MAC addresses learned on the T-UNI I-SID can be limited.
- IP traffic and control packets are transparently bridged over T-UNI endpoints.
- Untagged traffic ingressing on the T-UNI port will use port QoS. B-TAG and I-TAG priorities are derived from the port QoS.

- The 802.1p bits of the incoming traffic are used to derive the B-TAG and I-TAG priorities for tagged traffic.
- LACP and VLACP PDUs are extracted to the CP and all other control packets are transparently bridged over the T-UNI port or MLT.

T-UNI handles control PDUs in the following manner:

- All the Layer 2 and Layer 3 control packets are transparently bridged over the T-UNI port or MLT with the exception of LACP and VLACP PDUs. LACP PDUs and VLACP PDUs are not transparently bridged over the T-UNI port or MLT if LACP or VLACP is enabled on the port or MLT.
 - If an LACP MLT is associated with a T-UNI I-SID, LACP PDUs are extracted to CP and processed locally.
 - If LACP is not enabled globally and LACP MLT is not associated with the T-UNI I-SID, LACP PDUs are transparently bridged across the T-UNI port or MLT.
 - If a VLACP enabled port is added to a T-UNI I-SID, VLACP PDUs are extracted to the CP for local processing. If a port that is not VLACP enabled is added to the T-UNI I-SID, VLACP PDUs are transparently bridged across T-UNI port.
- The following list of control packet types are transparently bridged across the T-UNI I-SID:
 - SLPP
 - VRRP
 - OSPF
 - RIP
 - BGP
 - ISIS
 - CFM
 - STP
 - SONMP

*** Note:**

If you are configuring a T-UNI to terminate on a port or MLT on a switch in a vIST switch cluster, you must also configure the T-UNI I-SID on the other switch of the vIST switch cluster. You must configure the T-UNI I-SID on both switches of a vIST pair. It is not necessary to assign an actual port or MLT to the T-UNI on the second switch.

About this task

You must first create a type of service instance identifier (I-SID) to create the different types of services available. After you create an I-SID you can add members (ports or MLTs) to the I-SID to create end-points for the service.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.

2. Click **ISID**.
3. Click the **Service** tab.
4. To create a Transparent Port UNI service:
 - a. Click **Insert**.
 - b. Select **elan Transparent** in the **Type** field.
 - c. Enter the I-SID in the **Id** field.
5. To create a Switched UNI service:

*** Note:**

Flex UNI must be enabled to create a Switched UNI service.

- a. Click **Insert**.
 - b. Select **elan** in the **Type** field.
 - c. Enter the I-SID in the **Id** field.
6. Click **Insert**.

UNI field descriptions

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

Name	Description
I-SID ID	Specifies the service interface identifier (I-SID).
I-SID TYPE	Specifies the type of I-SID.
VLANID	Specifies the customer VLAN.
PORT INTERFACES	Specifies the port that is assigned to the I-SID
MLT INTERFACES	Specifies the MLT that is assigned to the I-SID

Associating a port and MLT with an ISID for Elan Transparent

Transparent Port UNI (T-UNI) maps a port or MLT to an I-SID. Transparent Port UNI configures a transparent port where all traffic is MAC switched on an internal virtual port using the assigned I-SID. Multiple ports on the same unit and on other Backbone Edge Bridges (BEBs) are switched on a common I-SID. No VLAN is involved in this process. The T-UNI port is not a member of any VLAN or STG.

Use the following procedure to associate a port and MLT with an ISID.

Before you begin

- You must configure Transparent Port UNI. For more information, see [Configuring Transparent UNI](#) on page 320.
- You must associate a T-UNI LACP MLT with a VLAN before mapping the LACP MLT to a T-UNI ISID.

⚠ Caution:

Ensure that a T-UNI LACP MLT is always associated with a VLAN (even if it is the default VLAN) before adding it to a T-UNI ISID. Otherwise, traffic is not forwarded on the T-UNI LACP MLT.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **ISID**.
3. Click on a row with type configured as elanTransparent.
4. Click **ELAN**.
5. Select port members.
6. Select MLT Ids.
7. Click **Apply**.

Elan Transparent field descriptions

Use the data in the following table to use the Elan Transparent tab.

Name	Description
PortMembers	The set of ports that are members of the elanTransparent service type. From the ports available, you can select single or multiple ports.
MltIds	The set of bits that represent the MLT Ids. From the MLTs available, you can select any, or all of the MLTs to be a part of elan transparent i-sid .

Viewing the ISID forwarding database

View the I-SID forwarding database (FDB).

*** Note:**

To view the T-UNI I-SID FDB entries filtered on a port that is part of an MLT, you must mention the MLT ID in the option for the port.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **ISID**.
3. Click the **FDB** tab.

Click **Filter** to filter rows based on specific filter criteria.

FDB field descriptions

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

Name	Description
IsidId	Specifies the service interface identifier (I-SID).
Address	Specifies the MAC address of the port assigned to the specific I-SID or C-MAC learned on the particular I-SID.
Status	Specifies the learning status of the associated MAC.
Port	Specifies the port on which the MAC is learned for the specific I-SID.
PortType	Specifies whether the MAC is a local or IST-peer or a remote MAC.
RemoteMacDestAddr	Specifies the virtual BMAC address or system-ID of the remote destination.
RemoteMacBVlanId	Specifies the BVLAN ID on which the remote destination was discovered.
RemoteMacDestSysName	Specifies the remote destination system name.

Associating a port and MLT with an I-SID for Elan

Shortest Path Bridging MAC (SPBM) supports Layer 2 Virtual Service Network (VSN) functionality where Switched UNIs are bridged over the SPBM core infrastructure.

Switched User Network Interface (S-UNI) allows the association of local endpoints to I-SIDs based on local port and VLAN together. With switched UNI, the same VLAN can be used on one port to create an endpoint to one I-SID, and on another port to create an endpoint to another I-SID.

Use the following procedure to associate a port and MLT with an I-SID.

About this task

You must configure Switched UNI.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **ISID**.
3. Click the row with type configured as elan.
4. Click **Switched Uni**.
5. Click **Insert**.
6. Enter the VLAN ID in the **Cvid** field.
7. Click **Port** or **Mlt** to update the interface index in the **Ifindex** field.
8. Click **Insert**.

Switched Uni field descriptions

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

Name	Description
Cvid	Specifies the customer VLAN identifier.
Ifindex	Specifies the interface index of the Elan end point.

Viewing the I-SID interface

View the I-SID interface.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **ISID**.
3. Click the **Interface** tab.

Click **Filter** to filter rows on specific filter criteria.

Interface field descriptions

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

Name	Description
Ifindex	Specifies the interface index.
Isid	Specifies the service interface identifier (I-SID).
Vlan	Specifies the platform VLAN.
Cvid	Specifies the customer VID.
Type	Specifies the type of service associated with the I-SID interface.
Origin	Specifies the origin of the service associated with the I-SID interface.
Bpdu	Specifies the BPDU forward option for the untagged traffic port.

Layer 2 VSN configuration examples

This section provides configuration examples to configure Layer 2 VSNs.

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Layer 2 VSN configuration example

The following figure shows a sample Layer 2 VSN deployment.




Figure 45: 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 223.

SwitchC

VLAN CONFIGURATION

```
vlan create 10 type port-mstprstp 1
vlan members 10 1/1 portmember
vlan i-sid 10 12990010
```

SwitchD

VLAN CONFIGURATION

```
vlan create 10 type port-mstprstp 1
vlan members 10 1/1 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.

SwitchC

```
SwitchC: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      SwitchD
12990010 f.30.13       4000   0015.e89f.e3df   config       SwitchC
=====
Total number of SPBM ISID entries configured: 1
Total number of SPBM ISID entries discovered: 1
Total number of SPBM ISID entries: 2
=====
```

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

```
=====
                        SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA           ISID   BVLAN   SYSID           HOST-NAME   OUTGOING-INTERFACES
```

```
-----
f3:30:14:c6:36:3a 12990010 4000 0014.0da0.13df SwitchD 1/1
f3:30:13:c6:36:3a 12990010 4000 0015.e89f.e3df SwitchC 1/30,1/1
-----
```

SwitchD

```
SwitchD: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        SwitchD
12990010 f.30.13       4000   0015.e89f.e3df  discover      SwitchC
=====
```

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

```
=====
                        SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA           ISID   BVLAN   SYSID           HOST-NAME   OUTGOING-INTERFACES
-----
f3:30:14:c6:36:3a 12990010 4000   0014.0da0.13df  SwitchD     MLT-1,1/1
f3:30:13:c6:36:3a 12990010 4000   0015.e89f.e3df  SwitchC     1/1
=====
```

SwitchC — verifying with CFM

```
SwitchC: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  SwitchC          00:15:e8:9f:e3:df -> SwitchG          00:0e:62:25:a3:df
2  SwitchG          00:0e:62:25:a3:df -> SwitchD          00:14:0d:a0:13:df
```

SwitchD — verifying with CFM

```
SwitchD: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  SwitchD          00:14:0d:a0:13:df -> SwitchG          00:0e:62:25:a3:df
2  SwitchG          00:0e:62:25:a3:df -> SwitchC          00:15:e8:9f:e3:df
```

SwitchC — verifying FDB

```
SwitchC:1# show vlan mac-address-entry 10
```

```
=====
                        Vlan Fdb
=====
VLAN   STATUS   MAC ADDRESS           INTERFACE           TUNNEL
-----
10     learned  00:00:00:00:00:01    Port-1/1           SwitchD
10     learned  00:00:00:00:00:02    Port-1/1           SwitchD
=====
```

```
2 out of 4 entries in all fdb(s) displayed.
```

```
SwitchC:1# show vlan remote-mac-table 10
```

```
=====
                        Vlan Remote Mac Table
=====
VLAN STATUS   MAC-ADDRESS           DEST-MAC           BVLAN   DEST-SYSNAME PORTS
-----
10   learned  00:00:00:00:00:02    00:14:0d:a0:13:df  0014.0da0.13df  SwitchD 1/30
=====
```

```
Total number of VLAN Remote MAC entries 1
```

SwitchD — verifying FDB

```
SwitchD:1# show vlan mac-address-entry 10
```

```
=====
```

Vlan Fdb				
VLAN ID	STATUS	MAC ADDRESS	INTERFACE	TUNNEL
10	learned	00:00:00:00:00:01	Port-1/1	SwitchC
10	learned	00:00:00:00:00:02	Port-1/1	SwitchC

```
=====
```

2 out of 4 entries in all fdb(s) displayed.

```
SwitchD:1# show vlan remote-mac-table 10
```

```
=====
```

Vlan Remote Mac Table						
VLAN	STATUS	MAC-ADDRESS	DEST-MAC	DEST-SYSID	DEST-SYSNAME	PORTS
10	learned	00:00:00:00:00:01	00:15:e8:9f:e3:df	0015.e89f.e3df	SwitchC	MLT-1

```
=====
```

```
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 223.




Figure 46: Layer 2 VSN with different VLAN IDs

The following sections show the steps required to configure the Layer 2 VSN parameters in this example.

SwitchC

```
VLAN CONFIGURATION
```

```
vlan create 9 type port 1
vlan members 9 1/1 portmember
vlan i-sid 9 9
```

SwitchD

```
VLAN CONFIGURATION
vlan create 19 type port 1
vlan members 19 1/1 portmember
vlan i-sid 19 9
```

Layer 2 VSN with IP Multicast over Fabric Connect configuration example

The example below shows the configuration steps to enable IP Multicast over Fabric Connect 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.

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.

Fabric Connect supports both IPv4 Shortcuts and IPv6 Shortcuts. Because IPv6 Shortcuts depend on IPv4 Shortcuts, you should understand how IPv4 Shortcuts work (see [SPBM IP shortcuts](#) on page 330) before jumping to the IPv6 section.

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 47: 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.

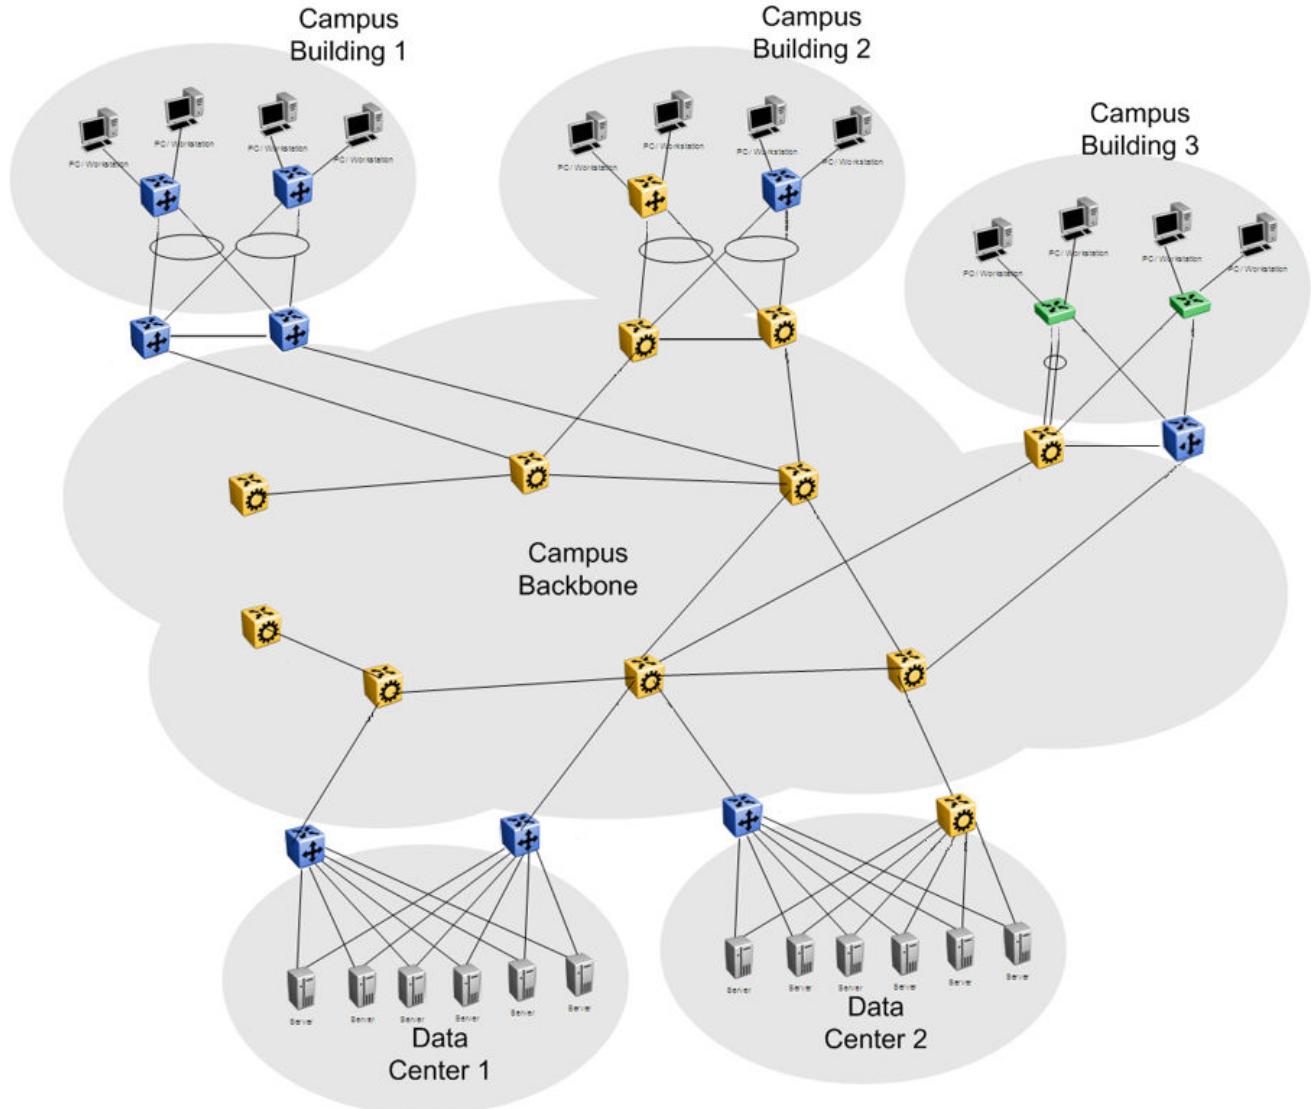


Figure 48: SPBM IP shortcuts in a campus

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. To advertise IPv6 routes from the BEBs into the SPBM network, you must enable route redistribution of IPv6 direct, IPv6 static, and OSPFv3 routes into IS-IS.

SPBM IPv6 Shortcuts

Both IPv4 and IPv6 Shortcuts use IS-IS as the Interior Gateway Protocol (IGP) and the link state packet (LSP) for reachability information. However, IPv4 Shortcuts use TLV 135 and IPv6 Shortcuts use TLV 236. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as

leaf nodes. IS-IS transports the IPv6 reachability information to remote BEBs and uses the shortest path, calculated by SPBM, for data forwarding.

*** Note:**

You only configure the IPv6 address information on the edges. There is no IPv6 in the SPBM cloud.

IS-IS transports the IPv6 routes through TLV 236 in the LSP advertisements. These routes are installed in the Global Routing Table (GRT) with the node from which the LSPs carrying the IPv6 routes are received as the next hop.

IPv6 Shortcuts dependency on IPv4 Shortcuts

IPv6 Shortcuts function in a very similar manner to IPv4 Shortcuts and depends on IPv4 Shortcuts for some functions. For example, IPv6 Shortcuts use the BMAC (local and remote) information created by IPv4 Shortcuts.

! Important:

IPv4 Shortcuts must be enabled before you enable IPv6 Shortcuts.

An error is displayed if you try to enable IPv6 Shortcuts but do not have IPv4 Shortcuts already enabled.

IPv6 Shortcuts alone can be disabled while leaving IPv4 Shortcuts enabled. When IPv4 Shortcuts is disabled without disabling IPv6 Shortcuts disabled first, a warning or error message is displayed indicating that IPv6 should be disabled first.

Circuitless IPv6 (CLIPv6)

To enable IPv6 Shortcuts on the BEBs and to advertise the local BEB to other IS-IS nodes, you must configure a circuitless IPv6 address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 236.

IPv6 Shortcuts support Circuitless IPv6 (CLIPv6), which ensures uninterrupted connectivity to the switch as long as there is an actual path to reach it. This route always exists and the circuit is always up because there is no physical attachment.

Migrating the GRT to IPv6 Shortcuts

Use the following steps to migrate the Global Router Table (GRT) to use IPv6 Shortcuts over the SPBM core:

- Identify the nodes that should be enabled with IPv6 Shortcuts. Apply these steps to all of these nodes.
- Activate and validate basic IPv6 Shortcuts. For information, see [SPBM IPv6 Shortcuts](#) on page 332.
- Configure IS-IS route preference to ensure that the IPv6 IGP protocol currently being used in the SPBM core is preferred over the IS-IS routes.
- Enable redistribution of direct and static IPv6 routes into IS-IS.
- Create route policies to permit only IPv6 IGP routes from the access side of the SPBM network.

- Configure redistribution of routes from the IPv6 route table from each of the IPv6 IGP protocols into IS-IS along with the appropriate route policy.
- Use the `show isis spbm ipv6-unicast-fib` command to check the IS-IS LSDB, IS-IS routes, and to verify that all the desired IPv6 routes are now in IS-IS.
- Configure redistribution of IS-IS routes from the IPv6 route table into each of the IPv6 IGP protocols in use. This redistribution does not require a route policy since IS-IS is only supported in the SPBM core.
- Change IS-IS route-preference to ensure that IS-IS routes are preferred over other IPv6 IGP routes.
- Disable/delete old IPv6 IGP in the SPBM core.

! Important:

Use only one IPv6 routing protocol in the SPBM core to prevent the possibility of routing loops.

IPv6 Shortcut limitations and considerations

The following features are not supported:

- Disabling and enabling alternate routes for IPv6 routes
- Redistribution of RIP into IS-IS
- BGP+ and its redistribution into IS-IS
- 6-in-4 tunnels are not supported when the tunnel destination IP is reachable via IPv4 Shortcuts route.
- IS-IS accept policies are applicable to IPv4 Shortcut routes only, not IPv6 Shortcut routes.
- Route policies are not supported with redistribution of IPv6 routes.

Keep the following considerations in mind when configuring IPv6 Shortcuts:

- IPv4 Shortcuts must be enabled before enabling IPv6 Shortcuts.
- IPv6 Shortcuts support Circuitless IPv6 (CLIPv6) with the following limitations:
 - Stateless address autoconfiguration (SLAAC) is not supported on IPv6 CLIP interfaces.
 - IPv6 CLIP does not support link-local address configuration.
 - To configure an IPv6 address with a prefix length from 65 to 127 on a CLIP interface, you must enable the IPv6 mode flag.

*** Note:**

This limitation does not apply to VSP 4000 switches.

- Neighbor discovery (ND) does not run on an IPv6 CLIP interface. Therefore, the system does not detect when you configure a duplicate IPv6 address.
- Multiple IPv6 address configuration on an IPv6 CLIP interface is not supported.
- You can configure a maximum of 64 IPv6 CLIP interfaces.
- IPv6 CLIP interface is enabled by default and it cannot be disabled.

- IPv6 with vIST provides the same support as IPv4 with vIST.
- To help with debugging, CFM provides full support for both IPv4 and IPv6 addresses for the `l2ping` and `l2tracert` commands.

ECMP with IS-IS

The ECMP feature supports and complements the IS-IS protocol.

With Equal Cost Multipath (ECMP), the Virtual Services Platform can determine up to four equal-cost paths on a VSP 4000 series device and up to eight equal-cost paths on a VSP 8000 or VSP 7000 series device to the same destination prefix.

You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP and IPv6 traffic. Equal Cost Multipath is formed using routes from the same protocol.

ECMP within IS-IS routes

Equal Cost Multipath (ECMP) allows the device to determine up to four equal cost paths for VSP 4000 series devices or eight equal cost paths for VSP 7000/8000 series devices 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 or IPv6 prefix that the device receives from a

BEB results in multiple ECMP paths. These paths may or may not be physically diverse. SPBM supports up to two B-VLANs; a primary B-VLAN and a secondary B-VLAN.

By default, when ECMP is enabled:

- up to four equal paths can exist to a destination for VSP 4000 series devices, but you can change this number to a value from 1 to 4
- up to eight equal paths can exist to a destination for VSP 7000 and 8000 series devices, 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, the device adds the route from the lowest system ID with the primary B-VLAN. In this example, the device adds BEB-A B-VLAN 1000.

*** Note:**

- ECMP is supported for IPv6 Shortcut routes.
- To add IS-IS equal cost paths in the routing table, you must enable IPv6 ECMP feature globally.

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 VSP switches:

- One router redistributes IP routes from Routing Information Protocol (RIP)/Open Shortest Path First (OSPF) into IS-IS (IP).
- The second router redistributes from IS-IS (IP) into RIP or OSPF.

The following figure illustrates this configuration.




Figure 49: Redundant OSPF or RIP network

In this scenario it is necessary to take extra care when redistributing through both VSP 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 (VSP 8200G and VSP 8200D) have an OSPF or a RIP route to 192.168.10.0/24 with the next-hop to VSP 8200A.

As soon as the VSP 8200G node redistributes that IP route into IS-IS, the VSP 8200D node learns the same route through IS-IS from VSP 8200G. (The VSP8200G node already has the route through OSPF or RIP). Because IS-IS has a higher preference, VSP 8200D replaces its 192.168.10.0 OSPF route with an IS-IS one that points at VSP 8200G as the next-hop. The following figure illustrates this scenario.

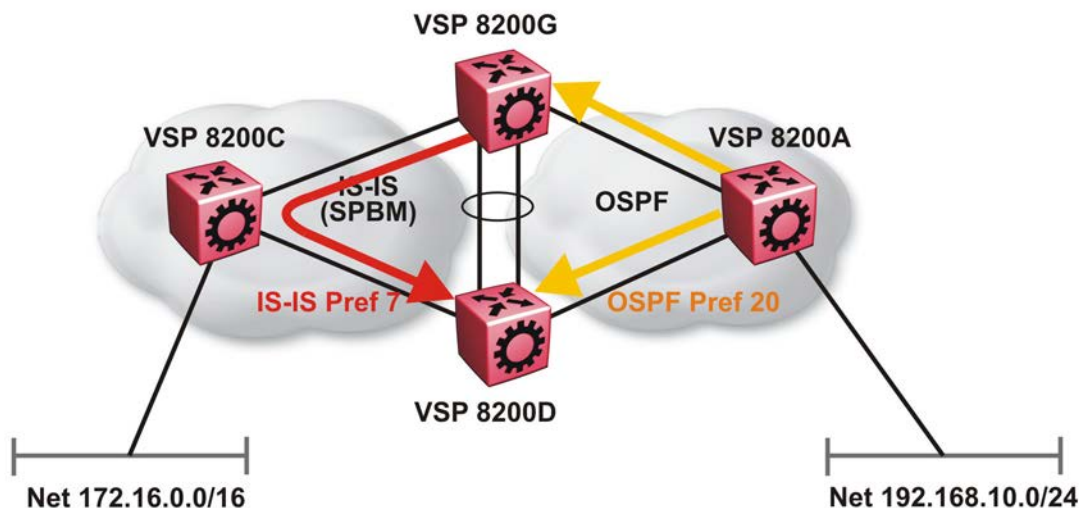


Figure 50: Redistributing routes into IS-IS

Clearly this is undesirable and care needs to be taken to ensure that the two redistributing nodes (VSP 8200G and VSP 8200D) do not accept redistributed routes from each other. With IS-IS accept policies, you can associate an IS-IS accept policy on VSP 8200D to reject all redistributed IP routes received from VSP 8200G, and VSP 8200G to reject all redistribute IP routes from VSP 8200D.

*** Note:**

IS-IS accept policies do not apply to IPv6 shortcut routes and only apply to IPv4 shortcut routes.

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
```

*** Note:**

For IPv6, the command is `ipv6 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:

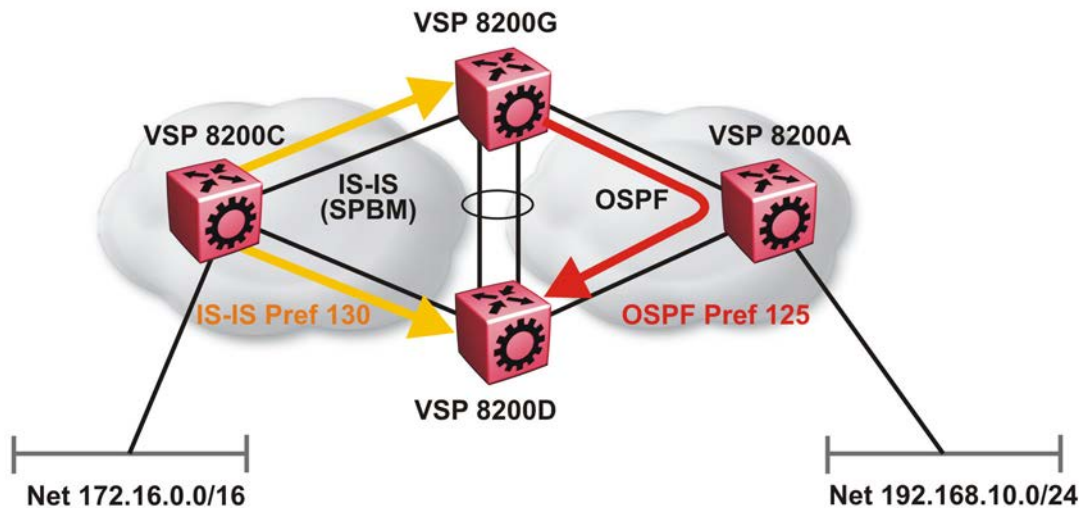


Figure 51: Redistributing routes into OSPF

In the preceding figure, both VSP 8200G and VSP 8200D have an IS-IS IP route for 172.16.0.0/16 with the next hop as VSP 8200C. As soon as the VSP 8200G redistributes the IS-IS route into OSPF, the VSP 8200D node learns that same route through OSPF from VSP 8200G. (The VSP 8200G node already has the route through IS-IS).

Because OSPF has a higher preference, VSP 8200D 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 8200D from accepting any AS External (LSA5) routes from VSP 8200G and prevent VSP 8200G from accepting any AS External (LSA5) routes from VSP 8200D. 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 that 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 8200A advertises 25000 OSPF routes to VSP 8200G and VSP 8200D, then both VSP 8200G and VSP 8200D install the 25000 routes as OSPF routes. Since VSP 8200D and VSP 8200G 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 8200D and VSP 8200G keep IS-IS learned routes as alternative routes.

If VSP 8200A withdraws its 25000 OSPF routes, VSP 8200G and VSP 8200D remove the OSPF routes. While the OSPF routes are removed the routing tables of VSP 8200G and VSP 8200D activate the alternative IS-IS routes for the same prefix. Since VSP 8200G and VSP 8200D have IS-IS to OSPF redistribution enabled, VSP 8200A learns these routes as OSPF and this causes a routing loop. Use the **no ip alternative-route** command to disable alternative routes on VSP 8200G and VSP 8200D to avoid routing loops.

In the preceding figure, you leveraged OSPF Accept policies, which can be configured to prevent VSP 8200D from accepting any AS External (LSA5) routes from VSP 8200G and prevent VSP 8200G from accepting any AS External (LSA5) routes from VSP 8200D. 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 52: Redistributing routes into RIP

VSP 8200G

```

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 1/11
ip rip default-supply enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

```


SPBM and IS-IS services configuration

```
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 8200A

RIP PORT CONFIGURATION

```
interface gigabitethernet 1/2
 ip rip default-listen enable
 exit
 interface gigabitethernet 1/3
 ip rip default-listen enable
 exit
```

VSP 8200D

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 2/11
 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 8200G and VSP 8200D supply the default route on their relevant interfaces, and not accept it on the same interfaces. Likewise, VSP 8200A will accept the default route on its interfaces to both VSP8200G and VSP8200D but it will not supply the default route back to them. This will prevent the default route advertised by VSP8200G from being installed by VSP8200D, and vice-versa.

The preceding example where IS-IS IP routes are aggregated into a single default route when redistributed into the RIP network also applies 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 policies and DvR

When you configure DvR in an SPB network, you can leverage IS-IS accept policies to control the DvR routes learned from the DvR backbone. The DvR backbone contains the master list of all the host routes learned from various DvR domains.

You can configure accept policies on a DvR Controller or a non-DvR BEB as a filter to determine which DvR host routes to accept into the routing table, from the DvR backbone. Accept policies apply to only those backbone (or inter-domain) host routes that are not part of the Controller's own DvR enabled subnets *and* do not have the same domain ID as that of the Controller.

For non-DvR BEBs, all the routes present in the backbone are learned, but you can still use the accept policies to filter specific routes.

For information on DvR, see *Configuring IPv4 Routing*.

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
- backbone-route-map
- a combination of route-map and backbone-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.

On DvR Controllers in a DvR domain, you can configure a backbone route policy to determine what host routes to accept from the DvR backbone, into the routing table. Also, just like on the route policy, you can configure match criteria, and set preferences on the backbone route policy.

To accept both IS-IS routes and host routes from the DvR backbone, you can configure both a route policy and a backbone route policy in the accept policy instance.

For more information on configuring route policies, see *Configuring IPv4 Routing*.

The following table describes IS-IS accept policy filters.

Filters into	Filter	Description
Global Routing Table (GRT)	accept route-map <i>WORD<1-64></i>	By default, the device accepts all routes into the GRT and VRF routing table. This is the default accept policy.

Table continues...

Filters into	Filter	Description
	accept route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	This is the default accept policy with configuration to accept specific DvR host routes from the DvR backbone.
	accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the list of I-SIDs. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
Virtual Routing and Forwarding (VRF) routing table	isis accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept i-sid <0-16777215> route-map <i>WORD</i> <1-64>	The device filters based on the I-SID, which represents a local or remote Layer

Table continues...

Filters into	Filter	Description
	backbone-route-map <i>WORD</i> <1-64>	3 VSN. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept adv-rtr <x.xx.xx> i-sid <0-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept adv-rtr <x.xx.xx> isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept route-map <i>WORD</i> <1-64> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the route policy. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.

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.
- 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-rtt` 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 the configuration of route-map, see *Configuring IPv4 Routing*.

IP Multicast over Fabric Connect within the GRT

IP Multicast over Fabric Connect within the GRT enables you to exchange IP multicast traffic with all or a subset of VLANs that are in the Global Routing Table (GRT). This restriction is called the *scope level*, which IP Multicast over Fabric Connect uses to constrain the multicast streams within the level in which they originate. For example, if a sender transmits a multicast stream to a BEB on a VLAN that is part of the GRT with IP Multicast over Fabric Connect enabled, only receivers that are part of the same GRT can receive that stream.

Applications that can use IP Multicast over Fabric Connect within the GRT include: Video surveillance, TV/Video/Ticker/Image distribution, VX-LAN.

Both **IP Shortcuts** and **IP Multicast over Fabric Connect within the GRT** use the GRT for the scope level to constrain multicast streams. However, they are separate features that work independently from each other.

Important:

You do not have to enable IP Shortcuts to support IP Multicast over Fabric Connect within the GRT.

With IP Multicast over Fabric Connect within the GRT, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP Multicast over Fabric Connect enabled. When you enable IP Multicast over Fabric Connect 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 Fabric Connect on all VLANs to which IP multicast senders and receivers attach. IP Multicast over Fabric Connect is typically configured only on BEBs.

*** Note:**

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).

I-SIDs

Unlike IP Shortcuts with unicast, a data I-SID (for mac-in-mac encapsulation of the multicast traffic) is required for IP Multicast over Fabric Connect within the GRT. 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, IP Multicast over Fabric Connect within the GRT does not have a scope I-SID to determine the scope of the multicast traffic. Instead the scope is the Global Routing Table.

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 Multicast over Fabric Connect within the GRT, 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`, 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 IP Multicast over Fabric Connect is enabled, an error message appears.

After you configure `ip spb-multicast enable` on each of the VLANs within the GRT that need to support IP multicast traffic, any IGMP functions required for IP Multicast over Fabric Connect within the GRT are automatically enabled. You do not need to configure anything IGMP related.

DvR

When you enable `ip spb-multicast` on the Controller nodes, the configuration is automatically pushed to all the Leaf nodes within the domain.

For more information on DvR, see *Configuring IPv4 Routing*.

IP Shortcuts configuration using the CLI

This section provides procedures to configure IP Shortcuts using the CLI.

Configuring SPBM IPv4 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 (CLIP) 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 and static routes into IS-IS.

*** Note:**

The loopback address on each switch or BEB must all be in different subnets to ensure connectivity between them. To do this, use a 32-bit mask with the CLIP address.

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 the Loopback Interface Configuration mode to Global Configuration mode:

```
exit
```

4. Log on to IS-IS Router Configuration mode:

```
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 {bgp | direct | ospf | rip | static | vrf  
WORD<1-16>}
```

Example

```
Switch:1> enable
```

```
Switch:1# configure terminal
```

```
Switch:1(config)# interface loopback 1
```

```
Switch:1(config-if)# ip address 10.0.0.2/8
```

```
Switch:1(config-if)# exit
```

```
Switch:1(config)# router isis
```

```
Switch:1(config-isis)# ip-source-address 10.0.0.2
```

```
Switch:1(config-isis)# spbm 1 ip enable
```

```
Switch:1(config-isis)# show isis spbm
```

```
show isis spbm
```

```
=====
                        ISIS SPBM Info
=====
```

SPBM INSTANCE	B-VID	PRIMARY VLAN	NICK NAME	LSDB TRAP	IP	IPV6	MULTICAST
1	4086-4087	4086	3.03.01	disable	enable	enable	disable

```
=====
                        ISIS SPBM SMLT Info
=====
```

SPBM INSTANCE	SMLT-SPLIT-BEB	SMLT-VIRTUAL-BMAC	SMLT-PEER-SYSTEM-ID
1	primary	00:00:03:03:03:03	0000.0303.0302

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

```
Switch:1(config-isis)# redistribute rip
```

```
Switch:1(config-isis)# redistribute rip enable
Switch:1(config-isis)# exit
Switch: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 <i>WORD</i> <0–64>	Configures the route policy to apply to redistributed routes. Type a name between 0 to 64 characters in length.

Table continues...

Variable	Value
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 use the `isis apply redistribute` command.

Variable	Value
{bgp direct ospf rip static}	Specifies the protocol.

Configuring SPBM IPv6 Shortcuts

! Important:

You must enable IPv4 Shortcuts before you enable IPv6 Shortcuts because IPv6 Shortcuts depend on IPv4 Shortcuts for some functions.

Configuring IPv6 Shortcuts is essentially the same as the IPv4 procedure except you use the following IPv6 commands instead of their IPv4 equivalents:

- Use `ipv6 interface address` to create a CLIPv6 interface with an IPv6 address.
- Use `ipv6 ipv6-source-address` to specify the CLIPv6 interface as the source address for IPv6 Shortcuts.
- Use `spbm ipv6 enable` to enable IPv6 Shortcuts.
- Use `ipv6 redistribute {direct | static | ospf } enable` to control the redistribution of GRT IPv6 routes into the SPBM IS-IS domain.
- Use `ipv6 route preference protocol spbm-level1` to change route preference values for IPv6 Shortcut routes learned through IS-IS.

To enable IPv6 Shortcuts on the BEBs, you must configure a circuitless IPv6 (CLIPv6) address (loopback address), and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 236. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct and static routes into IS-IS.

* Note:

The loopback address on each switch or BEB must all be in different subnets to ensure connectivity between them. To do this, use a 32-bit mask with the CLIP address, and the CLIPv6 address prefix must be 128.

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 IPv6 addresses and network masks.

Procedure

1. Enter Loopback Interface Configuration mode

```
enable
configure terminal
interface Loopback <1-256>
```
2. Configure a CLIPv6 interface to use as the source address for SPBM IPv6 Shortcuts:

```
ipv6 interface address WORD<0-255>
```
3. Exit the Loopback Interface Configuration mode to Global Configuration mode:

```
exit
```
4. Log on to IS-IS Router Configuration mode:

```
router isis
```
5. Specify the CLIPv6 interface as the source address for SPBM IPv6 Shortcuts:

```
ipv6-source-address WORD<0-46>
```
6. Enable SPBM IPv6 Shortcuts:

```
spbm <1-100> ipv6 enable
```
7. Display the status of SPBM IPv6 Shortcuts on the switch:

```
show isis spbm
```
8. Identify IPv6 routes on the local switch to be announced into the SPBM network.
 - a. To redistribute routes for directly connected subnets:

```
ipv6 redistribute direct enable
```
 - b. To redistribute static routes into IS-IS:

```
ipv6 redistribute static enable
```
 - c. To redistribute OSPFv3 routes into IS-IS, enter IS-IS Router Configuration mode and then enable IS-IS redistribution:

```
router isis
ipv6 redistribute ospf enable
```
9. If you want to delete the configuration, use the no option:

```
no redistribute {direct | static | ospf} enable
```
10. Exit to Global Configuration mode:

```
exit
```
11. **(Optional)** Change route preference values for IPv6 Shortcut routes learned through IS-IS:

```
ipv6 route preference protocol spbm-level1 <0-255>
```

12. Apply the configured redistribution:

```
isis apply redistribute {bgp | direct | ospf | rip | static | vrf}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 123
Switch:1(config-if)#ipv6 interface address 123::1/128
Switch:1(config-if)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#ipv6 ipv6-source-address <non-link-local ipv6-address>
Switch:1(config-isis)#spbm 1 ipv6 enable
Switch:1(config-isis)#show isis spbm
=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY    NICK      LSDB      IP        IPV6      MULTICAST
INSTANCE  VLAN
-----
1          4086-4087  4086       3.03.01   disable   enable    enable    disable
=====
                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID
INSTANCE
-----
1          primary              00:00:03:03:03:03      0000.0303.0302
-----
Total Num of SPBM instances: 1
=====
```

Variable definitions

Use the data in the following table to use the IPv6 Shortcuts commands.

Variable	Value
ipv6-source-address <i>WORD</i> <0-46>	Specifies the source IPv6 address for locally generated IPv6 packets whose egress port is an SPBM NNI port. The <i>WORD</i> <0-46> value must be a locally configured loopback IPv6 address (CLIPv6). Use the no option to delete the specified IPv6 address.
spbm<1-100> ipv6 enable	Enables or disables SPBM IPv6 Shortcuts. The default is disabled. Use the no or default options to disable SPBM IPv6 Shortcuts.
ipv6 route preference protocol spbm-level1 <0-255>	Sets the route preference value for IPv6 Shortcut routes learned through IS-IS. The default preference is 7.
ipv6 redistribute { <i>direct</i> <i>static</i> <i>ospf</i> enable	Specifies the GRT IPv6 route that you want to redistribute into the SPBM IS-IS domain.

Table continues...

Variable	Value
	The default is disabled. Use the no option to disable the redistribution.

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.

If DvR is enabled on your switch, and the switch is either a DvR Controller or a non-DvR BEB within the domain, you can configure IS-IS accept policies to accept specific host routes from the DvR backbone. For information on DvR, see *Configuring IPv4 Routing*.

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. You should 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.
- The number of unique Layer 3 VSN I-SIDs used on a BEB is limited to the number of VRFs supported on the switch. This includes the I-SID values used for Layer 3 VSNs and the I-SID values specified for the ISIS accept policy filters, which can be configured using the `ip isid-list [ISID#], accept i-sid <value>`, or `accept adv-rtr <isis nn> i-sid <value>` commands.

The switch supports 24 VRFs by default, so, in a default configuration, you cannot create an `ip isid-list` or `accept policy` with more than 24 unique I-SID entries. However, the configured VRFs take up an entry, so the formula to calculate the limit is: $[24 \text{ VRF Limit} - (\text{currently configured VRFs})]$. This gives the number of unique I-SIDs that can be used directly in the IS-IS accept policy filters, which you implement with the `ip isid-list` or `accept policy` command. The I-SIDs used for Layer 3 VSNs can be reused in IS-IS accept policy filters without affecting the limit.

If you increase the VRF scaling, you can create more Layer 3 VSNs. For more information about how to increase the number of supported VRFs, see *Configuring IPv4 Routing*. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Release Notes*.

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 IPv4 Routing*.
- Ensure that DvR is enabled on the switch before you configure an IS-IS accept policy with a backbone route policy, to accept host routes from the DvR backbone.

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
```

Configure IS-IS accept policies with a route policy or a backbone route policy or a combination of both, to determine which routes the IS-IS accept policy applies to.

Configure one of the following types of IS-IS accept policies.

- **An IS-IS accept policy with only the route policy:**

The IS-IS routes are selectively accepted based on the route policy. Since the backbone route policy is not configured, all host routes from the DvR backbone are *denied*.

If you do not configure a route policy, by default, all IS-IS routes are *accepted*.

- **An IS-IS accept policy with only the backbone route policy:**

The DvR host routes from the DvR backbone are selectively accepted based on the backbone route policy. Since the route policy is not configured, all IS-IS host routes are accepted.

If you do not configure a backbone route policy, all host routes from the DvR backbone are *denied*.

- **An IS-IS accept policy with both route policy and backbone route policy:**

IS-IS routes are selectively accepted based on the route policy and host routes from the DvR backbone are selectively accepted based on the backbone route policy.

5. Configure an IS-IS accept policy instance with a route policy.

Use one of the following options:

- a. 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>]
```

- b. 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>]
```

- c. **(Optional)** Delete an IS-IS accept policy instance:

```
no accept [adv-rtr <x.xx.xx>][i-sid <1-16777215>][isid-list WORD <1-32>]
```

- d. Specify an IS-IS route policy to apply to routes from all BEBs:

```
accept route-map WORD<1-64>
```

- e. Specify an IS-IS route policy to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx>[route-map WORD<1-64>]
```

- f. **(Optional)** Delete an IS-IS route policy:

```
no accept [adv-rtr <x.xx.xx>] [route-map]
```

- g. Enable an IS-IS route accept instance:

```
accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list WORD<1-32>]
```

- h. **(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>]
```

6. Configure an IS-IS accept policy instance with a backbone route policy to accept host routes from the DvR backbone:

*** Note:**

IS-IS accept policies typically apply to all IS-IS routes. However, to accept DvR host routes from the DvR backbone, you *must* explicitly configure the IS-IS accept policy with a backbone route policy.

Use one of the following options:

- a. Create the default IS-IS accept policy instance to accept host routes from the DvR backbone:

```
accept backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy instance with backbone route policy configuration:

```
no accept backbone-route-map
```

- c. Create an IS-IS accept policy instance to accept host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-
route-map WORD<1-64>
```

- d. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-
route-map
```

- e. Create an IS-IS accept policy instance to accept host routes from the DvR backbone and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> backbone-route-map WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to a specific advertising BEB

```
no accept adv-rtr <x.xx.xx> backbone-route-map
```

7. Configure an IS-IS accept policy with both route policy and backbone route policy, to selectively accept IS-IS routes as well as host routes from the DvR backbone.

- a. Create the default IS-IS accept policy instance with a route policy to accept IS-IS routes and a backbone route policy to accept host routes from the DvR backbone:

```
accept route-map WORD<1-32> backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy with route policy and backbone route policy configuration:

```
no accept route-map backbone-route-map
```

- c. Create an accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map
WORD<1-32> backbone-route-map WORD<1-64>
```

- d. **(Optional)** Delete an accept policy instance with route policy and backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map
backbone-route-map
```

- e. Create an IS-IS accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> route-map WORD<1-32> backbone-route-map
WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with route policy and backbone route policy configuration, which applies to a specific advertising BEB:

```
no accept adv-rtr <x.xx.xx> route-map backbone-route-map
```

- Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
isis apply accept [vrf WORD <1-16>]
```

- Exit IS-IS Router Configuration mode:

```
exit
```

You are in Global Configuration mode.

Example

Configure an I-SID based IS-IS accept policy with the route policy test:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map test 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit

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
```

Configuration of IS-IS accept policy to accept host routes from the DvR backbone

Example 1:

To accept host routes from the DvR backbone, you must configure a backbone route policy and apply it to the IS-IS accept policy.

- Configure a route policy for DvR:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map dvrmap1 1
Switch:1(route-map)#enable
```

- Configure an IS-IS accept policy for I-SID 10, and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap1
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
```

OR

Configure the default accept policy for IS-IS and DvR, and apply the route policy as a backbone route policy:

```
Switch:1(config)#route-map isismap1 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept route-map isismap1 backbone-route-map dvrmap1
```

3. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
Switch:1(config)#exit
```

4. Verify the configuration:

```
Switch:1#show ip isis accept

=====
                        Isis Accept - GlobalRouter
=====
ADV_RTR  I-SID    ISID-LIST                                ENABLE POLICY          BACKBONE
-----  -
-         10      -                                TRUE                    dvrmap1
-         -       -                                isismap1                dvrmap1

2 out of 2 Total Num of Isis Accept Policies displayed
```

Example 2:

Configure an IS-IS accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24
```

2. Create the route policy dvrmap2 to match the IP prefix list:

```
Switch:1(config)#route-map dvrmap2 1
Switch:1(route-map)#match network listPrefix
Switch:1(route-map)#enable
```

3. Create an IS-IS accept policy with I-SID 10 and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
```

4. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

The above command causes IS-IS to accept all routes with I-SID 10. To deny IS-IS routes and accept only DvR host routes, you can configure an additional IS-IS route policy as follows:

```
Switch:1(config)#route-map isismap2 1
Switch:1(route-map)#no permit
Switch:1(route-map)#enable

Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 route-map isismap2 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

5. Verify the configuration:

```
Switch:1(config)#exit
Switch:1#show ip isis accept

=====
Isis Accept - GlobalRouter
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
POLICY
-----
-         10       -                        TRUE  isismap2                dvrmap2
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configuration of IS-IS accept policies for a specific VRF instance

Example 1:

Configure IS-IS accept policies to accept host routes from the DvR backbone, for a specific VRF instance.

1. In the VRF green context, configure the route policy `dvrmap3` for DvR:

```
Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap3 1
Switch:1(router-vrf-routemap)#enable
```

2. Use one of the following options to configure an IS-IS accept policy, and apply the route policy as a backbone route policy:

Configure an IS-IS accept policy for a specific advertising BEB with nickname `1.11.11`:

```
Switch:1(router-vrf-routemap)#isis accept adv-rtr 1.11.11 backbone-route-map
dvrmap3
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#isis accept adv-rtr 1.11.11 enable
```

```
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
POLICY
-----
1.11.11  -        -                        TRUE                     dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

```
Switch:1(config)#show ip isis accept vrfids 2

=====
Isis Accept - VRF green
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
POLICY
-----
1.11.11  -        -                        TRUE                     dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for I-SID 10:

```
Switch:1(router-vrf)#isis accept i-sid 10 backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID      ISID-LIST                ENABLE POLICY      BACKBONE
POLICY
-----
-         10        -                            TRUE                dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for the I-SID list listisids:

```
Switch:1(router-vrf)#isis accept isid-list listisids backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID      ISID-LIST                ENABLE POLICY      BACKBONE
POLICY
-----
-         10        listisids                TRUE                dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure the default accept policy for IS-IS and DvR:

```
Switch:1(router-vrf)#route-map isismap3 1
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#isis accept route-map isismap3 backbone-route-map
dvrmap3
Switch:1(router-vrf)#
```

```
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID      ISID-LIST                ENABLE POLICY      BACKBONE
POLICY
-----
-         -         -                            TRUE isismap3        dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure the default accept policy for DvR:

```
Switch:1(router-vrf)#isis accept backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID      ISID-LIST                ENABLE POLICY      BACKBONE
POLICY
-----
-         -         -                            TRUE                dvrmap3
```

```
1 out of 1 Total Num of Isis Accept Policies displayed
```

Example 2:

Configure an accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24
```

2. For a specific VRF instance, create a route policy to match the IP prefix list:

```
Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap4 1
Switch:1(router-vrf-route-map)#match network listPrefix
Switch:1(router-vrf-route-map)#enable
Switch:1(router-vrf-route-map)#exit
Switch:1(router-vrf)#
```

3. Create an IS-IS accept policy with I-SID 10, and apply the route policy as the backbone route policy:

```
Switch:1(router-vrf)#accept i-sid 10 backbone-route-map dvrmap4
Switch:1(router-vrf)#accept i-sid 10 enable
```

4. Apply the IS-IS accept policy:

```
Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept
```

5. Verify the configuration:

```
Switch:1(config)#exit
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====

ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
POLICY
-----
-        -        -                        TRUE                     dvrmap4

1 out of 1 Total Num of Isis Accept Policies displayed
```

Variable definitions

Use the data in the following table to use the `ip isid-list` command.

Variable	Value
<code>WORD<1-32></code>	Creates a name for your I-SID list.
<code><1-16777215></code>	Specifies an I-SID number.
<code>list WORD<1-1024></code>	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.
backbone-route-map <i>WORD</i> <1-64>	Specifies the DvR backbone route map.
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 <i>WORD</i> <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 <i>WORD</i> <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 <i>WORD</i> <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>][vrfigs 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:

```
Switch:1#show ip route vrf test
```

```
=====
```

IP Route - VRF test									
DST	MASK	NEXT	NH VRF/ISID	COST	INTER 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	Switch13	GRT	10	1000	ISIS	0	IBSV	7
1.1.1.200	255.255.255.255	Switch200	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	Switch13	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	VRF ISID	DEST ISID	Destination	NH	BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST	IP ROUTE PREFERENCE
GRT	-	101	1.1.1.13/32	Switch13	1000	1/7	10	44	7	
GRT	-	101	1.1.1.13/32	Switch13	1001	1/7	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 <i>WORD</i> <1-16>	Displays I-SID list information for a particular VRF by name.
vrfids <i>WORD</i> <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 <i>WORD</i> <1-16>	Displays I-SID list information for a particular VRF by name.
vrfids <i>WORD</i> <0-512>	Displays I-SID list information for a particular VRF ID.
<i>WORD</i> <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 <i>WORD</i> <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 Forwarding Information Base (FIB) information.
id <1-16777215>	Displays IS-IS SPBM IP unicast FIB information by I-SID ID.

Table continues...

Variable	Value
spb-m-nh-as-mac	Displays the next hop B-MAC of the IP unicast FIB entry.

Configuring IP Multicast over Fabric Connect within the GRT

Use this procedure to configure IP Multicast over Fabric Connect within the GRT. 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 Fabric Connect 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 Fabric Connect 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 Fabric Connect within the GRT, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP Multicast over Fabric Connect enabled. When you enable IP Multicast over Fabric Connect on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must configure `ip spb-multicast enable` on each of the VLANs within the GRT that need to support IP multicast traffic. The default is disabled. After you enable IP Multicast over Fabric Connect on each of the VLANs within the GRT that need to support IP multicast traffic, any IGMP functions required for IP Multicast over Fabric Connect within the GRT are automatically enabled. You do not need to configure anything IGMP related.

If you only want to use IP Multicast over Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect routing does not depend on unicast routing, which allows for you to more easily migrate from a PIM environment to IP Multicast over Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect first and then migrate unicast separately or not at all.

The switch only supports IPv4 addresses with IP Multicast over Fabric Connect.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][, ...] or interface vlan <1-4059>
```

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create an IP interface on the VLAN:

```
ip address <A.B.C.D/X>
```

3. Enable IP Multicast over Fabric Connect:

```
ip spb-multicast enable
```

*** Note:**

After you configure `ip spb-multicast enable`, 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 IP Multicast over Fabric Connect is enabled, an error message appears.

*** Note:**

When you configure `ip spb-multicast enable` on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For information on DvR, see *Configuring IPv4 Routing*.

4. (Optional) Disable IP Multicast over Fabric Connect:

```
no ip spb-multicast enable
default ip spb-multicast enable
```

5. Ensure IP Multicast over Fabric Connect within the GRT is configured properly:

```
show ip igmp interface
```

If `routed-spb` appears under mode, IP Multicast over Fabric Connect within the GRT is properly enabled on the VLAN.

Example

Enable IP Multicast over Fabric Connect within the GRT:

```
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
=====
IF          QUERY  OPER  QUERY  WRONG  LASTMEM
INTVL  STATUS  VERS.  VERS  QUERIER  MAXRSPT  QUERY  JOINS  ROBUST  QUERY  MODE
-----
V500    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
<1-4059>	Specifies the VLAN ID.

Use the data in the following table to use the `interface GigabitEthernet` command.

Variable	Value
{slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

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 Fabric Connect services and IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect globally.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable

configure terminal

router isis
```

2. Configure the IP Multicast over Fabric Connect forward-cache timeout:

```
spbm <1-100> multicast fwd-cache-timeout(seconds) <10-86400>
```

3. **(Optional)** Configure the IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect 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 300
```

Variable definitions

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

Variable	Value
<1-100>	Specifies the SPBM instance. The switch only supports one instance.
<10-86400>	Specifies the IP Multicast over Fabric Connect forward-cache timeout in seconds. The default is 210 seconds.

Viewing IP Multicast over Fabric Connect within the GRT information

Use the following options to display IP Multicast over Fabric Connect within the GRT information to confirm proper configuration.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display all IP Multicast over Fabric Connect route information:
3. Display detailed IP Multicast over Fabric Connect route information:
4. Display the IP Multicast over Fabric Connect multicast group and source address information:
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 spbm ip-multicast-route [all]
```

```
show isis spbm ip-multicast-route [detail]
```

```
show isis spbm ip-multicast-route [group {A.B.C.D}] [source {A.B.C.D}] [source-beb WORD<0-255>]
```

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

Example

Display IP Multicast over Fabric Connect within the GRT 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 e12
routed GRT      501 192.0.2.1 233.252.0.2 5010 16300002 20 e12
routed GRT      501 192.0.2.1 233.252.0.3 5010 16300003 10 e12
routed GRT      501 192.0.2.1 233.252.0.4 5010 16300004 20 e12
routed GRT      501 192.0.2.1 233.252.0.5 5010 16300005 10 e12
routed GRT      501 192.0.2.1 233.252.0.6 5010 16300006 20 e12
routed GRT      501 192.0.2.1 233.252.0.7 5010 16300007 10 e12
routed GRT      501 192.0.2.1 233.252.0.8 5010 16300008 20 e12
routed GRT      501 192.0.2.1 233.252.0.9 5010 16300009 10 e12
routed GRT      501 192.0.2.1 233.252.0.10 5010 16300010 20 e12
=====

```

```

-----
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     1/3   V604:9/38  e12
192.0.2.10     233.252.0.2   16300002  20     1/2,1/3 V604:9/38  e12
192.0.2.10     233.252.0.3   16300003  10     1/3   V604:9/38  e12
192.0.2.10     233.252.0.4   16300004  20     1/2,1/3 V604:9/38  e12
192.0.2.10     233.252.0.5   16300005  10     1/3   V604:9/38  e12
192.0.2.10     233.252.0.6   16300006  20     1/2,1/3 V604:9/38  e12
192.0.2.10     233.252.0.7   16300007  10     1/3   V604:9/38  e12
192.0.2.10     233.252.0.8   16300008  20     1/2,1/3 V604:9/38  e12
192.0.2.10     233.252.0.9   16300009  10     1/3   V604:9/38  e12
192.0.2.10     233.252.0.10 16300010  20     1/2,1/3 V604:9/38  e12
=====

```

```

-----
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 ADDRESS I-SID  BVID  FRAG  NAME
-----
GRT    192.0.2.1 233.252.0.1 16300001 10 0x0 e12
GRT    192.0.2.1 233.252.0.3 16300003 10 0x0 e12
GRT    192.0.2.1 233.252.0.5 16300005 10 0x0 e12
GRT    192.0.2.1 233.252.0.7 16300007 10 0x0 e12
GRT    192.0.2.1 233.252.0.9 16300009 10 0x0 e12
GRT    192.0.2.1 233.252.0.2 16300002 20 0x0 e12
GRT    192.0.2.1 233.252.0.4 16300004 20 0x0 e12
GRT    192.0.2.1 233.252.0.6 16300006 20 0x0 e12
GRT    192.0.2.1 233.252.0.8 16300008 20 0x0 e12
GRT    192.0.2.1 233.252.0.10 16300010 20 0x0 e12

```


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 Fabric Connect route information.
detail	Displays detailed IP Multicast over Fabric Connect route information.
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 Fabric Connect route. If you select source it will also display the source IP address. Specifies the source BEB name.
vlan	Displays IP Multicast over Fabric Connect route information by VLAN.
vrf	Displays IP Multicast over Fabric Connect route information by VRF.
vsn-isid	Displays IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect summary for a given host-name.
lspid <xxxx.xxxx.xxxx.xx-xx>	Displays the IP Multicast over Fabric Connect 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: <ul style="list-style-type: none"> • routed—For GRT 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 Fabric Connect route.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.

Table continues...

Parameter	Description
VSN-ISID	Specifies the GRT because IP Multicast over Fabric Connect within the GRT does not use a VSN I-SID.
Data ISID	Specifies the data I-SID for the IP Multicast over Fabric Connect 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 Fabric Connect route.
Source-BEB	Specifies the source BEB for the IP Multicast over Fabric Connect 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 Fabric Connect route.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
Data ISID	Specifies the data I-SID for the IP Multicast over Fabric Connect 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 Fabric Connect route.
NNI Rcvrs	Specifies the NNI receivers.
UNI Rcvrs	Specifies the UNI receivers.
Source-BEB	Specifies the source BEB for the IP Multicast over Fabric Connect 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.

Table continues...

Parameter	Description
SOURCE ADDRESS	Specifies the group IP address for the IP Multicast over Fabric Connect route.
GROUP ADDRESS	Specifies the group IP address for the IP Multicast over Fabric Connect 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 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 multicast over Fabric Connect within the GRT

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[/sub-port]}[-slot/
port[/sub-port]}[,...]] [vlan <1-4059>] [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] [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 multicast over Fabric Connect within the GRT:

```
Switch:#enable
Switch:1#show ip igmp interface
```

```
=====
                          Igmp Interface - GlobalRouter
=====
```

IF	QUERY INTVL	STATUS	VERS.	OPER VERS	QUERIER	QUERY MAXRSPT	WRONG QUERY	JOINS	ROBUST	LASTMEM 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 ID	QUERY INTVL	QUERY MAX RESP	ROBUST	VERSION	LAST MEMB QUERY	PROXY SNOOP ENABLE	SNOOP ENABLE	SSM SNOOP ENABLE	FAST LEAVE ENABLE	FAST LEAVE PORTS
1	125	100	2	2	10	false	false	false	false	

VLAN ID	SNOOP QUERIER ENABLE	SNOOP QUERIER ADDRESS	DYNAMIC DOWNGRADE VERSION	COMPATIBILITY MODE	EXPLICIT HOST TRACKING
1	false	0.0.0.0	enable	disable	disable

```
Switch:1# show ip igmp sender
```

```
=====
                          IGMP Sender - GlobalRouter
=====
```

GRPADDR	IFINDEX	MEMBER	PORT/ 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[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vlan <1-4059>	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.
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

Table continues...

Parameter	Description
	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 Fabric Connect 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.
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.

Table continues...

Parameter	Description
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.

Table continues...

Parameter	Description
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.

The following table describes the fields for the `show ip igmp snoop-trace` command.

Parameter	Description
GROUP ADDRESS	Indicates the IP multicast group address for which this entry contains information.
SOURCE ADDRESS	Indicates the source of the multicast traffic.
IN VLAN	Indicates the incoming VLAN ID.
IN PORT	Indicates the incoming port number.
OUT VLAN	Indicates the outgoing VLAN ID.
OUT PORT	Indicates the outgoing port number.
TYPE	Indicates where the stream is learned. ACCESS indicates the stream is learned on UNI ports. NETWORK indicates the stream is learned over the SPBM network.

Viewing TLV information for IP Multicast over Fabric Connect within the GRT

Use the following commands to check TLV information.

For IP Multicast over Fabric Connect within the GRT, 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 <1-236> [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 <1-236> [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 {1, 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.
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 <1-236>	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.

Table continues...

Parameter	Description
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 a Circuitless IPv4 interface

About this task

You can use a circuitless IPv4 (CLIPv4) interface to provide uninterrupted connectivity to your system.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IP**.
3. Click the **Circuitless IP** tab.
4. Click **Insert**.
5. In the **Interface** field, assign a CLIP interface number.
6. Enter the IP address.
7. Enter the network mask.
8. Click **Insert**.
9. To delete a CLIP interface, select the interface and click **Delete**.

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.
Ip Address	Specifies the IP address of the CLIP.
Net Mask	Specifies the network mask.

Configuring a Circuitless IPv6 interface

About this task

You can use a circuitless IPv6 (CLIPv6) interface to provide uninterrupted connectivity to your system.

Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **IPv6**.
3. Click the **Circuitless IP** tab.
4. Click **Insert**.
5. In the **Interface** field, assign a CLIP interface number.
6. Type the IPv6 address and prefix length.

Circuitless IPv6 field descriptions

Use the data in the following table to use the **Circuitless IPv6** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
Addr	Specifies the IPv6 address to which this entry applies.
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.

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 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 and static 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.
- You must configure a circuitless IP (CLIP) interface:
 - To configure an IPv4 CLIP interface, see [Configuring a Circuitless IPv4 interface](#) on page 385
 - To configure an IPv6 CLIP interface, see [Configuring a Circuitless IPv6 interface](#) on page 386

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **IS-IS**.
3. From the **Globals** tab, in the **IpSourceAddress** field, specify the CLIP interface to use as the source address for SBPM IP Shortcuts.

*** Note:**

For IPv6 Shortcuts, select **ipv6** in the **Ipv6SourceAddressType** field, and then use the **Ipv6SourceAddress** field to specify the CLIPv6 interface to use as the source address for SBPM IPv6 Shortcuts.

4. Click **Apply**.
5. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
6. Click the **SPBM** tab.
7. In the **IpShortcut** field select **enable**.

*** Note:**

For IPv6 Shortcuts, select **enable** in the **Ipv6Shortcut** field.

8. Click **Apply**.
9. In the navigation pane, expand the **Configuration > IP** folders.
10. Click **Policy**.
11. Click the **Route Redistribution** tab.
12. Click **Insert** to identify routes on the local switch to be announced into the SPBM network.
13. 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.
14. 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 pane, expand the **Configuration > IP** folders.
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
DstVrflid	Specifies the destination Virtual Routing and Forwarding (VRF) ID used in the redistribution.
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrflid	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. 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.

Table continues...

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.

Enabling IPv6 IS-IS redistribution

Enable redistribution to announce routes of a certain source protocol type into the IPv6 IS-IS domain.

You can redistribute directly connected routes, IPv6 static routes, IPv6 BGP routes, OSPFv3 routes, and RIPng routes into IPv6 IS-IS.

Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **IS-IS**.
3. Click the **Redistribute** tab.
4. For the type of route source, double-click the cell in the **Enable** column to change the value.
5. Select **enable**.
6. Click **Apply**.

Redistribute field descriptions

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

Name	Description
DstVrflid	Shows the ID of the destination virtual router and forwarder (VRF). Because IPv6 is not virtualized, the value is 0 for the Global Router.
Protocol	Shows the routing protocol that receives the external routing information. In this case, the routing protocol is IPv6 IS-IS.
SrcVrflid	Shows the ID of the source VRF. Because IPv6 is not virtualized, the value is 0 for the Global Router.
RouteSource	Shows the source protocol from which to receive routes to insert into the IPv6 IS-IS domain. The possibilities are direct routes and static routes.
Enable	Configures the status of route redistribution. The default is disable.

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. You should 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 pane, expand the **Configuration > IP** folders.
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.
DefaultBackbonePolicyName	Specifies the backbone host 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 IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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. The value is 2.5 bytes in the format <x.xx.xx>.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.

Table continues...

Name	Description
PolicyName	Specifies a route policy. You must configure a policy earlier in a separate procedure.
BackbonePolicyName	Specifies the route policy for the backbone routes. 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 IPv4 Routing*.

Procedure

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

Table continues...

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.
BackbonePolicyName	Specifies the backbone 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 IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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 data 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. The value is 2.5 bytes in the format <x.xx.xx>.
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.
BackBonePolicyName	Specifies the backbone 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 pane, expand the **Configuration > IP** folders.
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.

Table continues...

Name	Description
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 IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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.
BackBonePolicyName	Specifies the backbone 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 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.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

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 pane, expand the **Configuration > IP** folders.
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	<p>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.</p> <p>The value is 2.5 bytes in the format <x.xx.xx>.</p>

Table continues...

Name	Description
Name	Specifies the name of the I-SID list used to 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 name.
BackBonePolicyName	Specifies a backbone route policy name.

Configuring IP Multicast over Fabric Connect on a VLAN within the GRT

Use this procedure to enable IP Multicast over Fabric Connect 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 Fabric Connect, see [Configuring IP Multicast over Fabric Connect on a Layer 3 VSN](#) on page 467.

* 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 Fabric Connect 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 Fabric Connect 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 Fabric Connect within the GRT, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP Multicast over Fabric Connect enabled. When you enable IP Multicast over Fabric Connect on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must enable IP Multicast over Fabric Connect on each of the VLANs within the GRT that need to support IP multicast traffic. After you enable IP Multicast over Fabric Connect on the VLANs, any IGMP functions required for IP Multicast over Fabric Connect within the GRT are automatically enabled. You do not need to configure anything IGMP related.

If you only want to use IP Multicast over Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect routing within the GRT does not depend on unicast routing. This allows for you to more easily migrate from a PIM environment to IP Multicast over Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect first and then migrate unicast separately or not at all.

The switch only supports IPv4 addresses with IP Multicast over Fabric Connect.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **VLANs**.
3. Choose a VLAN, and then click the **IP** button.
4. Click the **SPB Multicast** tab.

*** Note:**

After you enable IP Multicast over Fabric Connect, 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.

*** Note:**

When you enable IP Multicast over Fabric Connect on a Controller switch in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For information on DvR, see *Configuring IPv4 Routing*.

5. Click **Enable**.
6. Click **Apply**.

Configuring IP Multicast over Fabric Connect on a brouter port within the GRT

Use this procedure to enable IP Multicast over Fabric Connect on a brouter port IP interface. The default is enabled.

To configure a brouter port for a VRF with IP Multicast over Fabric Connect, see [Configuring IP Multicast over Fabric Connect on a brouter port for L3 VSN](#) on page 468.

*** 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 Fabric Connect 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 Fabric Connect 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 Fabric Connect within the GRT, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP Multicast over Fabric Connect enabled. When you

enable IP Multicast over Fabric Connect on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must enable IP Multicast over Fabric Connect on each of the VLANs within the GRT that need to support IP multicast traffic. After you enable IP Multicast over Fabric Connect on the VLANs, any IGMP functions required for IP Multicast over Fabric Connect within the GRT are automatically enabled.

If you only want to use IP Multicast over Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect routing does not depend on unicast routing, which allows for you to more easily migrate from a PIM environment to Multicast over Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect first, and then migrate unicast separately or not at all.

The switch only supports IPv4 addresses with IP Multicast over Fabric Connect.

Procedure

1. Select an enabled port on the Physical Device View.
2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
3. Click **IP**.
4. Click the **SPB Multicast** tab.
5. Click **Enable**.

Note:

When you enable IP Multicast over Fabric Connect on a DvR Controller switch in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For information on DvR, see *Configuring IPv4 Routing*.

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 pane, expand the **Configuration > IP** folders.
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 Fabric Connect globally.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
2. Click the **SPBM** tab.
3. Modify the **McastFwdCacheTimeout** value.
4. Click **Apply**.

IP Shortcuts SPBM configuration example

The following figure shows a sample IP Shortcuts over SPBM deployment.

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.




Figure 53: 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 223.

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 static routes). To inject IPv6 routes, you must enable route redistribution of IPv6 direct, IPv6 static, and OSPFv3 routes into IS-IS.
- No IP address needs to be configured on SwitchG.

The following sections show the steps required to configure the SPBM IP Shortcuts parameters in this example.

SwitchC

```
CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter
```

```
interface loopback 1
ip address 1 10.0.0.1/255.255.255.255
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 1/2 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
```

SwitchD

```
CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter
```

```
interface loopback 1
ip address 1 10.0.0.2/255.255.255.255
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 1/2
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 — SwitchC

```
SwitchC:1# show isis spbm ip-unicast-fib
```

```
=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
```

VRF	ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST
GRT	-	10.0.0.2/32	SwitchD	4000	1/30	20	1
GRT	-	10.0.14.1/24	SwitchD	4000	1/30	20	1

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

```
SwitchC:1# show ip route
```

```
=====
                        IP Route - GlobalRouter
=====
```

DST	MASK	NEXT	VRF	NH COST	INTER 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	SwitchD	Glob~	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	SwitchD	Glob~	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 Route,
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 — SwitchD

```
SwitchD:1# show isis spbm ip-unicast-fib
```

```
=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
```

VRF	ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST
GRT	-	10.0.0.1/32	SwitchC	4000	1/20	20	1
GRT	-	10.0.13.1/24	SwitchC	4000	1/20	20	1

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

```
SwitchD:1# show ip route
```

```
=====
                        IP Route - GlobalRouter
=====
DST          MASK          NEXT          VRF    NH    INTER
          COST  FACE  PROT  AGE  TYPE  PRF
-----
10.0.0.1    255.255.255.255  SwitchC      Glob~  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    SwitchC      Glob~  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 Route,
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 multicast over SPBM within the GRT configuration example

The following example 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 54: 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.

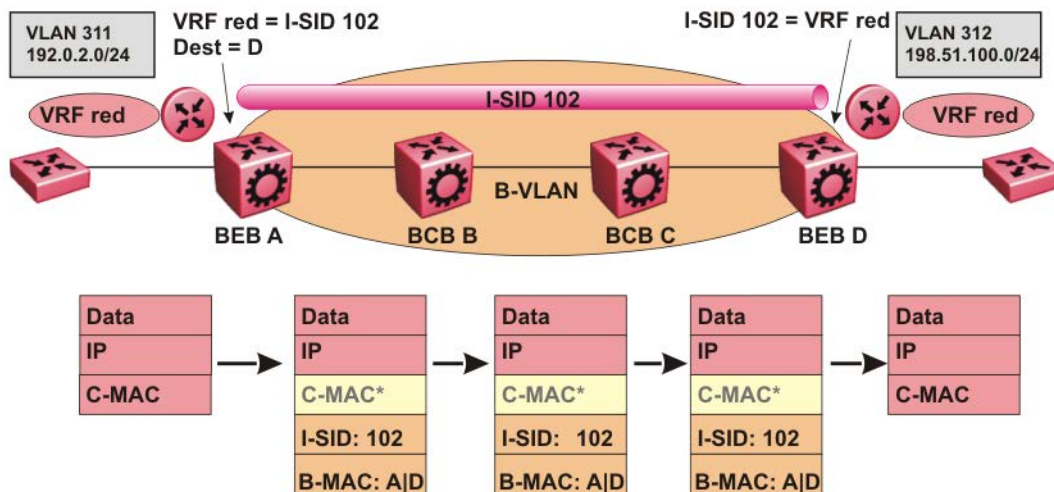


Figure 55: Packet forwarding in SPBM Layer 3 VSN

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 C-MAC 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 B-MAC 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 56: 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 57: 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:**

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 58: Redistributing routes into RIP

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 1/12
ip rip default-supply enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
    
```

Switch-A

```
RIP PORT CONFIGURATION

interface gigabitethernet 1/2
ip rip default-listen enable
exit

interface gigabitethernet 1/3
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 1/12
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 policies and DvR

When you configure DvR in an SPB network, you can leverage IS-IS accept policies to control the DvR routes learned from the DvR backbone. The DvR backbone contains the master list of all the host routes learned from various DvR domains.

You can configure accept policies on a DvR Controller or a non-DvR BEB as a filter to determine which DvR host routes to accept into the routing table, from the DvR backbone. Accept policies apply to only those backbone (or inter-domain) host routes that are not part of the Controller's own DvR enabled subnets *and* do not have the same domain ID as that of the Controller.

For non-DvR BEBs, all the routes present in the backbone are learned, but you can still use the accept policies to filter specific routes.

For information on DvR, see *Configuring IPv4 Routing*.

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
- backbone-route-map
- a combination of route-map and backbone-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.

On DvR Controllers in a DvR domain, you can configure a backbone route policy to determine what host routes to accept from the DvR backbone, into the routing table. Also, just like on the route policy, you can configure match criteria, and set preferences on the backbone route policy.

To accept both IS-IS routes and host routes from the DvR backbone, you can configure both a route policy and a backbone route policy in the accept policy instance.

For more information on configuring route policies, see *Configuring IPv4 Routing*.

The following table describes IS-IS accept policy filters.

Filters into	Filter	Description
Global Routing Table (GRT)	accept route-map <i>WORD</i> <1-64>	By default, the device accepts all routes into the GRT and VRF routing table. This is the default accept policy.
	accept route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	This is the default accept policy with configuration to accept specific DvR host routes from the DvR backbone.
	accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.

Table continues...

Filters into	Filter	Description
	accept isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the list of I-SIDs. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
Virtual Routing and Forwarding (VRF) routing table	isis accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept i-sid <0-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept adv-rtr <x.xx.xx> i-sid <0-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	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). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept adv-rtr <x.xx.xx> isid-list <i>WORD</i> <1-32> route-map	The device filters based on the specific advertising BEB and the list of I-SIDs.

Table continues...

Filters into	Filter	Description
	<code>WORD<1-64> backbone-route-map WORD<1-64></code>	The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	<code>isis accept route-map WORD<1-64> route-map WORD<1-64> backbone-route-map WORD<1-64></code>	The device filters based on the route policy. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.

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.
- 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 the configuration of route-map, see *Configuring IPv4 Routing*.

Layer 3 VSN with IP Multicast over Fabric Connect

IP Multicast over Fabric Connect supports Layer 3 VSN functionality where multicast traffic is bridged over the SPBM core infrastructure. Layer 3 VSN using IP Multicast over Fabric Connect 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 Fabric Connect 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 Fabric Connect 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 Fabric Connect is only configured on BEBs.

*** Note:**

- 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.
- If you only want to use IP Multicast over Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect 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 Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect 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 Fabric Connect, 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 Fabric Connect enabled. If a sender transmits a multicast stream to a BEB on a Layer 3 VSN with IP Multicast over Fabric Connect 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 Fabric Connect, 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.

DvR

On DvR Controllers in a DvR domain, you must manually configure IP multicast over Fabric Connect on Layer 3 VSNs (VRFs). This configuration is then automatically pushed to the Leaf nodes in the DvR domain.

For more information on DvR, see *Configuring IPv4 Routing*.

Enable/disable ICMP Response on VRFs/L3 VSNs

This feature supports VRFs/L3 VSNs to operate in stealth mode by disabling ICMP responses on specific VRFs/L3 VSNs.

If the ICMP response is disabled, the switch does not respond to any ICMP requests received on the VRFs/L3 VSNs.

If the ICMP response is enabled, the switch responds to ICMP requests received on the VRF/L3 VSNs.

Layer 3 VSN configuration using the CLI

This section provides a procedure to configure Layer 3 VSNs using the command line interface (CLI).

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 IPv4 Routing*.
- 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  11
```

Variable definitions

Use the data in the following table to configure the **show ip ipvpn** command.

Variable	Value
vrf <i>WORD<1-16></i>	Specifies the VRF name.
vrfids <i>WORD<0-512></i>	Specifies the VRF ID.

Use the data in the following table to configure the **i-sid** command.

Variable	Value
<i><0-16777215></i>	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
<i>{direct bgp ospf rip static}</i>	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.

Table continues...

Variable	Value
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.

If DvR is enabled on your switch, and the switch is either a DvR Controller or a non-DvR BEB within the domain, you can configure IS-IS accept policies to accept specific host routes from the DvR backbone. For information on DvR, see *Configuring IPv4 Routing*.

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. You should 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.
- The number of unique Layer 3 VSN I-SIDs used on a BEB is limited to the number of VRFs supported on the switch. This includes the I-SID values used for Layer 3 VSNs and the I-SID values specified for the ISIS accept policy filters, which can be configured using the `ip`

`isis-list [ISID#], accept i-sid <value>`, or `accept adv-rtr <isis nn> i-sid <value>` commands.

The switch supports 24 VRFs by default, so, in a default configuration, you cannot create an ip isid-list or accept policy with more than 24 unique I-SID entries. However, the configured VRFs take up an entry, so the formula to calculate the limit is: [24 VRF Limit – (currently configured VRFs)]. This gives the number of unique I-SIDs that can be used directly in the IS-IS accept policy filters, which you implement with the `ip isid-list` or `accept policy` command. The I-SIDs used for Layer 3 VSNs can be reused in IS-IS accept policy filters without affecting the limit.

If you increase the VRF scaling, you can create more Layer 3 VSNs. For more information about how to increase the number of supported VRFs, see *Configuring IPv4 Routing*. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Release Notes*.

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 IPv4 Routing*.
- Ensure that DvR is enabled on the switch before you configure an IS-IS accept policy with a backbone route policy, to accept host routes from the DvR backbone.

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
```

Configure IS-IS accept policies with a route policy or a backbone route policy or a combination of both, to determine which routes the IS-IS accept policy applies to.

Configure one of the following types of IS-IS accept policies.

- **An IS-IS accept policy with only the route policy:**

The IS-IS routes are selectively accepted based on the route policy. Since the backbone route policy is not configured, all host routes from the DvR backbone are *denied*.

If you do not configure a route policy, by default, all IS-IS routes are *accepted*.

- **An IS-IS accept policy with only the backbone route policy:**

The DvR host routes from the DvR backbone are selectively accepted based on the backbone route policy. Since the route policy is not configured, all IS-IS host routes are accepted.

If you do not configure a backbone route policy, all host routes from the DvR backbone are *denied*.

- **An IS-IS accept policy with both route policy and backbone route policy:**

IS-IS routes are selectively accepted based on the route policy and host routes from the DvR backbone are selectively accepted based on the backbone route policy.

5. Configure an IS-IS accept policy instance with a route policy.

Use one of the following options:

- a. 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>]
```

- b. 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>]
```

- c. **(Optional)** Delete an IS-IS accept policy instance:

```
no accept [adv-rtr <x.xx.xx>][i-sid <1-16777215>][isid-list WORD <1-32>]
```

- d. Specify an IS-IS route policy to apply to routes from all BEBs:

```
accept route-map WORD<1-64>
```

- e. Specify an IS-IS route policy to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx>[route-map WORD<1-64>]
```

- f. **(Optional)** Delete an IS-IS route policy:

```
no accept [adv-rtr <x.xx.xx>] [route-map]
```

- g. Enable an IS-IS route accept instance:

```
accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list WORD<1-32>]
```

- h. **(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>]
```

6. Configure an IS-IS accept policy instance with a backbone route policy to accept host routes from the DvR backbone:

*** Note:**

IS-IS accept policies typically apply to all IS-IS routes. However, to accept DvR host routes from the DvR backbone, you *must* explicitly configure the IS-IS accept policy with a backbone route policy.

Use one of the following options:

- a. Create the default IS-IS accept policy instance to accept host routes from the DvR backbone:

```
accept backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy instance with backbone route policy configuration:

```
no accept backbone-route-map
```

- c. Create an IS-IS accept policy instance to accept host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-  
route-map WORD<1-64>
```

- d. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-  
route-map
```

- e. Create an IS-IS accept policy instance to accept host routes from the DvR backbone and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> backbone-route-map WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to a specific advertising BEB

```
no accept adv-rtr <x.xx.xx> backbone-route-map
```

7. Configure an IS-IS accept policy with both route policy and backbone route policy, to selectively accept IS-IS routes as well as host routes from the DvR backbone.

- a. Create the default IS-IS accept policy instance with a route policy to accept IS-IS routes and a backbone route policy to accept host routes from the DvR backbone:

```
accept route-map WORD<1-32> backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy with route policy and backbone route policy configuration:

```
no accept route-map backbone-route-map
```

- c. Create an accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map  
WORD<1-32> backbone-route-map WORD<1-64>
```

- d. **(Optional)** Delete an accept policy instance with route policy and backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map
backbone-route-map
```

- e. Create an IS-IS accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> route-map WORD<1-32> backbone-route-map
WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with route policy and backbone route policy configuration, which applies to a specific advertising BEB:

```
no accept adv-rtr <x.xx.xx> route-map backbone-route-map
```

8. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
isis apply accept [vrf WORD <1-16>]
```

9. Exit IS-IS Router Configuration mode:

```
exit
```

You are in Global Configuration mode.

Example

Configure an I-SID based IS-IS accept policy with the route policy `test`:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map test 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit

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
```

Configuration of IS-IS accept policy to accept host routes from the DvR backbone

Example 1:

To accept host routes from the DvR backbone, you must configure a backbone route policy and apply it to the IS-IS accept policy.

1. Configure a route policy for DvR:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map dvrmap1 1
Switch:1(route-map)#enable
```

2. Configure an IS-IS accept policy for I-SID 10, and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap1
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
```

OR

Configure the default accept policy for IS-IS and DvR, and apply the route policy as a backbone route policy:

```
Switch:1(config)#route-map isismap1 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept route-map isismap1 backbone-route-map dvrmap1
```

3. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
Switch:1(config)#exit
```

4. Verify the configuration:

```
Switch:1#show ip isis accept
=====
Isis Accept - GlobalRouter
=====
ADV_RTR  I-SID    ISID-LIST          ENABLE POLICY      BACKBONE
POLICY
-----
-         10      -                  TRUE               dvrmap1
-         -       -                  isismap1          dvrmap1
2 out of 2 Total Num of Isis Accept Policies displayed
```

Example 2:

Configure an IS-IS accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24
```

2. Create the route policy dvrmap2 to match the IP prefix list:

```
Switch:1(config)#route-map dvrmap2 1
Switch:1(route-map)#match network listPrefix
Switch:1(route-map)#enable
```

3. Create an IS-IS accept policy with I-SID 10 and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
```

4. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

The above command causes IS-IS to accept all routes with I-SID 10. To deny IS-IS routes and accept only DvR host routes, you can configure an additional IS-IS route policy as follows:

```
Switch:1(config)#route-map isismap2 1
Switch:1(route-map)#no permit
Switch:1(route-map)#enable

Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 route-map isismap2 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

5. Verify the configuration:

```
Switch:1(config)#exit
Switch:1#show ip isis accept
```

```
=====
                        Isis Accept - GlobalRouter
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
-----
-         10      -                          TRUE  isismap2                dvrmap2
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configuration of IS-IS accept policies for a specific VRF instance**Example 1:**

Configure IS-IS accept policies to accept host routes from the DvR backbone, for a specific VRF instance.

1. In the VRF green context, configure the route policy `dvrmap3` for DvR:

```
Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap3 1
Switch:1(router-vrf-routemap)#enable
```

2. Use one of the following options to configure an IS-IS accept policy, and apply the route policy as a backbone route policy:

Configure an IS-IS accept policy for a specific advertising BEB with nickname `1.11.11`:

```
Switch:1(router-vrf-routemap)#isis accept adv-rtr 1.11.11 backbone-route-map
dvrmap3
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#isis accept adv-rtr 1.11.11 enable
```

```
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
```

```

-----
POLICY
1.11.11 - - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
Switch:1(config)#show ip isis accept vrfids 2
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
1.11.11 - - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

```

Configure an accept policy for I-SID 10:

```

Switch:1(router-vrf)#isis accept i-sid 10 backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
- 10 - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

```

Configure an accept policy for the I-SID list listisids:

```

Switch:1(router-vrf)#isis accept isid-list listisids backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
- 10 listisids TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

```

Configure the default accept policy for IS-IS and DvR:

```

Switch:1(router-vrf)#route-map isismap3 1
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#isis accept route-map isismap3 backbone-route-map
dvrmap3
Switch:1(router-vrf)#
Switch:1(router-vrf)#show ip isis accept vrf green
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE

```

```

-----POLICY
- - - TRUE isismap3 dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

Configure the default accept policy for DvR:

Switch:1(router-vrf)#isis accept backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
- - - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

```

Example 2:

Configure an accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24

```

2. For a specific VRF instance, create a route policy to match the IP prefix list:

```

Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap4 1
Switch:1(router-vrf-routemap)#match network listPrefix
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#

```

3. Create an IS-IS accept policy with I-SID 10, and apply the route policy as the backbone route policy:

```

Switch:1(router-vrf)#accept i-sid 10 backbone-route-map dvrmap4
Switch:1(router-vrf)#accept i-sid 10 enable

```

4. Apply the IS-IS accept policy:

```

Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept

```

5. Verify the configuration:

```

Switch:1(config)#exit
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
- - - TRUE dvrmap4

```


1 out of 1 Total Num of Isis Accept Policies displayed

Variable definitions

Use the data in the following table to use the `ip isid-list` command.

Variable	Value
<code>WORD<1-32></code>	Creates a name for your I-SID list.
<code><1-16777215></code>	Specifies an I-SID number.
<code>list WORD<1-1024></code>	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
<code>adv-rtr <x.xx.xx></code>	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.
<code>backbone-route-map WORD<1-64></code>	Specifies the DvR backbone route map.
<code>enable</code>	Enables an IS-IS accept policy.
<code>i-sid <1-16777215></code>	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).
<code>isid-list WORD<1-32></code>	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).
<code>route-map WORD<1-64></code>	Specifies a route policy by name.

Table continues...

Variable	Value
	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 <i>WORD</i> <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 routes and adds the IS-IS routes again. You should 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 IPv4 Routing*.

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
<code>WORD<1-32></code>	Creates a name for your I-SID list.
<code><0-16777215></code>	Specifies an I-SID value.
<code>list WORD<1-1024></code>	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
<code>adv-rtr <x.xx.xx></code>	Specifies a specific advertising BEB in which to apply the IS-IS accept policy to routes for a specific advertising BEB. <code>x.xx.xx</code> 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.
<code>enable</code>	Enables the IS-IS accept policy.
<code>i-sid <0-16777215></code>	Configures the I-SID to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).
<code>isid-list WORD<1-32></code>	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).
<code>route-map WORD <1-64></code>	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
<code>vrf WORD<1-16></code>	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:

```
Switch:1#show ip route vrf test
```

```
=====
```

IP Route - VRF test									
DST	MASK	NEXT	NH VRF/ISID	COST	INTER 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	Switch13	GRT	10	1000	ISIS	0	IBSV	7
1.1.1.200	255.255.255.255	Switch200	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	Switch13	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	ISID	DEST ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST	IP ROUTE PREFERENCE
GRT	-	101	1.1.1.13/32	Switch13	1000	1/7	10	44	7
GRT	-	101	1.1.1.13/32	Switch13	1001	1/7	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 <i>WORD</i> <1-16>	Displays I-SID list information for a particular VRF by name.
vrfids <i>WORD</i> <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 <i>WORD</i> <1-16>	Displays I-SID list information for a particular VRF by name.
vrfids <i>WORD</i> <0-512>	Displays I-SID list information for a particular VRF ID.
<i>WORD</i> <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 <i>WORD</i> <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 Forwarding 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 Fabric Connect

Use this procedure to configure IP Multicast over Fabric Connect for a Layer 3 VSN.

Configure the Layer 3 VSN (VRF) as a multicast VPN, and then enable IP Multicast over Fabric Connect on VRF VLANs to which IP multicast senders and receivers attach. After you enable IP Multicast over Fabric Connect on VRF VLANs, snooping and proxy on those VLANs is enabled. IGMPv2 at the VLAN level is the default setting. No configuration is required.

* Note:

On DvR Controllers in a DvR domain, you must manually configure IP multicast over Fabric Connect on Layer 3 VSNs (VRFs). This configuration is then automatically pushed to the Leaf nodes in the DvR domain.

For more information on DvR, see *Configuring IPv4 Routing*.

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 Fabric Connect globally.
- You must assign an I-SID for the IPVPN.

About this task

With Layer 3 VSN IP Multicast over Fabric Connect, 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.

The switch only supports IPv4 multicast traffic.

* Note:

You cannot enable IP PIM when IP Multicast over Fabric Connect 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 Fabric Connect for a particular VRF:

```
mvpn enable
```

The default is disabled.

3. **(Optional)** If you want to disable Layer 3 VSN IP Multicast over Fabric Connect, 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[/sub-port] [-slot/port[/sub-
port]][, ...]} or interface vlan <1-4059>
```

*** Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

6. Enable Layer 3 VSN IP Multicast over Fabric Connect 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 Fabric Connect 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
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

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

Variable	Value
GigabitEthernet{slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is

Variable	Value
	channelized, you must also specify the sub-port in the format slot/port/sub-port.

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

Variable	Value
access-list <i>WORD</i> <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. To use the default configuration, use the default option in the command: <code>default ip igmp compatibility-mode</code> , or use the no option to disable compatibility mode: <code>no ip igmp compatibility-mode</code>
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: <code>default ip igmp dynamic-downgrade-version</code> or use the no option to disable downgrade: <code>no ip igmp dynamic-downgrade-version</code>
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[/sub-port][-slot/port[/sub-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.

Table continues...

Variable	Value
	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. 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: <ul style="list-style-type: none"> • maxadvertinterval: 20 seconds • maxinitadvertinterval: 2 seconds • maxinitadvertisements: 3 • minadvertinterval: 15 seconds • neighdeadinterval: 60 seconds
mrouter {slot/port[/sub-port][/-slot/port[/sub-port]][,...]}	Adds multicast router ports.
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). <p>! Important:</p> 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. <p>! Important:</p> To maximize network performance, configure this parameter according to the version of IGMP currently in use: <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable

Table continues...

Variable	Value
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[/sub-port][/-slot/port[/sub-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[/sub-port][/-slot/port[/sub-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.
stream-limit-group {slot/port[/sub-port][/-slot/port[/sub-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 Fabric Connect services and IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect globally.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

2. Configure the IP Multicast over Fabric Connect forward-cache timeout:

```
spbm <1-100> multicast fwd-cache-timeout(seconds) <10-86400>
```

3. **(Optional)** Configure the IP Multicast over Fabric Connect 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 IP Multicast over Fabric Connect 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 300
```

Variable definitions

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

Variable	Value
<1-100>	Specifies the SPBM instance. The switch only supports one instance.
<10-86400>	Specifies the IP Multicast over Fabric Connect forward-cache timeout in seconds. The default is 210 seconds.

Viewing Layer 3 VSN with IP Multicast over Fabric Connect information

Use the following options to display Layer 3 VSN with IP Multicast over Fabric Connect 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 Fabric Connect route information:

```
show isis spbm ip-multicast-route [all][detail]
```

4. Display IP Multicast over Fabric Connect 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 Fabric Connect route information by VRF:

```
show isis spbm ip-multicast-route [vrf WORD<1-16>] [group {A.B.C.D}]
```

6. Display IP Multicast over Fabric Connect route information by VLAN:

```
show isis spbm ip-multicast-route [vlan <1-4059>][detail][group
{A.B.C.D}]
```

7. Display IP Multicast over Fabric Connect 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 Fabric Connect information:

```
Switch:1>enable
Switch:1#show ip vrf mvpn

                Vrf name : green
                mvpn : enable
                fwd-cache-timeout(seconds) : 210

                Vrf name : 4
                mvpn : enable
                fwd-cache-timeout(seconds) : 210

                Vrf name : blue
                mvpn : enable
                fwd-cache-timeout(seconds) : 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
-----
routed GRT      501   192.0.2.1   233.252.0.1   5010      16300001  10     e12
routed GRT      501   192.0.2.1   233.252.0.2   5010      16300002  20     e12
routed GRT      501   192.0.2.1   233.252.0.3   5010      16300003  10     e12
routed GRT      501   192.0.2.1   233.252.0.4   5010      16300004  20     e12
routed GRT      501   192.0.2.1   233.252.0.5   5010      16300005  10     e12
routed GRT      501   192.0.2.1   233.252.0.6   5010      16300006  20     e12
routed GRT      501   192.0.2.1   233.252.0.7   5010      16300007  10     e12
routed GRT      501   192.0.2.1   233.252.0.8   5010      16300008  20     e12
routed GRT      501   192.0.2.1   233.252.0.9   5010      16300009  10     e12
routed GRT      501   192.0.2.1   233.252.0.10  5010      16300010  20     e12
-----
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     e12
192.0.2.10     233.252.0.2   16300002  20     e12
192.0.2.10     233.252.0.3   16300003  10     e12
192.0.2.10     233.252.0.4   16300004  20     e12
192.0.2.10     233.252.0.5   16300005  10     e12
192.0.2.10     233.252.0.6   16300006  20     e12
192.0.2.10     233.252.0.7   16300007  10     e12
192.0.2.10     233.252.0.8   16300008  20     e12
192.0.2.10     233.252.0.9   16300009  10     e12
192.0.2.10     233.252.0.10  16300010  20     e12

```

```

-----
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    e12
routed GRT    501  192.0.2.1  233.252.0.2  5010   16300002  20    e12
routed GRT    501  192.0.2.1  233.252.0.3  5010   16300003  10    e12
routed GRT    501  192.0.2.1  233.252.0.4  5010   16300004  20    e12
routed GRT    501  192.0.2.1  233.252.0.5  5010   16300005  10    e12
routed GRT    501  192.0.2.1  233.252.0.6  5010   16300006  20    e12
routed GRT    501  192.0.2.1  233.252.0.7  5010   16300007  10    e12
routed GRT    501  192.0.2.1  233.252.0.8  5010   16300008  20    e12
routed GRT    501  192.0.2.1  233.252.0.9  5010   16300009  10    e12
routed GRT    501  192.0.2.1  233.252.0.10 5010   16300010  20    e12

```

```

-----
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     e12
192.0.2.1       233.252.0.3   16300003  10     e12
192.0.2.1       233.252.0.4   16300004  20     e12
192.0.2.1       233.252.0.5   16300005  10     e12
192.0.2.1       233.252.0.6   16300006  20     e12
192.0.2.1       233.252.0.7   16300007  10     e12
192.0.2.1       233.252.0.8   16300008  20     e12
192.0.2.1       233.252.0.9   16300009  10     e12
192.0.2.1       233.252.0.10  16300010  20     e12

```

```

-----
Total Number of SPBM IP multicast ROUTE Entries: 10
-----

```



```
Switch:1# show isis spb-mcast-summary
```

```
=====
                        SPB multicast - Summary
=====
```

SCOPE I-SID	SOURCE ADDRESS	GROUP ADDRESS	DATA I-SID	BVID	LSP FRAG	HOST NAME
5010	192.0.2.1	233.252.0.1	16300001	10	0x0	e12
5010	192.0.2.1	233.252.0.3	16300003	10	0x0	e12
5010	192.0.2.1	233.252.0.5	16300005	10	0x0	e12
5010	192.0.2.1	233.252.0.7	16300007	10	0x0	e12
5010	192.0.2.1	233.252.0.9	16300009	10	0x0	e12
5010	192.0.2.1	233.252.0.2	16300002	20	0x0	e12
5010	192.0.2.1	233.252.0.4	16300004	20	0x0	e12
5010	192.0.2.1	233.252.0.6	16300006	20	0x0	e12
5010	192.0.2.1	233.252.0.8	16300008	20	0x0	e12
5010	192.0.2.1	233.252.0.10	16300010	20	0x0	e12

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 Fabric Connect route information.
detail	Displays detailed IP Multicast over Fabric Connect route information.
group{A.B.C.D}	Displays information on the group IP address for the IP Multicast over Fabric Connect route.
vlan<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrfWORD<1-16>	Displays IP Multicast over Fabric Connect route information by VRF.
vsn-isid<1-16777215>	Displays IP Multicast over Fabric Connect route information by I-SID.

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

Variable	Value
host-nameWORD<0-255>	Displays the IP Multicast over Fabric Connect summary information by host-name.

Table continues...

Variable	Value
lspid<xxxx.xxxx.xxxx.xx-xx>	Displays the IP Multicast over Fabric Connect 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 (in seconds) 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 Fabric Connect 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 Fabric Connect route.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect 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 Fabric Connect 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 Fabric Connect route.
Source-BEB	Specifies the source BEB for the IP Multicast over Fabric Connect 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 Fabric Connect route.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect 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 Fabric Connect route.
NNI Rcvrs	Specifies the NNI receivers.
UNI Rcvrs	Specifies the UNI receivers.
Source-BEB	Specifies the source BEB for the IP Multicast over Fabric Connect 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 Fabric Connect route.
GROUP ADDRESS	Specifies the group IP address for the IP Multicast over Fabric Connect route.
DATA I-SID	Specifies the data I-SID for the IP Multicast over Fabric Connect 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[/sub-port]}[-slot/
port[/sub-port]][,...]] [vlan <1-4059>[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 Fabric Connect:

```
Switch:#enable
Switch:1#show ip igmp interface vrf green
=====
                        Igmp Interface - GlobalRouter
=====
IF      QUERY      OPER      QUERY      WRONG      LASTMEM
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    RESP            MEMB  SNOOP  ENABLE  SNOOP  LEAVE  LEAVE
                    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
-----
501   false    0.0.0.0        enable     disable        disable
```

SPBM and IS-IS services configuration

```
Switch:1# show ip igmp sender vrf green
```

```
=====
IGMP Sender - GlobalRouter
=====
GRPADDR      IFINDEX      MEMBER      MLT      PORT/
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[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vlan <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the

Table continues...

Variable	Value
	system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrf <i>WORD</i> <1–16>	Specifies the VRF by name.
vrfids <i>WORD</i> <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 <i>WORD</i> <1–16>	Specifies the VRF by name.
vrfids <i>WORD</i> <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 { <i>A.B.C.D</i> }	Specifies the group address.
member-subnet { <i>A.B.C.D/X</i> }	Specifies the IP address and network mask.
vrf <i>WORD</i> <1–16>	Displays the multicast route configuration for a particular VRF by name.
vrfids <i>WORD</i> <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 { <i>A.B.C.D</i> }	Specifies the group address.
member-subnet { <i>A.B.C.D/X</i> }	Specifies the IP address and network mask.
vrf <i>WORD</i> <1–16>	Displays the multicast route configuration for a particular VRF by name.
vrfids <i>WORD</i> <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.

Table continues...

Parameter	Description
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.
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 Fabric Connect 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.
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.

Table continues...

Parameter	Description
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 Fabric Connect

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 <1-236> [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 <1-236>] [sub-tlv <1-3>] [detail]
```

Example

Display TLV information for a Layer 3 VSN with IP Multicast over Fabric Connect:

```
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: e12
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 {1, 12, 112}	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.
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 <1-236>	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.
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 IPv4 Routing*.
- You must create the Customer VLANs and add slots/ports.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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 pane, expand the **Configuration > IP** folders.
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 **DstVrflid** 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 **SrcVrflid** 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 pane, expand the **Configuration > IP** folders.
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
DstVrflid	Specifies the destination Virtual Routing and Forwarding (VRF) ID used in the redistribution.
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrflid	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. 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. You should 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 pane, expand the **Configuration > IP** folders.
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.
DefaultBackbonePolicyName	Specifies the backbone host 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 IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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. The value is 2.5 bytes in the format <x.xx.xx>.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.

Table continues...

Name	Description
PolicyName	Specifies a route policy. You must configure a policy earlier in a separate procedure.
BackbonePolicyName	Specifies the route policy for the backbone routes. 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 IPv4 Routing*.

Procedure

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

Table continues...

Name	Description
	The value is 2.5 bytes in the format <x.xx.xx>.
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.
BackBonePolicyName	Specifies the backbone 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 pane, expand the **Configuration > IP** folders.
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.

Table continues...

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 IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
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.
BackBonePolicyName	Specifies the backbone 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 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.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

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 pane, expand the **Configuration > IP** folders.
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. The value is 2.5 bytes in the format <x.xx.xx>.

Table continues...

Name	Description
Name	Specifies the name of the I-SID list used to filter.
Enable	Enables or disables the SPBM nicksanme advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy name.
BackBonePolicyName	Specifies a backbone route policy name.

Enabling MVPN for a VRF

Use this procedure to enable MVPN for a particular VRF. IP Multicast over Fabric Connect, 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 IP Multicast over Fabric Connect globally.

Procedure

- In the navigation pane, expand the **Configuration > IP** folders.
- Click **IP-MVPN**.
- Click the **MVPN** tab.
- Double-click in the **Enable** field in the table.
- Select **Enable** from the drop down menu.
- Double-click in the **FwdCacheTimeout** field in the table, and then type the VRF timeout value.
- 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 Fabric Connect 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 Fabric Connect on a VLAN for Layer 3

Use this procedure to enable IP Multicast over Fabric Connect for a Layer 3 VSN. The default is disabled.

To configure a VLAN for IP Shortcuts with IP Multicast over Fabric Connect, see [Configuring IP Multicast over Fabric Connect on a VLAN](#) on page 397.

* Note:

On DvR Controllers in a DvR domain, you must manually configure IP multicast over Fabric Connect on Layer 3 VSNs (VRFs). This configuration is then automatically pushed to the Leaf nodes in the DvR domain.

For more information on DvR, see *Configuring IPv4 Routing*.

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 IP Multicast over Fabric Connect 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 Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect 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 Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect first and then migrate unicast separately or not at all.

The switch only supports IPv4 address with IP Multicast over Fabric Connect.

* Note:

You cannot enable IP PIM when IP Multicast over Fabric Connect is enabled on the VLAN.

Procedure

1. In the navigation pane, expand the **Configuration > VRF Context View** folders.
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. In the navigation pane, expand the **Configuration > VLAN** folders.

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 Fabric Connect on a router port for a Layer 3 VSN

Use this procedure to enable IP Multicast over Fabric Connect on a router port. The default is disabled.

To configure a router port for IP Shortcuts with IP Multicast over Fabric Connect, see [Configuring IP Multicast over Fabric Connect on a router port for IP Shortcuts](#) on page 398.

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 IP Multicast over Fabric Connect 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 Fabric Connect on each of the VLANs that need to support IP multicast traffic.

If you only want to use IP Multicast over Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect routing does not depend on unicast routing, which allows for you to more easily migrate from a PIM environment to Multicast over Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect first, and then migrate unicast separately or not at all.

The switch only supports IPv4 address with IP Multicast over Fabric Connect.

Procedure

1. In the navigation pane, expand the **Configuration > VRF Context View** folders.
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. In the navigation pane, expand the **Configuration > Edit > Port** folders.
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 Fabric Connect 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 Fabric Connect 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. In the navigation pane, expand the **Configuration > VRF Context View** folders.
2. Click **Set VRF Context View**.
3. Choose a VRF name.
4. Click **Launch VRF Context View**.
5. In the navigation pane, expand the **Configuration > VLAN** folders.
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	<p>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.</p> <p>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.)</p> <p>! Important: You must configure this value lower than the QueryInterval.</p>
Robustness	<p>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.</p> <p>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.</p>
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.

Table continues...

Name	Description
	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. Configure this parameter to values greater than 3. If you do not require a fast leave process, 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. 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

Table continues...

Name	Description
	<p>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.</p> <p>Enable Layer 2 Querier on only one node in the VLAN.</p>
SnoopQuerierAddr	<p>Specifies the pseudo IP address of the IGMP Snoop Querier. The default IP address is 0.0.0.0.</p> <p>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.</p>

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 Fabric Connect globally.

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
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 pane, expand the **Configuration > IP** folders.
2. Click **IGMP**.
3. Click the **Interface** tab.

Layer 3 VSN configuration example

The following figure shows a sample Layer 3 VSN deployment.




Figure 59: 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.

VRF green configuration

The following figure shows the green VRF in this Layer 3 VSN example.




Figure 60: Layer 3 VSN — VRF green

The following sections show the steps required to configure the green VRF parameters in this example.

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 1/2 portmember
interface Vlan 101
```

SPBM and IS-IS services configuration

```
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 1/2 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 61: Layer 3 VSN — VRF red

The following sections show the steps required to configure the red VRF parameters in this example.

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 1/2 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 1/2 portmember
```

SPBM and IS-IS services configuration

```
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.

Switch-C

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

```
=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   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 1/3    20    1    7
GRT   -    -    10.0.14.0/24 Switch-D 4000 1/3    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   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 1/2    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   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 1/3    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	VRF ISID	DEST ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST	IP ROUTE PREFERENCE
GRT	-	-	10.0.0.2/32	Switch-D	4000	1/3	20	1	7
GRT	-	-	10.0.14.0/24	Switch-D	4000	1/3	20	1	7
green	-	13990001	10.1.102.0/24	Switch-D	4000	1/3	20	1	7
red	-	13990002	10.2.202.0/24	Switch-D	4000	1/3	20	1	7

```
-----
Total number of SPBM IP-UNICAST FIB entries 4
-----
```

Switch-D

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

```
=====
```

VRF	VRF ISID	DEST ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST	IP ROUTE PREFERENCE
GRT	-	-	10.0.0.1/32	Switch-C	4000	1/2	20	1	7
GRT	-	-	10.0.13.0/24	Switch-C	4000	1/2	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	VRF ISID	DEST ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST	IP ROUTE PREFERENCE
green	-	13990001	10.1.101.0/24	Switch-C	4000	1/2	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	VRF ISID	DEST ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST	IP ROUTE PREFERENCE
red	-	13990002	10.2.201.0/24	Switch-C	4000	1/2	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	VRF ISID	DEST ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST	IP ROUTE PREFERENCE
GRT	-	-	10.0.0.1/32	Switch-C	4000	1/2	20	1	7

SPBM and IS-IS services configuration

```
GRT - - 10.0.13.0/24 Switch-C 4000 1/2 20 1 7
green - 13990001 10.1.101.0/24 Switch-C 4000 1/2 20 1 7
red - 13990002 10.2.201.0/24 Switch-C 4000 1/2 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
=====
DST                MASK                NEXT                NH                INTER
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 Route,
```

```
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
=====
DST                MASK                NEXT                NH                INTER
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 Route,
```

```
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
=====
DST                MASK                NEXT                NH                INTER
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 Route,
```

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
=====
DST                MASK                NEXT                NH                INTER
                   MASK                VRF/ISID            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 Route,
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 Fabric Connect configuration example

The example below shows the configuration to enable IP Multicast over Fabric Connect 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
```

Enable/disable ICMP Response on VRFs/L3 VSNs

This feature supports VRFs/L3 VSNs to operate in stealth mode by disabling ICMP responses on specific VRFs/L3 VSNs.

If the ICMP response is disabled, the switch does not respond to any ICMP requests received on the VRFs/L3 VSNs.

If the ICMP response is enabled, the switch responds to ICMP requests received on the VRF/L3 VSNs.

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 62: Inter-VSN routing

Inter-VSN routing provides a routing hub for Layer 2 Virtual Services Network edge devices, Layer 3 devices, routers, or hosts connected 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 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 route traffic between VLAN 11 (I-SID 2011) and VLAN 12 (I-SID 2012).

IP interfaces are where the routing instance exists. In this case, on Switch-20.

*** Note:**

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

Inter-VSN routing configuration using the CLI

This section provides a procedure to configure Inter-VSN routing using the CLI.

Configuring SPBM Inter-VSN Routing

Inter-VSN allows you to route between IP networks on Layer 2 VLANs with different I-SIDs. Inter-VSN routing is typically used only when you have to extend a VLAN as a Layer 2 Virtual Services Network (VSN) for applications such as vMotion. Normally, it is recommended to use 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 IP interface where the routing instance exists.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Follow the procedures below on the Backbone Edge Bridges (BEBs) containing the VSNs you want to route traffic between.

- a. Create a customer VLAN (C-VLAN) by port:

```
vlan create <2-4059> type port-mstprstp <0-63>
```

- b. Add ports in the C-VLAN:

```
vlan members add <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}
```

- c. Map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID):

```
vlan i-sid <1-4059> <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.

3. On the Backbone Core Bridge (BCB), create a VRF and add a VLAN for each VSN:

- a. Create a VRF:

```
ip vrf WORD<1-16> vrfid <1-511>
```

- b. Create a VLAN to associate with each VSN:

```
vlan create <2-4059> type port-mstprstp <0-63>
```

- c. Enter VLAN Interface Configuration mode:

```
interface vlan <1-4059>
```

- d. Add a VLAN to the VRF you created in step a:

```
vrf WORD<1-16>
```

- e. Associate an I-SID with the VLAN:

```
vlan i-sid <1-4059> <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.

The switch reserves I-SID 0x00ffffff. The switch 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.

- f. Configure an IP address for the VLAN:

```
ip address {A.B.C.D/X}
```

- g. Repeat steps b to f 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
<2-4059>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999.
type port-mstprstp <0-63> [color <0-32>]	Creates a VLAN by port: <ul style="list-style-type: none"> • <0-63> is the STP instance ID. • <i>color</i> <0-32> is the color of the VLAN.

Use the data in the following table to use the `vlan members add` command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the

Table continues...

Variable	Value
	system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port][/-slot/port[/sub-port]][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Use the data in the following table to use the `vlan i-sid` command.

Variable	Value
vlan i-sid <1-4059> <0-16777215>	Specifies the customer VLAN (C-VLAN) to associate with the I-SID. Use the no or default options to remove the I-SID from the specified VLAN.

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 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 an IP address for the VLAN.

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 allows you to route between IP networks on Layer 2 VLANs with different I-SIDs. Inter-VSN routing is typically used only when you have to extend a VLAN as a Layer 2 Virtual Services Network (VSN) for applications such as vMotion. Use 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) that contain 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 pane, expand the **Configuration > VLAN** folders.
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 pane, expand the **Configuration > VLAN** folders.
14. Click **VLANs**.
15. Click the **Advanced** tab.

- To map a C-VLAN to an I-SID, in the **I-sid** field, specify the I-SID to associate with the specified VLAN.

The switch reserves I-SID 0x00ffffff. The switch 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.

- 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.

- Configure the Backbone Core Bridge (BCB) for Inter-VSN Routing. For more information, see [Configuring BCBs for Inter-VSN routing](#) on page 489.

Basic field descriptions

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

Name	Description
Id	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
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: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-vlan • private
MstpInstance	Identifies the MSTP instance.
VrfId	Indicates the Virtual Router to which the VLAN belongs.

Table continues...

Name	Description
VrfName	Indicates the name of the Virtual Router to which the VLAN belongs.
PortMembers	Specifies the slot/port/sub-port of each VLAN member. The sub-port only appears for channelized ports.
ActiveMembers	Specifies the slot/port/sub-port of each VLAN member. The sub-port only appears for channelized ports.
StaticMembers	Specifies the slot/port/sub-port of each static member of a policy-based VLAN. The sub-port only appears for channelized ports.
NotAllowToJoin	Specifies the slot/ports/sub-port that are never allowed to become a member of the policy-based VLAN. The sub-port only appears for channelized ports.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers of remote function call (RFC). If the VLAN type is port-based, none is displayed in the Basic tab ProtocolId field.

Advanced field descriptions


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

Name	Description
Id	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. If you enable VRF scaling and SPBM mode, the system also reserves VLAN IDs 3500 to 3999. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
Name	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.
Type	Specifies the type of VLAN: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-vlan
I-sid	Specifies the I-SID number assigned to a customer VLAN (C-VLAN). The range is 0 — 16777215. The

Table continues...

Name	Description
	default value is 0, which indicates that no I-SID is assigned.
ProtocolId	<p>Specifies the network protocol for protocol-based VLANs.</p> <p>If the VLAN type is not protocol-based, None is displayed in the Basic tab ProtocolId field.</p>
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.
Vlan Operation Action	<p>Performs an operation on the VLAN. The values are:</p> <ul style="list-style-type: none"> • 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 not perform the Snoop-related actions. <p>The default is none.</p>
Result	Specifies the result code after you perform an action. The default is none.
NlbMode	Enables or disables Microsoft Network Load Balancing (NLB) on the VLAN.

Table continues...

Name	Description
	<p> Note:</p> <p>This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see <i>Release Notes</i>.</p>
SpbMcast	Enables or disables Multicast over Fabric Connect. The default is disabled.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
IpssecEnable	Enables or disables IP security (IPsec) on the interface. The default is disabled.
Ipv6FhsSnoopDhcpEnable	Enables or disables IPv6 dhcp snooping on a VLAN. The default is disabled.
Ipv6FhsNDInspectionEnable	Enables or disables neighbor discovery (ND) inspection on a VLAN. The default is disabled.
DvrEnable	
DvrGwIpv4Addr	

Configuring BCBs for Inter-VSN routing

Inter-VSN allows you to route between IP networks on Layer 2 VLANs with different I-SIDs. Inter-VSN routing is typically used only when you have to extend a VLAN as a Layer 2 Virtual Services Network (VSN) for applications such as vMotion. Use IP Shortcuts 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.
- 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 485.

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 pane, expand the **Configuration > IP** folders.
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 pane, expand the **Configuration > VLAN** folders.
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 pane, expand the **Configuration > VLAN** folders.
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 pane, expand the **Configuration > IP** folders.
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	Description
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 maximum value varies per platform so refer to the <i>Release Notes</i> for platform-specific scaling information. The default value is 10000.
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: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-vlan • private

Table continues...

Name	Description
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/sub-port of each VLAN member. The sub-port only appears for channelized ports.
ActiveMembers	Specifies the slot/port/sub-port of each VLAN member. The sub-port only appears for channelized ports.
StaticMembers	Specifies the slot/port/sub-port of each static member of a policy-based VLAN. The sub-port only appears for channelized ports.
NotAllowToJoin	Specifies the slot/port/sub-ports that are never allowed to become a member of the policy-based VLAN. The sub-port only appears for channelized ports.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers of remote function call (RFC). If the VLAN type is port-based, none is displayed in the Basic tab ProtocolId field.

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: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-vlan
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.

Table continues...


Name	Description
	<p> Note:</p> <p>The switch reserves I-SID 0x00ffffff.</p>
ProtocolId	<p>Specifies the network protocol for protocol-based VLANs.</p> <p>If the VLAN type is not protocol-based, None is displayed in the Basic tab ProtocolId field.</p>
AgingTime	<p>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.</p>
MacAddress	<p>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 or ARP replies.</p>
Vlan Operation Action	<p>Performs an operation on the VLAN. The values are:</p> <ul style="list-style-type: none"> • 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 not perform the Snoop-related actions. • flushSnoopMemb: This action is not supported. • flushSnoopMRtr: This action is not supported. <p>The default is none.</p>
Result	<p>Specifies the result code after you perform an action. the default is none.</p>

Table continues...

Name	Description
NibMode	Enables or disables Microsoft Network Load Balancing (NLB) on the VLAN. * Note: This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see <i>Release Notes</i> .
SpbMcast	Enables or disables Multicast over Fabric Connect. The default is disabled.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
IpsecEnable	Enables or disables IP security (IPsec) on the interface. The default is disabled.
Ipv6FhsSnoopDhcpEnable	Enables or disables IPv6 dhcp snooping on a VLAN. The default is disabled.
Ipv6FhsNDInspectionEnable	Enables or disables neighbor discovery (ND) inspection on a VLAN. The default is disabled.
DvrEnable	
DvrGwIpv4Addr	

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.

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Inter-VSN routing with SPBM configuration example

The following figure shows a sample Inter-VSN deployment.




Figure 63: 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 223.

Note that the IP interfaces are configured where the routing instance exists, namely, on SwitchG.

SwitchC

VLAN CONFIGURATION

```
vlan create 11 type port-mstprstp 1
vlan members 11 1/2 portmember
vlan i-sid 11 12990011
```

SwitchG

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
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
ip address 10.100.12.1 255.255.255.0
exit
```

SwitchD

VLAN CONFIGURATION

```
vlan create 12 type port-mstprstp 1
```

```
vlan members 12 1/2 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.

SwitchG

```
SwitchG:1# show ip route vrf blue
```

```
=====
                        IP Route - VRF blue
=====
```

DST	MASK	NEXT	NH VRF	COST	INTER 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

```
SwitchG: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

SwitchG

```
SwitchG:1# show vlan mac-address-entry 11
```

```
=====
                        Vlan Fdb
=====
```

VLAN ID	STATUS	MAC ADDRESS	INTERFACE	TUNNEL
11	learned	00:00:00:00:01:02	Port-1/2	SwitchC
11	self	00:0e:62:25:a2:00	Port-cpp	-

2 out of 4 entries in all fdb(s) displayed.

```
SwitchG:1# show vlan mac-address-entry 12
```

```
=====
                        Vlan Fdb
=====
```

VLAN ID	STATUS	MAC ADDRESS	INTERFACE	TUNNEL
12	learned	00:00:00:00:02:02	Port-1/2	SwitchD
12	self	00:0e:62:25:a2:01	Port-cpp	-

2 out of 4 entries in all fdb(s) displayed.

SwitchC

SwitchC:1# show vlan mac-address-entry 11

```

=====
                                Vlan Fdb
=====
VLAN          MAC
ID   STATUS   ADDRESS                INTERFACE          TUNNEL
-----
11   learned  00:00:00:00:01:02     Port-1/2          SwitchD
11   learned  00:0e:62:25:a2:00     Port-1/2          SwitchD
2 out of 2 entries in all fdb(s) displayed.

```

SwitchD

SwitchD:1# show vlan mac-address-entry 12

```

=====
                                Vlan Fdb
=====
VLAN          MAC
ID   STATUS   ADDRESS                INTERFACE          TUNNEL
-----
12   learned  00:00:00:00:02:02     Port-1/2          SwitchC
12   learned  00:0e:62:25:a2:01     Port-1/2          SwitchC
2 out of 2 entries in all fdb(s) displayed.

```

Chapter 5: Operations and Management

CFM fundamentals

The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of ping and traceroute. To support troubleshooting of the SPBM cloud, the switch supports a subset of CFM functionality. Configure CFM on all SPBM VLANs.

CFM is based on the IEEE 802.1ag standard.

IEEE 802.1ag Connectivity Fault Management (CFM) provides OAM tools for the service layer, which allows you to monitor and troubleshoot an end-to-end Ethernet service instance. CFM is the standard for Layer 2 ping, Layer 2 traceroute, and the end-to-end connectivity check of the Ethernet network.

The 802.1ag feature divides or separates a network into administrative domains called Maintenance Domains (MD). Each MD is further subdivided 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)

CFM supports three kinds of standard CFM messages: Continuity Check Message (CCM), Loopback Message (LBM), and Linktrace Message (LTM). Messages are sent between Maintenance Points (MP) in the system.

On the switch, CFM is implemented using the LBM and LTM features only to debug SPBM. CCM messages are not required or supported.

You can assign maintenance levels for each CFM SPBM MEP and MIP to each SPBM B-VLAN individually or you can assign maintenance levels and global MEPs for all SPBM VLANs.

Autogenerated CFM and explicitly configured CFM

The switch simplifies CFM configuration with autogenerated CFM. With autogenerated CFM, you use the commands `cfm spbm enable` and `cfm cmac enable` and the switch creates default MD, MA, MEPs, and MIPs for SPBM B-VLANs and C-VLANs respectively.

If you choose to configure CFM explicitly, you must configure an MD, MA, MEPs, and MIPs.

- For SPBM B-VLANs, the switch provides two methods to configure CFM, namely, autogenerated and explicitly configured. You cannot use both.
- For C-VLANs, you can only use autogenerated CFM.

Important:

Only the VSP 4000 Series switch supports CFM configuration on C-VLANs.

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 be supported. However, if you want to enable the autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLANs. The switch only supports one type of MEP or MIP for each SPBM B-VLAN.

For information on autogenerated CFM configuration using the CLI see:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 509
- [Configuring autogenerated CFM on C-VLANs](#) on page 511

For information on autogenerated CFM configuration using the EDM see:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 536
- [Configuring autogenerated CFM on C-VLANs](#) on page 538

Explicitly configured CFM

If you choose to explicitly configure CFM, you must configure an MD, MA, MEPs, and MIPs. You can configure explicit CFM only on SPBM B-VLANs.

For explicit configuration information for CLI see [Configuring explicit mode CFM](#) on page 514.

For explicit configuration information for EDM see [Configuring explicit CFM in EDM](#) on page 539.

Using CFM

For SPBM B-VLANs, the autogenerated MEPs and MIPs respond to `12 ping`, `12 traceroute`, and `12 tracertree` in the same manner as the MEPs and MIPs created explicitly. For C-VLANs, the autogenerated MEPs and MIPs respond to `12 ping` and `12 traceroute`, but not to `12 tracertree` 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 `12 ping` and `12 traceroute` requests.

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 64: MD level assignment

Maintenance association endpoints (MEP)

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.

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.

l2ping

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.

This provides a simpler command syntax than the standard LBM commands, which require the user 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 in the customer domain. SPBM B-VLANs support the SMLT virtual option for the source mode.

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. The switch 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 given 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 above.
 - 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.

I2traceroute

The `12 traceroute` command is a proprietary command that allows you to trigger an LTM message. Use this command as follows:

- For B-VLANs, specify either the destination MAC address or node name.
- For C-VLANs, specify the destination MAC address.

This command provides a simpler command syntax than the standard LTM commands, which require the user to specify the MD, MA, and MEP ID information. The `12 traceroute` command provides a trace equivalent at Layer 2 for use with nodes on the SPBM B-VLAN in the customer domain.

! Important:

Only the VSP 4000 Series switch supports CFM configuration on C-VLANs.

* Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. After 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 `12 ping` and `12 traceroute` requests.

I2 traceroute with IP address

The `12 traceroute` command allows you to specify an IP address as the destination address. In this case, the IP address can be either a C-VLAN or a B-VLAN in the SPBM cloud.

The `12 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, `12 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 I2 traceroute and linktrace messages

For C-VLANs, CFM uses the following destination MAC addresses for the corresponding maintenance domain (MD) levels for `12 traceroute` and `linktrace` messages.

The switch supports both `12 traceroute` and `linktrace` for C-VLANs, but It is recommended that you use `12 traceroute`.

Table 15: 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

I2 tracetree

The `12 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.

Note:

The Avaya Virtual Services Platform 4000 Series does not support the `12 tracetree` command on C-VLANs because no multicast tree exists on C-VLANs.

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.

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 given B-VLAN. Each node (chassis) has a given 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 10 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 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.

The switch 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.

General CFM

- A single switch has a limit of one MEP and one MIP on a C-VLAN or B-VLAN.
- The maintenance level for MEPs and MIPs on a given B-VID (in a network) must be configured to the same level for them to respond to a given CFM command.
- You can configure global CFM at only one MD level for each switch for each VLAN type.
- All nodal MEPs and MIPs are restricted to SPBM B-VIDs.
- SMLT Virtual MAC for C-VLAN does not exist, so the switch does not support this option for the `12 ping` and `12 traceroute` commands.

Autogenerated CFM

- Autogenerated MEPs are not unique across the entire network unless you configure the global MEP ID on each switch to a different value. You must configure a unique MEP ID at a global level, for CFM.
- A single switch can have only one autogenerated MEP or MIP for each B-VLAN or C-VLAN.

Explicit CFM

- Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to be supported. However, if you want to enable autogenerated CFM you must first remove the existing MEP and MIP on the SPBM B-VLAN.
- You can assign maintenance levels for each CFM SPBM MEP and MIP to each SPBM B-VLAN individually or you can assign maintenance levels and global MEPs for all SPBM VLANs by following the appropriate procedure:
 - [Assigning a MEP MIP level to an SPBM B-VLAN](#) on page 517
 - [Assigning MEP MIP levels to SPBM B-VLANs globally](#) on page 519
 - [Configuring CFM nodal MEP](#) on page 542

C-VLAN versus SPBM B-VLAN considerations

Important:

Only VSP 4000 Series supports CFM configuration on C-VLANs.

CFM breaks the network into sections, called MEPs, so you can determine exactly where the problem exists.

The 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.

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 `12 ping` to learn the C-MAC address.
- For B-VLANs, you do not have to trigger an `12 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-VLAN 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 0x64. 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.

CFM configuration using CLI

This section provides procedures to configure and use Connectivity Fault Management (CFM) using Command Line Interface (CLI). The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. To support troubleshooting of the SPBM cloud, the switch supports a subset of CFM functionality.

*** Note:**

When you enable CFM in an SBPM network, it is recommended 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.

You can configure CFM using one of two modes: simplified or explicit. Both modes are described in the following sections, but the simplified mode is recommended.

*** Note:**

If you enable the `cfm spbm enable` command, you cannot assign a MEP/MIP level to an individual SPBM B-VLAN or configure CFM MD maintenance levels individually.

Regardless of whether you have chosen to configure individually or globally, there is one MEP per SPBM B-VLAN and one MIP level per SPBM B-VLAN.

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.

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

*** Note:**

Configuring CFM on a C-VLAN is not supported on all hardware platforms. If you do not see this command in the command list, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

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. However, if you want to enable the autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN.

The switch only supports one MEP and one MIP, either autogenerated or explicitly configured, on the SPBM B-VLAN. Similarly, the switch 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 information on configuring autogenerated CFM using the CLI, see:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 509
- [Configuring autogenerated CFM on C-VLANs](#) on page 511

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:

```
Switch>enable
Switch#configure terminal
Switch(config)#cfm spbm level 6
Switch(config)#cfm spbm mepid 4
```

```
Switch(config)#cfm spbm enable
Switch(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.

Note:

For C-VLANs, you can only use autogenerated CFM MEPs.

Configuring autogenerated CFM on a C-VLAN is not supported on all hardware platforms. If you do not see this command in the command list, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

! 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:

```
Switch>enable
Switch#configure terminal
Switch(config)#cfm cmac level 0
Switch(config)#cfm cmac mepid 4
Switch(config)#cfm cmac enable
Switch(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 mode CFM

In the explicit mode of configuring CFM, you can manually configure an MD, MA, MEP and then associate the MEP to a B-VLAN and assign a MIP level to a B-VLAN.

*** Note:**

If you use autogenerated CFM, these steps are unnecessary.

Configuring CFM MD

Use this procedure to configure the 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).

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create the CFM MD:

```
cfm maintenance-domain WORD<0-22> [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<0-22>
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# cfm maintenance-domain mdl index 99 maintenance-level 3
Switch:1(config)# show cfm maintenance-domain
```

Maintenance Domain			
Domain Name	Domain Index	Level	Domain Type
mdl	99	3	NODAL

```
Total number of Maintenance Domain entries: 1.
Switch:1(config)# no cfm maintenance-domain mdl
Switch: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
<code>WORD<0-22></code>	Specifies the maintenance domain name.
<code>index <1-2147483647></code>	Specifies a maintenance domain entry index.
<code>maintenance-level <0-7></code>	Specifies the MD maintenance level when creating the MD. The default is 4.
<code>level <0-7></code>	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). 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.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create the CFM MA:

```
cfm maintenance-association WORD<0-22> WORD<0-22> [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<0-22> WORD<0-22>
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# cfm maintenance-association md1 ma1 index 98
```

```
Switch:1(config)# show cfm maintenance-association
```

```
=====
Maintenance Association Status
=====
Domain Name          Assn Name          Domain Idx  Assn Idx
-----
mdl                  mal                1           98

Total number of Maintenance Association entries: 1.
=====
Maintenance Association config
=====
Domain Name          Assn Name
-----
mdl                  mal

Total number of MA entries: 1.
```

Variable definitions

Use the data in the following table to use the `cfm maintenance-association` command.

Variable	Value
<code>WORD<0-22> WORD<0-22></code>	Creates the CFM MA. The first parameter, specifies the MD name. The second parameter, specifies the MA short name.
<code>index <1-2147483647></code>	Specifies a maintenance association entry index.

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.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Create the CFM MEP:


```
cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
[state <enable>]
```
3. Enable an existing CFM MEP:


```
cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
```
4. Disable an existing CFM MEP:


```
no cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
```
5. Display the CFM MEP configuration:

```
show cfm maintenance-endpoint
```

6. Delete an existing CFM MEP:

```
no cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191>
```

Example

```
Switch:1> enable
```

```
Switch:1# configure terminal
```

```
Switch:1(config)# cfm maintenance-endpoint mdl ma1 1 state enable
```

```
Switch:1(config)# show cfm maintenance-endpoint
```

```
=====
Maintenance Endpoint Config
=====
DOMAIN          ASSOCIATION      MEP  ADMIN
NAME            NAME             ID
-----
mdl             ma1              1    enable

Total number of MEP entries: 1.

=====
Maintenance Endpoint Service
=====
DOMAIN_NAME     ASSN_NAME        MEP_ID TYPE    SERVICE_DESCRIPTION
-----
mdl             ma1              1     nodal   Vlan 1, Level 4

Total number of MEP entries: 1.
```

Variable definitions

Use the data in the following table to use the `cfm maintenance-endpoint` command.

Variable	Value
<code>WORD<0-22></code>	The first parameter, specifies the MD name.
<code>WORD<0-22></code>	The second parameter, specifies the MA short name.
<code><1-8191></code>	Specifies the MEP ID.
<code>enable</code>	Enables an existing MEP. Use this parameter with the <code>no</code> option to disable an existing MEP.
<code>state {enable disable}</code>	Enables or disables the MEP when creating the MEP. The default is disabled.

Assigning a MEP/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 given B-VLAN. 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.

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 given

MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

1. Add nodal MEPs to the B-VLAN:

```
vlan nodal-mep <1-4059> WORD<0-22> WORD<0-22> <1-8191>
```

2. Display the nodal MEP configuration:

```
show vlan nodal-mep <1-4059>
```

3. Remove the nodal MEPs from the B-VLAN:

```
no vlan nodal-mep <1-4059> WORD<0-22> WORD<0-22> <1-8191>
```

4. Add nodal MIP level to the B-VLAN:

```
vlan nodal-mip-level <1-4059> WORD<0-15>
```

5. Display the nodal MIP level configuration:

```
show vlan nodal-mip-level [<1-4059>]
```

6. Remove the nodal MIP level from the B-VLAN:

```
no vlan nodal-mip-level <1-4059> WORD<0-15>
```

Example

```
Switch:1> enable
```

```
Switch:1# configure terminal
```

```
Switch:1(config)# vlan nodal-mep 100 md1 ma1 2
```

```
Switch:1(config)# show vlan nodal-mep
```

```
=====
                                Vlan Nodal Mep
=====
VLAN_ID    DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
-----
100        spbm.100.6
200        spbm.200.6
=====
```

```
Switch:1(config)# vlan nodal-mip 100 6
```

```
Switch:1(config)# show vlan nodal-mip
```

```
=====
                                Vlan Nodal Mip Level
=====
VLAN_ID    NODAL_MIP_LEVEL_LIST
-----
1
100        6
=====
```

```
216
304
41000
1001
```

Variable definitions

Use the data in the following table to use the `vlan nodal-mep` command.

Variable	Value
<1-4059>	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
<1-4059>	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.

Assigning MEP/MIP levels to SPBM B-VLANs globally

* Note:

If you enable the `cfm spbm enable` command, you cannot assign a MEP/MIP level to an individual SPBM B-VLAN or configure CFM MD maintenance levels individually.

About this task

Enables the global CFM MEP and MIPs for all SPBM B-VLANs.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Enable simplified CFM configuration for SPBM VLANs:


```
cfm spbm enable
```
3. Enter the CFM SPBM level:


```
cfm spbm level <0-7>
```
4. Enter the CFM SPBM MEPID level:


```
cfm spbm mepid <1-8191>
```

Example

```
Switch:1(config)# cfm spbm level 7
Switch:1(config)# cfm spbm mepid 12
Switch:1(config)# cfm spbm enable
```

Variable definitions

Use the data in the following table to use the simplified CFM commands.

Variable	Value
spbm level <0-7>	Configures the maintenance level for every CFM SPBM MEP and MIP level on all SPBM B-VLANs. The default is 4.
mepid <1-8191>	Assigns a global MEP ID for all CFM SPBM MEPs. The default is 1.
no cfm spbm enable	Disables global configuration of CFM SPBM MEP and MIP levels on all SPBM B-VLANs.
default cfm spbm level	Returns maintenance level to default for all CFM SPBM MEP and MIP level on all SPBM B-VLANs.
default cfm spbm mepid	Returns MEP ID for all CFM SPBM MEPs to default.
show cfm spbm	Displays the global CFM MEP configuration for SPBM B-VLANs.

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.

Procedure

Trigger the loopback test:

```
loopback WORD<0-22> WORD<0-22> <1-8191> <0x00:0x00:0x00:0x00:0x00:0x00>
[burst-count <1-200>] [data-tlv-size <0-400>] [frame-size <64-1500>]
[interframe-interval <msecs>] [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

```
Switch:1# loopback md1 4001 13 00:14:0D:A2:B3:DF burst-count 10 priority 3
time-out 5
```

```
Result of LBM from mep: spbm.bvlan1000.8 to MAC address: 00:66:00:66:00:66 :
Sequence number of the first LBM is 150404162
The total number of LBMs sent out is 1
The number of LBRs received is 1
```

```

The number of LBRs lost is 0
The percentage of LBMs lost is 0.00%
The RTT Min is 15071 microseconds, Max is 15071 microseconds, Average is 15071.00 microseconds
The Standard Deviation of RTT is 0.00 microseconds

```

Variable definitions

Use the data in the following table to use the `loopback` command.

Variable	Value
<code>WORD<0–22></code>	The first parameter, specifies the MD name.
<code>WORD<0–22></code>	The second parameter, specifies the MA name.
<code><1–8191></code>	Specifies the MEP ID.
<code><0x00:0x00:0x00:0x00:0x00:0x00></code>	Specifies the remote MAC address to reach the MEP/MIP.
<code>burst-count <1–200></code>	Specifies the burst-count.
<code>data-tlv-size <0–400></code>	Specifies the data TLV size.
<code>frame-size <64–1500></code>	Specifies the frame-size. The default is 0.
<code>priority <0–7></code>	Specifies the priority. The default is 7.
<code>source-mode{nodal noVlanMac smltVirtual}</code>	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> • 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 B-MAC-SA. • smltVirtual—Use this value with B-VLANs only. <p>The default is nodal.</p>
<code>testfill-pattern {all-zero all-zero-crc pseudo-random-bit-sequence pseudo-random-bit-sequence-crc}</code>	<p>Specifies the testfill pattern:</p> <ul style="list-style-type: none"> • 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. <p>A cyclic redundancy check is a code that detects errors.</p> <p>The default is 1: all-zero.</p>
<code>time-out <1–10></code>	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 virtual 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 VLAN.

Procedure

Trigger the linktrace:

```
linktrace WORD<0-22> WORD<0-22> <1-8191> <0x00:0x00:0x00:0x00:0x00:0x00>
[detail] [priority <0-7>] [source-mode <nodal|noVlanMac|smltVirtual>]
[ttl-value <1-255>]
```

Example

```
Switch: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<0-22>	The first parameter, specifies the MD name.
WORD<0-22>	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: <ul style="list-style-type: none"> • 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.

Table continues...

Variable	Value
	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 to trigger a Layer 2 ping, inside an SPBM cloud or network, which acts like native **ping**. This feature enables CFM to debug Layer 2. It can also help you debug IP shortcuts and the record for the shortcuts' ARP.

Before you begin

- You must have a MEP that is associated with a VLAN.

Procedure

Trigger a Layer 2 ping:

```
l2 ping {vlan <1-4059> routernodename WORD<0-255> | vlan <1-4059> 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

```
Switch:1# l2 ping vlan 2 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(68) bytes of data
1 packets transmitted, 0 packets received, 100.00% packet loss
```

```
Switch:1# l2 ping vlan 2 routernodename MONTIO
```

```
Please wait for l2ping to complete or press any key to abort
----00:14:0d:a2:b3:df L2 PING Statistics---- 0(68) 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
```

```
Switch: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
(us)                                PKTS  PKTS  LOSS    MIN/MAX/AVE
```

```
=====
2 SHAMIM (00:1a:8f:08:53:df) 1 0 100.00% 0/0/0.00
=====
```

Variable definitions

Use the data in the following table to configure the L2 ping parameters.

Variable	Value
{vlan <1-4059> routernodename WORD<0-255> } {vlan <1-4059> mac <0x00:0x00:0x00:0x00:0x00:0x00> } {ip-address WORD<0-255> }	Specifies the destination for the L2 ping: <ul style="list-style-type: none"> • <1-4059> — 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.
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: <ul style="list-style-type: none"> • 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. <p>A cyclic redundancy check is a code that detects errors.</p> <p>The default is all-zero.</p>
priority <0-7>	Specifies the priority. The default is 7.
time-out <1-10>	Specifies the interval in seconds. The default is 3.
source-mode<nodal noVlanMac smltVirtual>	Specifies the source mode: <ul style="list-style-type: none"> • 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. <p>The default is 1: nodal.</p>
vrf WORD<1-16>	Specifies the VRF name.

Triggering a Layer 2 traceroute

Use this procedure to trigger a Layer 2 traceroute, which acts like native `traceroute`. This feature enables CFM to debug Layer 2 in an SPBM cloud or network. It can determine the path used by IS —IS to get from one MEP to another, by showing all the hops between. Therefore, it can show where connectivity is lost. It can also work for IP shortcuts.

! 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.

`linktrace` traces the path up to the closest device to that MAC address that supports CFM.

Before you begin

- You must have a MEP that is associated with a VLAN.

Procedure

Trigger a Layer 2 traceroute:

```
l2 traceroute {<vlan <1-4059> routernodename WORD<0-255> | <vlan <1-4059>
mac <0x00:0x00:0x00:0x00:0x00:0x00>} [priority <0-7>] [source-mode
<nodal|noVlanMac|smltVirtual>] [ttl <1-255>]
```

```
l2 traceroute {ip-address WORD<0-255>} [priority <0-7>][source-mode
<nodal|noVlanMac|smltVirtual>][ttl <1-255>] [vrf WORD<1-16>]
```

Example

```
Switch:1# l2 traceroute vlan 2 routernodename Switch-MONTIO
```

```
Please wait for l2traceroute to complete or press any key to abort
```

```
l2traceroute to Switch-MONTIO (00:14:0d:a2:b3:df), vlan 2
0 Switch-PETER4 (00:15:9b:11:33:df)
1 Switch-MONTIO (00:14:0d:a2:b3:df)
```

```
Switch: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
```

```
=====
Switch-SHAMIM (00:1a:8f:08:53:df), vlan 2
0 Switch-PETER4 (00:15:9b:11:33:df)
1 Switch-MONTIO (00:14:0d:a2:b3:df)
```

Variable definitions

Use the data in the following table to use the `l2 traceroute` command.

Variable	Value
{vlan <1-4059> routernodename WORD<0-255> } {vlan <1-4059> mac <0x00:0x00:0x00:0x00:0x00:0x00> } {ip-address WORD<0-255> }	Specifies the destination for the L2 traceroute: <ul style="list-style-type: none"> • <1-4059> — 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
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: <ul style="list-style-type: none"> • 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 B-MAC-SA. • 2: smltVirtual—Use this value with B-VLANs only. The default is 1: 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:

This command is supported 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

Trigger a Layer 2 tracetree:

```
12 tracetree {<1-4059> <1-16777215> [routernodename WORD<0-255> |
<1-4059> <1-16777215>] [mac <0x00:0x00:0x00:0x00:0x00:0x00>]} [priority
<0-7>] [source-mode <nodal|noVlanMac|smltVirtual>] [ttl-value <1-255>]
```

Example

```
Switch:1# l2 tracetree 500 1
Switch: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  Switch-PETER4          00:15:9b:11:33:df -> Switch-MONTI0          00:14:0d:a2:b3:df
2  Switch-MONTI0         00:14:0d:a2:b3:df -> Switch-LEE2           00:15:e8:b8:a3:df
```

Variable definitions

Use the data in the following table to use the `l2 tracetree` command.

Variable	Value
{ <1-4059><1-16777215> routernodename WORD<0-255> <1-4059><1-16777215> mac <0x00:0x00:0x00:0x00:0x00:0x00>}	<ul style="list-style-type: none"> • <1-4059> — 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.
tll-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 noVlanMac smltVirtual>	Specifies the source mode: <ul style="list-style-type: none"> • 1: nodal • 2: smltVirtual The default is nodal.

Triggering a Layer 2 tracemroute

Use this procedure to debug the IP Multicast over Fabric Connect 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
<1-4059>[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 PETER4 00:03:00:00:00:00 -> LEE1 00:14:0d:bf:a3:df
2 LEE1 00:14:0d:bf:a3:df -> 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 12 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 PETER4 00:03:00:00:00:00 -> LEE1 00:14:0d:bf:a3:df
2 LEE1 00:14:0d:bf:a3:df -> LEE2 00:15:e8:b8:a3:df

```

Variable definitions

Use the data in the following table to use the `12 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 `12 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 `12 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 CLI, instrumentation, 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), it is recommended that 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

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# clear trace
Switch:1(config)# trace cfm level 3
Switch:1(config)# trace shutdown
Switch:1(config)# show trace cfm
=====
                          CFM Tracing Info
=====
Status      : Enabled
Level       : VERBOSE
Switch:1(config)#trace level 120 3
Switch:1(config)# save trace
Switch:1(config)# trace grep error
Switch:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

Use the data in the following table to use the `trace` command.

Variable	Value
cfm level [<i><0-4></i>]	Starts the trace by specifying the level. <ul style="list-style-type: none"> • <i><0-4></i> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
filter	Configures a filter trace for a file or module.
flags	Configures trace flags for IS-IS or OSPF.
grep [<i>WORD<0-128></i>]	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level <i><0-215></i> [<i><0-4></i>]	Starts the trace by specifying the module ID and level. <ul style="list-style-type: none"> • <i><0-215></i> specifies the module ID. • <i><0-4></i> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
route-map	Enables or disables the trace route-map. The values are on and off.
screen {disable enable}	Enables the display of trace output to the screen.
shutdown	Stops the trace operation.
spbm isis level [<i><0-4></i>]	Starts the trace by specifying the level. <ul style="list-style-type: none"> • <i><0-4></i> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose. The default is 1, very terse.

Use the data in the following table to use the **save trace** command.

Variable	Value
file <i>WORD<1-99></i>	Specifies the file name in one of the following formats: <ul style="list-style-type: none"> • a.b.c.d:<file> • 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 CPU. /mnt/extflash is the external flash of the CPU.

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 CLI, instrumentation, 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 CLI, instrumentation, 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), it is recommended that 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

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# clear trace
Switch:1(config)# trace spbm isis level 3
Switch:1(config)# trace shutdown
Switch:1(config)# show trace spbm isis
=====
                        SPBM ISIS Tracing Info
=====
Status      : Enabled
Level       : VERY_TERSE
Flag Info   :
Switch:1(config)#trace level 125 3
Switch:1(config)#trace level 119 3
Switch:1(config)# save trace
Switch:1(config)# trace grep error
Switch:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

Use the data in the following table to use the `trace` command.

Variable	Value
cfm level [<code><0-4></code>]	Starts the trace by specifying the level. <ul style="list-style-type: none"> <code><0-4></code> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
filter	Configure a filter trace for a file or module.

Table continues...

Variable	Value
flags	Configure trace flags for IS-IS or OSPF.
grep [WORD<0-128>]	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level <0-215>[<0-4>]	Starts the trace by specifying the module ID and level. <ul style="list-style-type: none"> • <0-215> 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 verbose, 4 is very verbose.
route-map	Enables or disables the trace route-map. The values are on and off.
screen {disable enable}	Enables the display of trace output to the screen.
shutdown	Stops the trace operation.
spbm isis level [<0-4>]	Starts the trace by specifying the level. <ul style="list-style-type: none"> • <0-4> specifies the trace level from 0-4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose. <p>The default is 1, very terse.</p>

Use the data in the following table to use the **save trace** command.

Variable	Value
file WORD<1-99>	Specifies the file name in one of the following formats: <ul style="list-style-type: none"> • a.b.c.d:<file> • x:x:x:x:x:x:x <file> • /intflash/<file> • /extflash/<file> • /usb/<file> • /mnt/intflash/ <file> • /mnt/extflash/ <file> <p>/mnt/intflash is the internal flash of the CPU. /mnt/extflash is the external flash of the CPU.</p>

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, it is recommended 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.

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

*** Note:**

Configuring CFM on a C-VLAN is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function. However, if you want to enable the autogenerated commands, you must first remove the existing MEP and MIP on the SPBM B-VLAN. The switch only supports one MEP or MIP on the SPBM B-VLAN, either explicitly configured or autogenerated.

For autogenerated CFM configuration information for EDM, see the following tasks:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 536
- [Configuring autogenerated CFM on C-VLANs](#) on page 538

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 configuration 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 538.

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. In the navigation pane, expand the **Configuration > Edit > Diagnostics** folders.
2. Click **CFM**.
3. Click the **Global** tab.
4. Select **enable** next to **SpbmAdminState**.
5. Click **Apply**.
6. To verify the values assigned to MA, MD, and MEP, perform the following steps:
 - a. Click the **MD** tab.
 - b. Select **SPBM**, and then check the MA and MEP values.

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. 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. This field does not appear for all hardware platforms.
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. This field does not appear for all hardware platforms.

Table continues...

Name	Description
CmacMepId	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. This field does not appear for all hardware platforms.
Bmac	Displays the B-MAC address of the node. This field does not appear for all hardware platforms.
Cmac	Displays the C-MAC address of the node. This field does not appear for all hardware platforms.

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 536.

* Note:

For C-VLANs, you can only use autogenerated CFM MEPs.

Configuring autogenerated CFM on a C-VLAN is not supported on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

! Important:

CFM supports one MEP or MIP on 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 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. In the navigation pane, expand the **Configuration > Edit > Diagnostics** folders.
2. Click **CFM**.
3. Click the **Global** tab.

4. Select **enable** next to **CmacAdminState**.
5. In the fields provided, specify a maintenance level and a MEP ID.
6. 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. 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. This field does not appear for all hardware platforms.
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. This field does not appear for all hardware platforms.
CmacMepId	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. This field does not appear for all hardware platforms.
Bmac	Displays the B-MAC address of the node. This field does not appear for all hardware platforms.
Cmac	Displays the C-MAC address of the node. This field does not appear for all hardware platforms.

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 536. 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.

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).

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics** folders.
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.

Before you begin

- You must configure a CFM MD.

Procedure

- In the navigation pane, expand the **Configuration > Edit > Diagnostics** folders.
- Click **CFM**.
- Click the **MD** tab.
- Highlight an existing MD, and then click **MaintenanceAssociation**.
- In the **MA** tab, click **Insert**.
- In the fields provided, specify an index value and name for the MA.
- 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.

Procedure

- In the navigation pane, expand the **Configuration > Edit > Diagnostics** folders.
- Click **CFM**.
- Click the **MD** tab.
- Highlight an existing MD, and then click **MaintenanceAssociation**.
- In the **MA** tab, highlight an existing MA, and then click **MaintenanceEndpoint**.
- Click **Insert**.
- In the fields provided, specify the ID and the administrative state of the MEP.
- 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: <ul style="list-style-type: none"> • trunk • sg • endpt • vlan • port • endptClient • nodal • remotetrunk • remotesg • remoteendpt • remoteVlan • remotePort • remoteEndptClient
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 given B-VLAN. 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.

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 given MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **VLANs**.
3. Click the **Advanced** tab.
4. Select a VLAN with a type of spbm-bvlan.
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 on a per level per VLAN basis.

Configuring Layer 2 ping

Use this procedure to configure a Layer 2 ping inside an SPBM cloud or network. This feature enables CFM to debug Layer 2. It can also help you debug IP shortcuts and the record for the shortcuts' ARP.

* Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

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. In the navigation tree, expand the **Configuration > Edit > Diagnostics** folders.
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 backbone 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	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Ping messages. • abort: the service aborted or is about to abort the L2Ping messages. <p>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.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the L2Ping Messages will be (or have been) sent. • false: the L2Ping Messages will not be sent. <p>The default is true.</p>
Priority	<p>Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame.</p> <p>The default is 7.</p>
TimeoutInt	<p>Specifies the interval to wait for an L2Ping time-out. The default value is 3 seconds.</p>
TestPattern	<p>Specifies the test pattern to use in the L2Ping PDU:</p> <ul style="list-style-type: none"> • allZero: null signal without cyclic redundancy check

Table continues...

Name	Description
	<ul style="list-style-type: none"> 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. <p>A cyclic redundancy check is a code that detects errors. The default value is allZero.</p>
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	<p>Specifies the source mode of the transmit loopback service:</p> <ul style="list-style-type: none"> 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. <p>The default is nodal.</p>
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 to trigger a Layer 2 traceroute. This feature enables CFM to debug Layer 2 in an SPBM cloud or network. It can determine the path used by IS—IS to get from one MEP to another, by showing all the hops between. Therefore, it can show where connectivity is lost. It can also work for IP shortcuts.

* Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

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 configuring tracetable, see [Configuring Layer 2 tracetable](#) on page 562.

! 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.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

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. 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 **L2 Traceroute/TraceTree** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the Backbone VLAN (B-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: <ul style="list-style-type: none"> • If you configure IsTraceTree to false then EDM performs Traceroute on the unicast path.

Table continues...

Name	Description
	<ul style="list-style-type: none"> If you configure IsTraceTree to true then EDM performs TraceTree on the multicast tree.
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> 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. <p>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.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> true: the L2Trace messages will be (or have been) sent. false: the L2Trace messages will not be sent. <p>The default is true.</p>
Ttl	<p>Specifies the number of hops remaining to this L2Trace.</p> <p>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.</p> <p>The default value is 64.</p>
SourceMode	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> 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. <p>The default is 1: nodal.</p>
SeqNumber	<p>Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.</p>

Table continues...

Name	Description
Flag	<p>L2Trace result flag that indicates L2Trace status or error code:</p> <ul style="list-style-type: none"> • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled in order 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 • 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.

Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

About this task

You can display Layer 2 tracetrace 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 tracetrace results](#) on page 565.

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/Tracetree Result field descriptions

Use the data in the following table to use the **L2 Traceroute/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 L2Trace's response of the L2Trace is going to be returned. The default is 0.
Hop	The number of hops away from L2Trace initiator.
ReceiveOrder	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	Time-to-Live (TTL) field value for a returned L2Trace response.
SrcMac	MAC address of the MP that responds to the L2Trace request for this L2TraceReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	The host name of the node that forwarded the L2Trace to the responding node.

Configuring Layer 2 IP ping

Use this procedure to configure Layer 2 IP ping

*** Note:**

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

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.
- 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 IPv4 Routing*.

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 then 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 per MAC/VLAN pair. Range is 1–200. The default is 1.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • 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. <p>This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: L2IpPing Messages will be or have been sent. • false: L2IpPing Messages will not be sent. <p>The default is true.</p>

Table continues...

Name	Description
TimeoutInt	Specifies the interval to wait for an L2IPing time-out with a range of 1–10 seconds with a default value of 3 seconds.
TestPattern	<p>Specifies the test pattern to use in the L2IPing PDU:</p> <ul style="list-style-type: none"> • 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. <p>A cyclic redundancy check is a code that detects errors.</p> <p>The default value is allZero.</p>
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:**

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

*** 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	The address type of the destination IP Address.
IpAddr	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.
Vrflid	Specifies the VRF ID.
VlanId	Specifies the VLAN ID found from the Layer 3 lookup and used for transmission.
DestMacAddress	An indication of the target MAC Address transmitted.
PortNum	Either the value '0', or the port number of the port used for the L2 IP ping.
DestHostName	The host name of the responding node.
Size	The number of bytes of data sent.
PktsTx	Number of Packets transmitted for this VLAN/MAC.
PktsRx	Number of Packets received for this VLAN/MAC.
PercentLossWhole	Percentage of packet loss for this VLAN/MAC.
PercentLossFract	Percentage of packet loss for this VLAN/MAC.
MinRoundTrip	Minimum time for round-trip for this VLAN/MAC in us.
MaxRoundTrip	Maximum time for round-trip for this VLAN/MAC in us.
RttAvgWhole	Average time for round-trip for this VLAN/MAC in us.
RttAvgFract	Fractional portion of average time for round-trip.
Flag	Result flag indicating status or error code: <ul style="list-style-type: none"> • 1 - No error • 2 - Internal error • 3 - Invalid IP • 4 - L2Trace completed • 5 - Lookup failure for IP (no VLAN/MAC entries)

Configuring Layer 2 IP traceroute

Use this procedure to configure Layer 2 IP traceroute.

*** Note:**

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

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
- 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 IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics** folders.
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: <ul style="list-style-type: none"> • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Trace messages.

Table continues...

Name	Description
	<ul style="list-style-type: none"> • abort: the service is aborted or about to abort the L2Trace messages. <p>This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time. The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the Trace Messages will be or have been sent. • false. the Trace Messages will not be sent <p>The default is true.</p>
PathsFound	<p>Specifies the number of paths found to execute the L2trace. The default is 0.</p>

Viewing Layer 2 IP traceroute results

Use this procedure to view Layer 2 IP traceroute results.

*** Note:**

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

*** 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.

Table continues...

Name	Description
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 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	L2Trace result flag indicating status or error code: <ul style="list-style-type: none"> • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed

Table continues...

Name	Description
	<ul style="list-style-type: none"> I2TraceLookupFailure (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

- 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 pane, expand the **Configuration > Edit > Diagnostics** folders.
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.

Table continues...

Name	Description
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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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) msec and a default value of 500 msec. The value of 0 msec indicates to send the frames as fast as possible. The default is 500.
TestPattern	Specifies the testfill pattern: <ul style="list-style-type: none"> • 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.

Table continues...

Name	Description
FrameSize	Specifies the frame-size. The default is 0.
Sourcemode	<p>Specifies the source mode of the transmit loopback service:</p> <ul style="list-style-type: none"> • 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 B-MAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. <p>The default is nodal.</p>

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

- 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 pane, expand the **Configuration > Edit > Diagnostics** folders.
 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.
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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 1 - No error

Table continues...

Name	Description
	<ul style="list-style-type: none"> • 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 in order to perform LTM • 7 - No LTR response received • 8 - Linktrace to own MEP MAC is not sent • 9 - Endpoint must be enabled in order to perform LTM • 10 - Pbt-trunk must be enabled in order to perform LTM • 11 - LTM completed • 12 - LTM leaf node
SourceMode	<p>Specifies the source mode of the transmit loopback service:</p> <ul style="list-style-type: none"> • 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. <p>The default is nodal.</p>

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 pane, expand the **Configuration > Edit > Diagnostics** folders.
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.
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.

Table continues...

Name	Description
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.
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.

* Note:

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

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:

This command is supported 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.
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	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the L2Tracetree Messages will be (or have been) sent. • false: the L2Tracetree Messages will not be sent <p>The default is true.</p>

Table continues...

Name	Description
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 mode of the transmit loopback service: <ul style="list-style-type: none"> • 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.
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: <ul style="list-style-type: none"> • 1 — No error • 2 — L2Tracetree internal error • 3 — Invalid MAC address • 4 — MEP must be enabled in order 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 nonprimary 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.

*** Note:**

Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. For more information, see *Release Notes*. As an alternative, use CLI.

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 Traceroute/Tracetree Result field descriptions

Use the data in the following table to use the **L2 Traceroute/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, that indicates 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.

Table continues...

Name	Description
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:

- Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) are not supported in EDM. As an alternative, use the command line interface.
- If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context.

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. In the navigation pane, expand the **Configuration > Edit > Diagnostics > L2Ping/ L2Trace Route** folders.
2. Click the **L2MCAST Traceroute** tab.
3. Click **Insert** to insert the L2 MCAST Traceroute.
4. Enter the **SrcIpAddr**.
5. Enter 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 568.

7. In the **ServiceId** field, 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 **Start**.
13. Update the Layer 2 tracemroute by clicking **Refresh** .
14. To stop the Layer 2 tracemroute, click **Stop** .
15. To see the result, click **Result**.

L2 MCAST Traceroute field descriptions

Use the data in the following table to use the **L2MCAST Traceroute** tab.

Name	Description
SrclpAddrType	Specifies the source IP address type as IPv4.
SrclpAddr	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.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • 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. <p>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.</p>
ResultOK	<p>Specifies the result of the operation:</p> <ul style="list-style-type: none"> • true: The trace messages will be or have been sent. • false: The trace messages will not be sent.

Table continues...

Name	Description
Flag	<p>Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case.</p> <ul style="list-style-type: none"> • 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:

- Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. As an alternative, use the CLI.
- 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 566.

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. In the navigation pane, expand the **Configuration > VRF Context View > Set VRF Context View** folders.
2. Select a VRF and click the **Launch VRF Context View** tab.
3. In the navigation pane, 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. Enter the **SrclpAddr**.
7. Enter 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 **Start**.
15. Update the Layer 2 tracemroute by clicking **Refresh** .
16. To stop the Layer 2 tracemroute, click **Stop** .
17. To see the result, click **Result**.

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.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • 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. <p>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.</p>

Table continues...

Name	Description
ResultOK	<p>Specifies the result of the operation:</p> <ul style="list-style-type: none"> • true: The trace messages will be or have been sent. • false: The trace messages will not be sent.
Flag	<p>Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case.</p> <ul style="list-style-type: none"> • 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.

*** Note:**

- Troubleshooting using ping and traceroute (including Layer 2 ping and Layer 2 traceroute) is not supported on EDM. As an alternative, use the CLI.
- If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Diagnostics > L2Ping/L2Trace Route** folders.
2. Click the **L2 MCAST Traceroute** tab.
3. To view the CFMI2 trace multicast route results, highlight an entry and click **Result**.

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 I2 tracemroute response is going to be returned.

Table continues...

Name	Description
Hop	Specifies the number of hops away from the I2 tracemroute initiator.
ReceiveOrder	Specifies an index to distinguish among multiple I2 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 I2 tracemroute response.
SrcMac	Specifies the MAC address of the MP that responds to the I2 tracemroute request for this I2 tracemrouteReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the I2 tracemroute to the responding node.
LastHostName	Specifies the host name of the node that forwarded the I2 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 configure CFM.

Switch 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 CONFIGURATION

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

```

Switch 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 CONFIGURATION

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

```

Switch:1# show isis adjacencies
=====
                        ISIS Adjacencies
=====
INTERFACE          IP ADDR          L STATE          UPTIME          PRI          HOLDTIME          SYSID
-----
Port1/3             44.17.10.33      1 UP             00:37:37        127          19
0014.0dbf.a3df
Port1/19            44.17.10.36      1 UP             1d 05:09:16     127          21
0014.0da2.b3df
-----
2 out of 2 interfaces have formed an adjacency
=====
    
```

I2 ping with vlan

```

Switch: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(68) bytes of data
1 packets transmitted, 0 packets received,  100.00% packet loss
    
```

I2 ping with vlan

```

Switch:1# l2 ping vlan 500 routernodename MONTI0

Please wait for l2ping to complete or press any key to abort
    
```

```

----00:14:0d:a2:b3:df    L2 PING Statistics---- 0(68) 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

```

I2 traceroute with vlan

```
Switch:1# l2 traceroute vlan 500 routernodename MONTI0
```

Please wait for l2traceroute to complete or press any key to abort

```

l2traceroute to MONTI0 (00:14:0d:a2:b3:df),  vlan 500
0   PETER4          (00:15:9b:11:33:df)
1   MONTI0          (00:14:0d:a2:b3:df)

```

I2 tracetree with vlan

```
Switch: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   PETER4          00:15:9b:11:33:df -> MONTI0          00:14:0d:a2:b3:df
2   MONTI0          00:14:0d:a2:b3:df -> 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.

I2 ping with IP address

```
Switch: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  SHAMIM          (00:1a:8f:08:53:df)  1     0     100.00%  0/0/0.00

```

I2 ping with IPv6 address

```
Switch:1# l2 ping ip-address 49:0:0:0:0:0:0:11
```

Please wait for l2ping to complete or press any key to abort

```

L2 PING  Statistics : IP 49:0:0:0:0:0:0:11, paths found 1, paths attempted 1
=====
          TX      RX      PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                                PKTS  PKTS  LOSS      MIN/MAX/AVE (us)
=====
41   SHAMIM          (00:49:00:01:00:11)  1     1     0.00%  11876/11876/11876.00

```

I2 traceroute with IP address

```
Switch: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
=====
SHAMIM (00:1a:8f:08:53:df), vlan 500
0 PETER4 (00:15:9b:11:33:df)
1 MONTIO (00:14:0d:a2:b3:df)
```

I2 traceroute with IPv6 address

```
Switch:1# l2 traceroute ip-address 49:0:0:0:0:0:11
Please wait for l2trace to complete or press any key to abort
L2 Trace Statistics : IP 49:0:0:0:0:0:11, paths found 1
=====
SHAMIM (00:49:00:01:00:11), vlan 41
0 4K-DUT7 (00:49:00:01:00:17)
1 9k-2 (00:49:00:01:00:92)
2 8K-1 (00:49:00:08:00:81)
3 4K-DUT1 (00:49:00:01:00:11)
```

show cfm maintenance-domain

```
Switch:1#show cfm maintenance-domain
=====
Maintenance Domain
=====
Domain Name          Domain Index   Level Domain Type
-----
mdl                   99             3      NONE
Total number of Maintenance Domain entries: 1.
```

show cfm maintenance-association

```
Switch:1#show cfm maintenance-association
=====
Maintenance Association Status
=====
Domain Name          Assn Name      Domain Idx  Assn Idx
-----
mdl                   mal             1           98
Total number of Maintenance Association entries: 1.
=====
Maintenance Association config
=====
Domain Name          Assn Name
-----
mdl                   mal
Total number of MA entries: 1.
```

show cfm maintenance-endpoint

```
Switch:1#show cfm maintenance-endpoint
=====
Maintenance Endpoint Config
=====
DOMAIN              ASSOCIATION    MEP  ADMIN
NAME                NAME           ID
-----
mdl                   mal             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

```
Switch: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

```
Switch: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
```

Chapter 6: Resources

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.

Documentation

See *Documentation Reference* for a list of documentation for all VOSS products.

For installation and initial setup information of the Open Networking Adapter (ONA), refer to the Quick Install Guide that came with your ONA.

 **Note:**

The ONA works only with the Avaya Virtual Services Platform 4000 Series.

Training

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

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.

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.

Subscribing to e-notifications

Subscribe to e-notifications to receive an email notification when documents are added to or changed on the Avaya Support website.

About this task

You can subscribe to different types of general notifications, for example, Product Correction Notices (PCN), which apply to any product or a specific product. You can also subscribe to specific types of documentation for a specific product, for example, Application & Technical Notes for Virtual Services Platform 7000.

Procedure

1. In an Internet browser, go to <https://support.avaya.com>.
2. Type your username and password, and then click **Login**.
3. Under **My Information**, select **SSO login Profile**.
4. Click **E-NOTIFICATIONS**.
5. In the GENERAL NOTIFICATIONS area, select the required documentation types, and then click **UPDATE**.

GENERAL NOTIFICATIONS

1/5 Notifications Selected

End of Sale and/or Manufacturer Support Notices	<input type="checkbox"/>
Product Correction Notices (PCN)	<input checked="" type="checkbox"/>
Product Support Notices	<input type="checkbox"/>
Security Advisories	<input type="checkbox"/>
Services Support Notices	<input type="checkbox"/>

UPDATE >>

6. Click **OK**.
7. In the **PRODUCT NOTIFICATIONS** area, click **Add More Products**.

PRODUCT NOTIFICATIONS

Add More Products

Show Details

1 Notices

8. Scroll through the list, and then select the product name.
9. Select a release version.
10. Select the check box next to the required documentation types.

The screenshot shows a web interface with two main panels. The left panel, titled 'PRODUCTS', has a 'My Notifications' link in the top right. It contains a list of product names: Virtual Services Platform 7000, Virtualization Provisioning Service, Visual Messenger™ for OCTEL® 250/350, Visual Vectors, Visualization Performance and Fault Manager, Voice Portal, Voice over IP Monitoring, W310 Wireless LAN Gateway, WLAN 2200 Series, and WLAN Handset 2200 Series. The right panel, titled 'VIRTUAL SERVICES PLATFORM 7000', features a 'Select a Release Version' dropdown menu set to 'All and Future'. Below this are several items with checkboxes: Administration and System Programming, Application Developer Information, Application Notes, Application and Technical Notes (checked), Declarations of Conformity, and Documentation Library (checked). A red 'SUBMIT >>' button is located at the bottom right of the right panel.

11. Click **Submit**.

Appendix A: SPBM reference architectures

Reference architectures

SPBM has a straightforward architecture that simply forwards encapsulated C-MACs across the backbone. Because the B-MAC header stays the same across the network, there is no need to swap a label or perform a route lookup at each node. This architecture allows the frame to follow the most efficient forwarding path from end to end.

The following reference architectures illustrate SPBM with multiple switches in a network.

For information about solution-specific architectures like Video Surveillance or Data Center implementation using the VSP switch, see [Solution-specific reference architectures](#) on page 590.

The following figure shows the MAC-in-MAC SPBM domain with BEBs on the boundary and BCBs in the core.




Figure 65: SPBM basic architecture

Provisioning an SPBM core is as simple as enabling SPBM and IS-IS globally on all the nodes and on the core facing links. To migrate an existing edge configuration into an SPBM network is just as simple.

The boundary between the MAC-in-MAC SPBM domain and the 802.1Q domain is handled by the BEBs. At the BEBs, VLANs or VRFs are mapped into I-SIDs based on the local service provisioning.

Services (whether Layer 2 or Layer 3 VSNs) only need to be configured at the edge of the SPBM backbone (on the BEBs). There is no provisioning needed on the core SPBM nodes.

The following figure illustrates an existing edge that connects to an SPBM core.




Figure 66: Access to the SPBM Core

For Layer 2 virtualized bridging (Layer 2 VSN), identify all the VLANs that you want to migrate into SPBM and assign them to an I-SID on the BEB.

For Layer 3 virtualized routing (Layer 3 VSN), map IPv4-enabled VLANs to VRFs, create an IP VPN instance on the VRF, assign an I-SID to the VRF, and then configure the desired IP redistribution of IP routes into IS-IS.

All BEBs that have the same I-SID configured can participate in the same VSN. That completes the configuration part of the migration and all the traffic flows return to normal operation.

Campus architecture

For migration purposes, you can add SPBM to an existing network that has SMLT configured. In fact, if there are other protocols already running in the network, such as Open Shortest Path First (OSPF), you can leave them in place too. SPBM uses IS-IS, and operates independently from other protocols. However, it is recommended that you eventually eliminate SMLT in the core and eliminate other unnecessary protocols. This reduces the complexity of the network and makes it much simpler to maintain and troubleshoot.

Whether you configure SMLT in the core, the main point to remember is that SPBM separates services from the infrastructure. For example, in a large campus, a user may need access to other

sites or data centers. With SPBM you can grant that access by associating the user to a specific I-SID. With this mechanism, the user can work without getting access to confidential information of another department.

The following figure depicts a topology where the BEBs in the edge and data center distribution nodes are configured in SMLT clusters. Prior to implementing SPBM, the core nodes would also have been configured as SMLT clusters. When migrating SPBM onto this network design, it is important to note that you can deploy SPBM over the existing SMLT topology without network interruption. After the SPBM infrastructure is in place, you can create VSN services over SPBM or migrate them from the previous end-to-end SMLT-based design.




Figure 67: SPBM campus without SMLT

After you migrate all services to SPBM, the customer VLANs (C-VLANs) will exist only on the BEB SMLT clusters at the edge of the SPBM network. The C-VLANs will be assigned to an I-SID instance and then associated with either a VLAN in an Layer 2 VSN or terminated into a VRF in an Layer 3 VSN. You can also terminate the C-VLAN into the default router, which uses IP shortcuts to IP route over the SPBM core.

In an SPBM network design, the only nodes where it makes sense to have an SMLT cluster configuration is on the BEB nodes where VSN services terminate. These are the SPBM nodes where C-VLANs exist and these C-VLANs need to be redundantly extended to non-SPBM devices such as Layer 2 edge stackable switches. On the BCB core nodes where no VSNs are terminated and no Layer 2 edge stackables are connected, there is no longer any use for the SMLT clustering functionality. Therefore, in the depicted SPBM design, the SMLT/vIST configuration can be removed

from the core nodes because they now act as pure BCBs that simply transport VSN traffic and the only control plane protocol they need to run is IS-IS.

Because SMLT BEB nodes exist in this design (the edge BEBs) and it is desirable to use equal cost paths to load balance VSN traffic across the SPBM core, all SPBM nodes in the network are configured with the same two B-VIDs.

Where the above figure shows the physical topology, the following two figures illustrate a logical rendition of the same topology. In both of the following figures, you can see that the core is almost identical. Because the SPBM core just serves as a transport mechanism that transmits traffic to the destination BEB, all the provisioning is performed at the edge.

In the data center, VLANs are attached to Inter-VSNs that transmit the traffic across the SPBM core between the data center on the left and the data center on the right. A common application of this service is VMotion moving VMs from one data center to another.

The following figure uses IP shortcuts that route VLANs. There is no I-SID configuration and no Layer 3 virtualization between the edge distribution and the core. This is normal IP forwarding to the BEB.

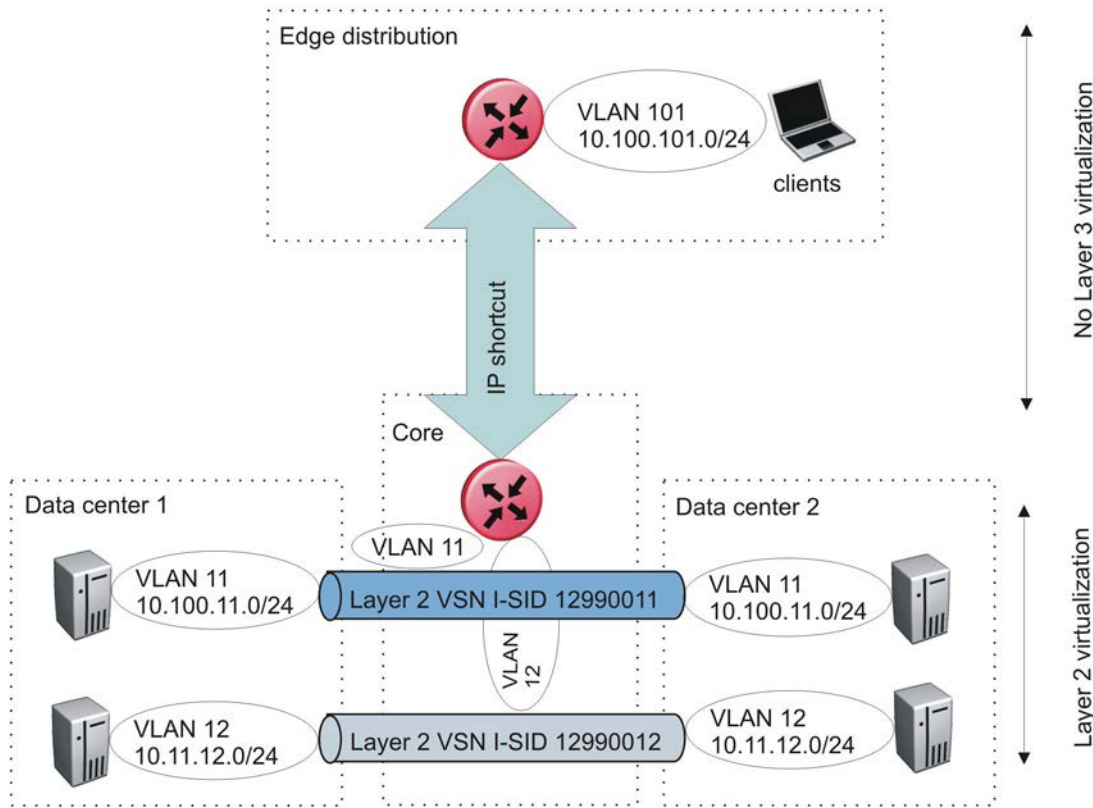


Figure 68: IP shortcut scenario to move traffic between data centers

The following figure uses Layer 3 VSNs to route VRFs between the edge distribution and the core. The VRFs are attached to I-SIDs and use Layer 3 virtualization.




Figure 69: VRF scenario to move traffic between data centers

Large data center architecture

SPBM supports data centers with IP shortcuts, Layer 2 VSNs, or Layer 3 VSNs. If you use vMotion, you must use Layer 2 between data centers (Layer 2 VSN). With Layer 2 VSNs, you can add IP addresses to the VLAN on both data centers and run Virtual Router Redundancy Protocol (VRRP) between them to allow the ESX server to route to the rest of the network.

The following figure shows an SPBM topology of a large data center. This figure represents a full-mesh data center fabric using SPBM for storage over Ethernet. This topology is optimized for storage transport because traffic never travels more than two hops.

* Note:

It is recommended that you use a two-tier, full-mesh topology for large data centers.




Figure 70: SPBM data center—full mesh

Traditional data center routing of VMs:

In a traditional data center configuration, the traffic flows into the network to a VM and out of the network in almost a direct path.

The following figure shows an example of a traditional data center with VRRP configured. Because end stations are often configured with a static default gateway IP address, a loss of the default gateway router causes a loss of connectivity to the remote networks. VRRP eliminates the single point of failure that can occur when the single static default gateway router for an end station is lost.




Figure 71: Traditional routing before moving VMs

A VM is a virtual server. When you move a VM, the virtual server is moved as is. This action means that the IP addresses of that server remain the same after the server is moved from one data center to the other. This in turn dictates that the same IP subnet (and hence VLAN) exist in both data centers.

In the following figure, the VM moved from the data center on the left to the data center on the right. To ensure a seamless transition that is transparent to the user, the VM retains its network connections through the default gateway. This method works, but it adds more hops to all traffic. As you can see in the figure, one VM move results in a complicated traffic path. Multiply this with many moves and soon the network look like a tangled mess that is very inefficient, difficult to maintain, and almost impossible to troubleshoot.




Figure 72: Traditional routing after moving VMs

Optimized data center routing of VMs:

Two features make a data center optimized:

- VLAN routers in the Layer 2 domain (green icons)
- VRRP BackupMaster

The VLAN routers use lookup tables to determine the best path to route incoming traffic (red dots) to the destination VM.

VRRP BackupMaster solves the problem of traffic congestion on the vIST. Because there can be only one VRRP Master, all other interfaces are in backup mode. In this case, all traffic is forwarded over the vIST link towards the primary VRRP switch. All traffic that arrives at the VRRP backup interface is forwarded, so there is not enough bandwidth on the vIST link to carry all the aggregated riser traffic. VRRP BackupMaster overcomes this issue by ensuring that the vIST trunk is not used in such a case for primary data forwarding. The VRRP BackupMaster acts as an IP router for packets destined for the logical VRRP IP address. All traffic is directly routed to the destined subnetwork and not through Layer 2 switches to the VRRP Master. This avoids potential limitation in the available vIST bandwidth.

The following figure shows a solution that optimizes your network for bidirectional traffic flows. However, this solution turns two SPBM BCB nodes into BEBs where MAC and ARP learning will be enabled on the Inter-VSN routing interfaces. If you do not care about top-down traffic flows, you can omit the Inter-VSN routing interfaces on the SPBM BCB nodes. This makes the IP routed paths top-down less optimal, but the BCBs remain pure BCBs, thus simplifying core switch configurations.




Figure 73: Optimized routing before moving VMs

In the traditional data center, chaos resulted after many VMs were moved. In an optimized data center as shown in the following figure, the incoming traffic enters the Layer 2 domain where an edge switch uses Inter-VSN routing to attach an I-SID to a VLAN. The I-SID bridges traffic directly to the destination. With VRRP BackupMaster, the traffic no longer goes through the default gateway; it takes the most direct route in and out of the network.




Figure 74: Optimized routing after moving VMs

Related links

[Solution-specific reference architectures](#) on page 590

Solution-specific reference architectures

The following sections describe solution-specific reference architectures, like for example for Video Surveillance or Data Center implementation, using the VSP 4000.

Multi-tenant — fabric connect

This fabric connect-based solution leverages the fabric capabilities of the VSP platforms: a VSP 7000 core and a VSP 4000 edge. This solution provides the ability to run, by default, up to 24 VRFs for each wiring closet and is well suited for multi-tenant applications. The zero-touch core is enabled by the fabric connect endpoint provisioning capabilities.

*** Note:**

You can increase VRF scaling to run more than 24 VRFs. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Release Notes*.

If this solution must support IPv6, then a central router-pair routes all IPv6 traffic. The IPv6 traffic is tunneled from each wiring closet to the IPv6 routers by extending Layer 2 VSNs to the q-tagged router interfaces.

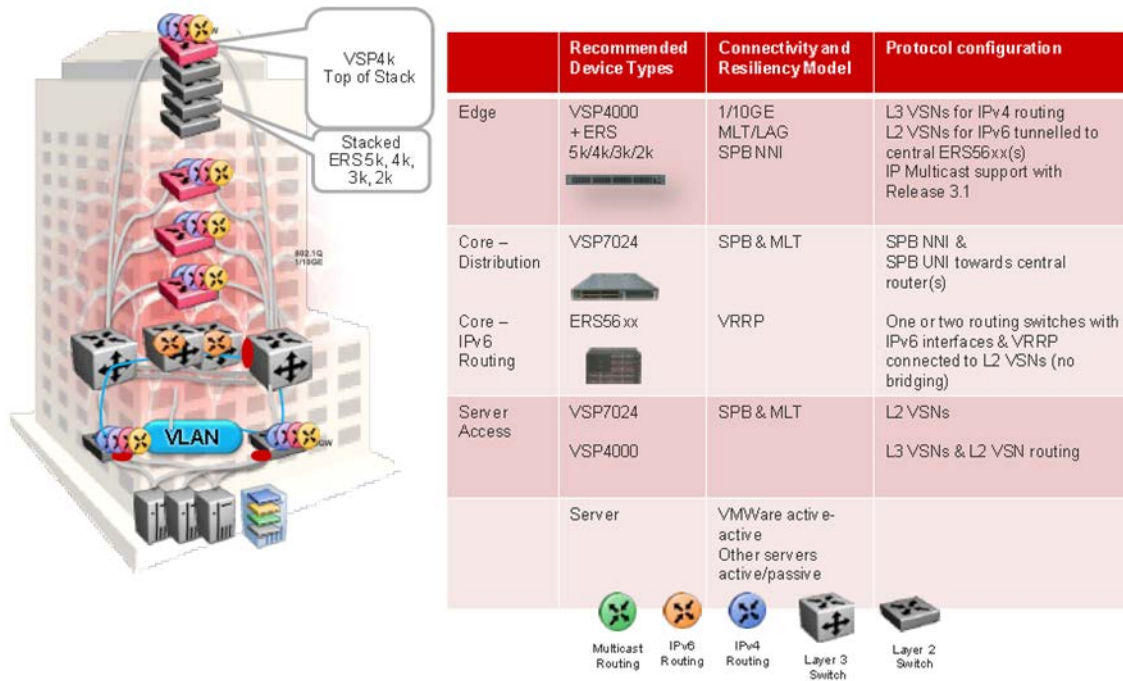


Figure 75: Small core — multi-tenant

The following list outlines the benefits of the fabric connect-based solution:

- Endpoint provisioning
- Fast failover
- Simple to configure
- Layer 2 and Layer 3 virtualized

Hosted data center management solution — E-Tree

In some hosted data center solutions, the hosting center operating company takes responsibility for managing customer servers. For this shared management, shown in the following figure, servers that control the operating system level of the production servers, such as the patch level, are deployed. Because customer production servers do not communicate with each other, a distributed private VLAN solution based on fabric connect is deployed to manage all production servers. This solution builds a distributed set of E-Trees for each management domain.

The VSP switches as access, provide an elegant network-wide E-Tree solution. Spokes, or managed servers, cannot communicate to each other over this network, but the shared management servers on the hub ports can access all spokes. Because of the Layer 2 – E-Tree nature of this setup, the managed servers do not require any route entries, and only require one IP interface in this management private VLAN. This solution supports tagged and untagged physical and virtual (VM) servers.

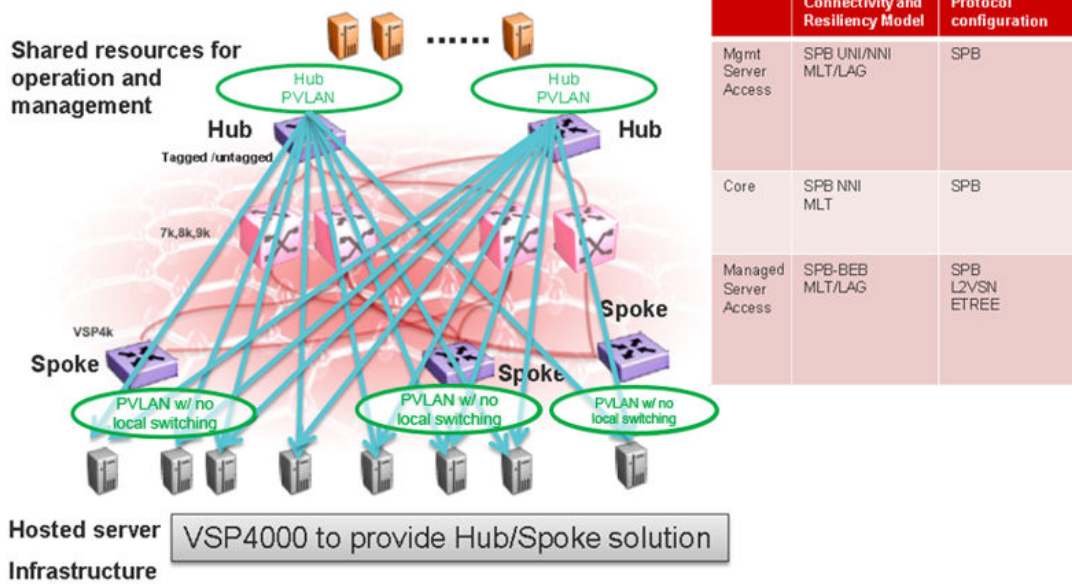


Figure 76: Data center hosting private VLAN

The following list outlines the benefits of the hosted data center management solution:

- Easy endpoint provisioning
- Optimal resiliency
- Secure tenant separation

Video surveillance — bridged

In a video surveillance solution, optimal traffic forwarding is a key requirement to ensure proper operation of the camera and recorder solutions. However, signaling is also important to ensure quick channel switching. This is achieved by deploying a fabric connect based IP multicast infrastructure that is optimized for multicast transport, so that the cameras can be selected quickly, and so that there is no unnecessary traffic sent across the backbone.




Figure 77: Deployment scenario — bridged video surveillance and IP camera deployment for transportation, airports, and government

The following list outlines the benefits of the bridged video surveillance solution:

- Easy end-point provisioning
- sub second resiliency and mc forwarding
- secure tenant separation
- quick camera switching

Video surveillance — routed

In a video surveillance solution, optimal traffic forwarding is a key requirement to ensure proper operation of the camera and recorder solutions. However, signaling is also important to ensure quick channel switching. This is achieved by deploying an IP multicast infrastructure that is optimized for multicast transport, so that the cameras can be selected quickly, and so that there is no unnecessary traffic sent across the backbone. In the topology shown in the following figure, each camera is attached to its own IP subnet. In a larger topology, this can reduce network overhead. To increase network scalability, you can attach a set of cameras to a Layer 2 switch that has IGMP, and then connect the cameras to the fabric edge (BEB) which has a routing instance.

In many customer scenarios, surveillance must be separated from the rest of the infrastructure. This can be achieved by deploying a Layer 3 VSN for the surveillance traffic to keep the surveillance traffic isolated from any other tenant.

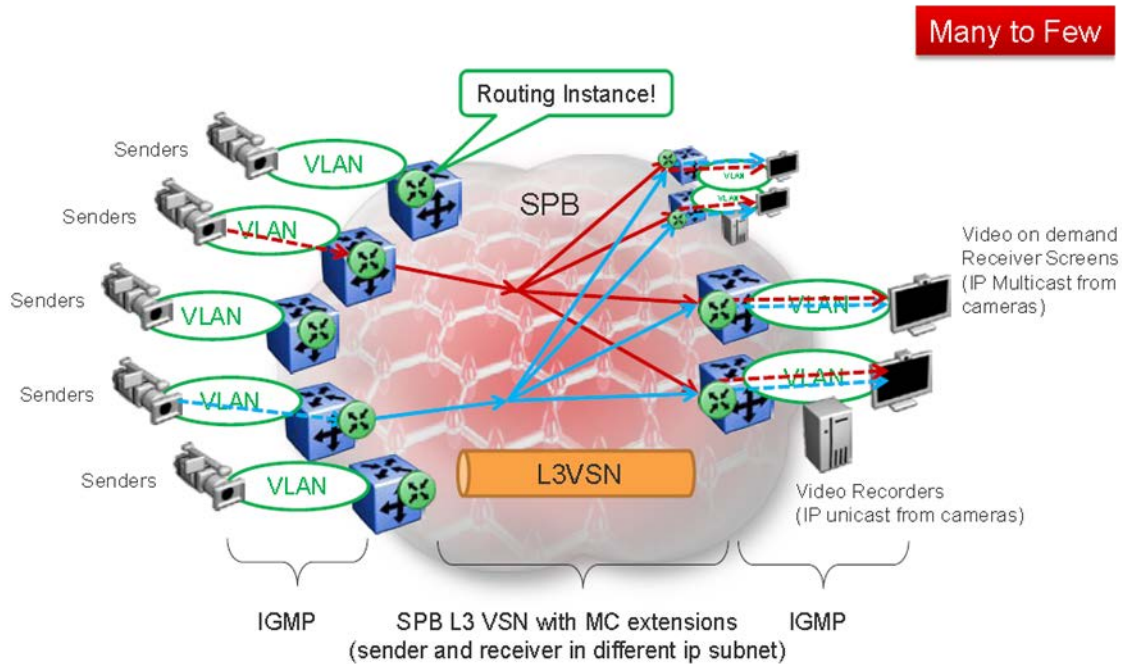


Figure 78: Deployment scenario — Routed video surveillance and IP camera deployment for transportation, airports, and government

The following list outlines the benefits of the routed video surveillance solution:

- Easy endpoint provisioning
- Optimal resiliency and mc forwarding
- Secure tenant separation
- Rapid channel/camera switching

Metro-Ethernet Provider solution

VSP switches provide an end-to-end Metro-Ethernet Provider solution. Leveraging fabric connect throughout the infrastructure enables a scalable and flexible wholesale provider infrastructure.

This use case extends the Transparent Port UNI functionality to transparently forward any customer VLAN across the services.




Figure 79: Metro ring access solution

The following list outlines the benefits of the Metro-Ethernet Provider solution:

- Easy endpoint provisioning
- Optimal resiliency
- Secure tenant separation

Related links

[Reference architectures](#) on page 581

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 Autonomous System, and using an EGP to route packets to other Autonomous Systems.
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.
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 Path 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 B-MAC-DA and B-MAC-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.

BGP+	BGP+ is an extension of BGPv4 to support IPv6. BGP+ carries routing information for multiple network layer protocol address families, for example, IPv6 address family and for IP multicast routes.
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).
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 multicast 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.
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)	<p>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).</p> <p>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.</p>

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.
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 send 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.
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. <ul style="list-style-type: none"> • 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).

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.

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 Fabric Connect, 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 Dijkstra'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 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 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.