



Configuring Fabric Connect on Ethernet Routing Switch 4900 and 5900 Series

Release 7.8.1
9036740-00 Rev. AA
July 2020

© 2017-2020, Extreme Networks, Inc.
All Rights Reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see:
www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>

Contents

Chapter 1: About this Document	6
Purpose.....	6
Conventions.....	6
Text Conventions.....	6
Documentation and Training.....	8
Getting Help.....	9
Providing Feedback.....	10
Chapter 2: New in this document	11
Chapter 3: SPBM and IS-IS Configuration	12
SPBM and IS-IS Fundamentals.....	12
MAC-in-MAC Encapsulation.....	13
I-SID.....	13
BCBs and BEBs.....	14
Basic SPBM Layer 2 VSN topology.....	14
IS-IS.....	16
Standard TLVs.....	16
Fabric Attach Client Discovery and Disconnect Traps.....	18
IS-IS Hierarchies.....	18
IS-IS PDUs.....	19
IS-IS configuration parameters.....	19
NNI-NNI forwarding.....	22
SPBM B-VLAN.....	22
Pre-populated FIB.....	23
RPFC.....	23
SPBM FIB.....	23
SPBM Script.....	24
SPBM and EAP.....	25
SPBM with EAP MHMA-MV.....	26
SPBM and DAI.....	26
SPBM and IPSG.....	26
IP Multicast over Fabric Connect.....	26
E-Tree and Private VLAN.....	31
IPv6 Management over SPB.....	33
Fabric Attach.....	34
SPBM and IS-IS infrastructure configuration using CLI.....	55
Running the SPBM script.....	55
Configuring minimum SPBM and IS-IS parameters.....	56
Displaying global SPBM parameters.....	62
Displaying global IS-IS parameters.....	64

Enabling IP Multicast over Fabric Connect globally.....	66
Displaying IP Multicast over Fabric Connect information.....	68
Displaying IS-IS areas.....	69
Configuring optional SPBM parameters.....	69
Configuring optional IS-IS global parameters.....	72
Configuring optional IS-IS interface parameters.....	76
Displaying IS-IS interface parameters.....	78
Displaying the multicast FIB, unicast FIB, and unicast tree.....	79
Displaying IS-IS LSDB and adjacencies.....	82
Displaying IS-IS statistics and counters.....	85
Configuring I-SIDs for Private VLANs.....	86
Fabric Attach configuration using Command Line Interface	87
SPBM and IS-IS infrastructure configuration using EDM.....	106
Configuring required SPBM and IS-IS parameters.....	106
Configuring IP Multicast over Fabric Connect globally.....	111
Modifying IP Multicast over Fabric Connect globally.....	112
Displaying the SPBM I-SID information.....	113
Displaying Level 1 Area information.....	113
Enabling or disabling SPBM globally.....	114
Configuring SPBM parameters.....	115
Displaying SPBM nicknames.....	116
Configuring interface SPBM parameters.....	116
Configuring SPBM on an interface.....	117
Displaying the unicast FIB.....	118
Displaying the multicast FIB.....	118
Displaying LSP summary information.....	119
Displaying IS-IS adjacencies.....	120
Configuring IS-IS globally.....	121
Configuring system level IS-IS parameters.....	122
Configuring IS-IS interfaces.....	123
Configuring IS-IS interface level parameters.....	125
Configuring an IS-IS Manual Area.....	126
Displaying IS-IS system statistics.....	126
Displaying IS-IS interface counters.....	127
Displaying IS-IS interface control packets.....	128
Fabric Attach configuration using Enterprise Device Manager.....	129
Chapter 4: Layer 2 VSN Configuration.....	141
Layer 2 VSN configuration fundamentals.....	141
SPBM Layer 2 VSN.....	141
SPBM Layer 2 VSN sample operation.....	143
Layer 2 VSN IP Multicast over Fabric Connect.....	149
SPBM IP shortcuts.....	150
Layer 2 VSN configuration using CLI.....	157

- Configuring a SPBM Layer 2 VSN C-VLAN..... 157
- Configuring Layer 2 VSN IP Multicast over Fabric Connect..... 159
- Viewing Layer 2 VSN IP Multicast over Fabric Connect information..... 160
- Viewing IGMP information for Layer 2 VSN multicast..... 161
- Viewing TLV information for Layer 2 VSN IP Multicast over Fabric Connect..... 162
- Configuring the IP Multicast over Fabric Connect forward cache timeout value..... 165
- Configuring the loopback port..... 166
- Viewing the loopback port settings..... 167
- Configuring a SPBM Layer 2 VSN Switched UNI..... 168
- Displaying C-VLAN and Switched UNI I-SID information..... 169
- Managing the switch via Layer 2..... 171
- SPBM IP Shortcuts configuration using CLI..... 172
- Layer 2 VSN configuration using EDM..... 184
 - Configuring SPBM Layer 2 VSN C-VLANs..... 184
 - Displaying the MAC address table for a C-VLAN..... 184
 - Configuring IP Multicast over Fabric Connect on a Layer 2 VSN..... 185
 - Displaying IS-IS redistribution..... 187
 - Configuring SPBM switched UNIs..... 187
 - Managing the switch via Layer 2..... 188
 - SPBM IP Shortcuts configuration using EDM..... 189
- Glossary..... 192**

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

This document provides instructions to configure Fabric Connect on the following platforms:

- Extreme Networks Ethernet Routing Switch 4900 Series
- Extreme Networks Ethernet Routing Switch 5900 Series

Fabric Connect includes Shortest Path Bridging (SPB, the MAC-in-MAC variant of IEEE 802.1aq), Intermediate System to Intermediate System (IS-IS), and Connectivity Fault Management (CFM).

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notice Icons




Icon	Alerts you to...
 Important:	A situation that can cause serious inconvenience.
 Note:	Important features or instructions.
 Tip:	Helpful tips and notices for using the product.

Table continues...




Icon	Alerts you to...
 Danger:	Situations that will result in severe bodily injury; up to and including death.
 Warning:	Risk of severe personal injury or critical loss of data.
 Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	<p>Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.</p> <p>If the command syntax is <code>cfm maintenance-domain maintenance-level <0-7></code>, you can enter <code>cfm maintenance-domain maintenance-level 4</code>.</p>
Bold text	<p>Bold text indicates the GUI object name you must act upon.</p> <p>Examples:</p> <ul style="list-style-type: none"> • Click OK. • On the Tools menu, choose Options.
Braces ({ })	<p>Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command.</p> <p>For example, if the command syntax is <code>ip address {A.B.C.D}</code>, you must enter the IP address in dotted, decimal notation.</p>
Brackets ([])	<p>Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command.</p> <p>For example, if the command syntax is <code>show clock [detail]</code>, you can enter either <code>show clock</code> or <code>show clock detail</code>.</p>
Ellipses (...)	<p>An ellipsis (...) indicates that you repeat the last element of the command as needed.</p> <p>For example, if the command syntax is <code>ethernet/2/1 [<parameter> <value>]...</code>, you enter <code>ethernet/2/1</code> and as many parameter-value pairs as you need.</p>

Table continues...

Convention	Description
<i>Italic Text</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
Plain Courier Text	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. Examples: <ul style="list-style-type: none"> • <code>show ip route</code> • <code>Error: Invalid command syntax</code> <code>[Failed][2013-03-22 13:37:03.303</code> <code>-04:00]</code>
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation tree, expand the Configuration > Edit folders.
Vertical Line ()	A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. For example, if the command syntax is <code>access-policy by-mac action { allow deny }</code> , you enter either <code>access-policy by-mac action allow</code> or <code>access-policy by-mac action deny</code> , but not both.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and software compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other resources](#) such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

[Extreme Portal](#) Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

[The Hub](#) A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

[Call GTAC](#) For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form (all fields are required).
3. Select the products for which you would like to receive notifications.



Note:

You can modify your product selections or unsubscribe at any time.

4. Select **Submit**.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Chapter 2: New in this document

The following sections detail what is new in this document.

Fabric Attach Zero Touch

This release supports updates to the operation of Fabric Attach Zero Touch. For more information, see [FA Zero Touch](#) on page 43.

Chapter 3: SPBM and IS-IS Configuration

This chapter provides conceptual and procedural information related to the configuration and management of Shortest Path Bridging MAC (SPBM) and Intermediate-System-to-Intermediate-System (IS-IS).

SPBM and IS-IS Fundamentals

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

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).

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 deployments to the enterprise market without the complexity of alternative technologies traditionally used in carrier deployments, for example, Multiprotocol Label Switching (MPLS).

The switch also supports IP Shortcuts in SPBM setups. With IP Shortcuts, I-SIDs are no longer needed to forward traffic between BEBs.

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 computation of 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 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:

- For a Layer 2 VSN, the device associates the I-SID with a customer VLAN, which the device then virtualizes across the backbone.

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

MAC-in-MAC Encapsulation

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

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

The encapsulation of customer MAC addresses in backbone MAC addresses greatly improves network scalability, as no end-user MAC learning is required in the backbone, and also significantly improves network robustness, as customer-introduced network loops have no effect on the backbone infrastructure.

Note:

By default, the chassis MAC becomes the B-MAC address for the switch. This address can be used, but it is highly recommended to change the B-MAC to an easy-to-recognize value.

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

*** Note:**

I-SID configuration is required only for virtual services such as Layer 2 VSN.

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.

*** Note:**

If you use IP Shortcuts in the SPBM setup, you must configure it only on the BEBs. No changes are made on the BCBs.

Basic SPBM Layer 2 VSN topology

The following figure shows a basic SPBM network topology, specifically a Layer 2 VSN. Switches B and C are the Backbone Core Bridges (BCB) that form the core of the SPBM network. Switches A and D are the Backbone Edge Bridges (BEB) where the services such as Layer 2 VSNs are provisioned. Only bridges A and B perform both customer MAC (C-MAC) and B-MAC learning and forwarding while bridges B and C only perform B-MAC learning and forwarding.

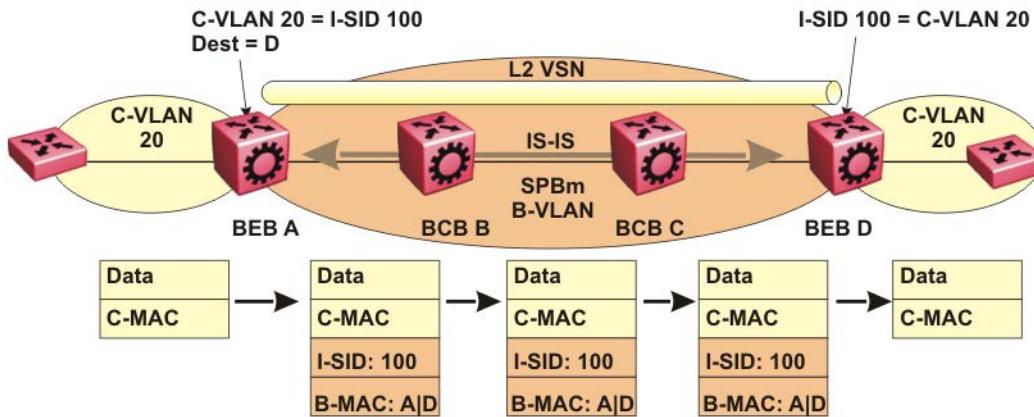


Figure 1: SPBM Layer 2 VSN

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. The C-VLAN ID is only of local significance, as the I-SID defines the service identifier

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

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

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

IS-IS

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 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).

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

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

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 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.

Extreme Networks is in full compliance with the IEEE 802.1aq standard. The following table lists the TLVs that the switch supports.

Figure 2: Standard TLVs

TLV	Description	Usage
1	Area addresses — The Area Addresses TLV contains the area addresses to which the IS-IS is connected.	IS-IS area
22	Extended IS reachability — The Extended IS Reachability TLV contains information about adjacent neighbors.	SPBM link metric Sub TLV (type 29) is carried within this TLV.
129	Protocol 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	TE IP reachability — The Extended IP Reachability TLV 135 is used to distribute IP reachability between IS-IS peers.	SPBM uses this existing IS-IS TLV to carry IP Shortcut routes through the SPBM core.
137	Host name.	
144	Multi-topology Capability (MT-Capability) TLV. This TLV carries the SPB instance ID in a multiple SPB instance environment. This TLV is carried within LSPs.	This TLV carries the following Sub TLVs: <ul style="list-style-type: none"> • SPB instance Sub TLV (type 1): This Sub TLV contains a unique SPSourceID (nickname) to identify the SPBM node within this SPB topology. • SPB Service ID Sub TLV (type 3): This Sub TLV carries service group membership (I-SIDs) for a particular SPB BVLAN.
185	IPVPN multicast TLV with IPMC sub TLV — The IPVPN multicast TLV contains information about the scope I-SID.	TLV 185 on the BEB bridge, where the source is located, displays the multicast source and group addresses and has the transmit (Tx) bit set. Each multicast group has its own data I-SID that maps to the source and group addresses. As part of the IPVPN TLV, sub-TLVs define IPv4 unicast and IPv4

Table continues...

TLV	Description	Usage
		multicast information. Layer 2 VSN IP multicast over SPBM and Layer 3 VSN IP multicast over SPBM (using VRF) use TLV 185.
186	IP multicast TLV (GRE) — TLV 186 on the BEB bridge, where the source is located, displays the multicast source and group addresses and has the transmit (Tx) bit set. Each multicast group has its own data I-SID that maps to the source and group addresses.	IP Shortcuts with IP multicast over SPBM uses TLV 186. All multicast streams are constrained within the level in which they originate, which is called the scope level.

Fabric Attach Client Discovery and Disconnect Traps

An SNMP Trap can be generated by FA Proxy and FA Server devices when a FA Client is discovered and/or when a FA Client disconnects (through timer-based element expiration or link termination). Trap generation is controlled through the existing SNMP Server Notification Control mechanism.

This feature introduces the following traps:

- avFabricAttachDiscoveredElement
- avFabricAttachExpiredElement

Both traps are global (not port-based) and can be enabled or disabled individually.

FA Client discovery and expiration trap generation is enabled by default.

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. IS-IS uses a two-level hierarchy, dividing the domain into multiple Level 1 areas and one Level 2 area. The Level 2 area serves as backbone of the domain, connecting to all the Level 1 areas.

 **Important:**

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

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 within the adjacency holdtime (hello-interval x hello-multiple) seconds, the node tears down the adjacency. In the current release, IIH carries TLV 143 and SPB-B-VLAN Sub-TLV (among other sub-TLVs). For two nodes to form an adjacency the B-VLAN pairs for primary B-VLAN and secondary B-VLAN must match.

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

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

IS-IS configuration parameters

The following sections describe 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 default value.
- Manual area — The manual area or area ID is up to 13 bytes long. The first byte of the area number (for example, 49) is the Authority and Format Indicator (AFI). The next bytes are the assigned domain (area) identifier, which is up to 12 bytes (for example, 49.0102.0304.0506.0708.0910.1112). IS-IS supports a maximum of three manual areas, but the current release only supports one manual area.
- NSEL — The last byte (00) is the n-selector. In the switch 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.

- All IS-IS enabled routers must be in the same area in order to form adjacencies.

PSNP interval

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

CSNP periodic and interval rate

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

Parameters for the link state packet (LSP)

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

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

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

Point-to-point mode

All SPBM links are point-to-point links. 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

The passwords for all authentications are saved as cleartext in the configuration file on the switch. The passwords for simple and HMAC-MD5 are displayed in cleartext through CLI. The HMAC-MD5 packet is encrypted when transmitted over the network.

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

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

Hellos

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

- Interface Hello interval
- Interface Hello multiplier

Interface Hello interval

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

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

Hello multiplier

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

Link metric

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

In this release, only the wide metric is supported.

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

The default value for wide metrics is 10.

Disabling IS-IS

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

Overload bit

If the overload bit parameter is configured, the switch sets the overload bit in its LSP. The `overload` parameter works in conjunction with the `overload-on-startup` parameter. When the `overload-on-startup` timer expires, the SPBM node clears the overload bit and re-advertises its LSP.

When an LSP with an overload bit is received from a neighboring transit-capable SPBM device, the switch ignores the LSP in its SPF calculation so that the transit traffic will not go through the overloaded node. The overloaded node can still receive traffic that is destined for the node itself. The overload bit is usually enabled on stub nodes, which are not used for traversing traffic. By default, overload is set to false on the switch and can be modified.

Stack to Standalone transition

IS-IS settings are kept on non-base unit after base unit fails only if stack forced-mode was enabled on the former stack.

NNI-NNI forwarding

When the switch advertises the IS-IS Link State Database (LSDB) without overload to other nodes the forwarding path is calculated through this switch. When the forward path is established the switch forwards the traffic that was meant for the other nodes. This feature allows forwarding of all unicast and multicast traffic for both scope ISIDs and data ISIDs.

NNI-NNI forwarding is enabled by default but it does not override previous settings after an upgrade to maintain backward compatibility. Overload must be disabled on the switch in order for NNI-NNI forward to work.

SPBM B-VLAN

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

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

Ports cannot be added to a B-VLAN manually, IS-IS takes care of adding ports to the B-VLAN. Ports assigned by IS-IS into B-VLAN are automatically tagged and port state is not restored after IS-IS is disabled.

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.

*** Note:**

When configuring a VLAN ID (VID) for a B-VLAN, some VIDs might be unavailable due to other system features. For example, the STP tagged PBDUs default VID range is 4001–4008. Tagged BPDUs cannot use the same VID as an active B-VLAN.

For more information, see [Configuring VLANs, Spanning Tree, and MultiLink Trunking on Ethernet Routing Switch 4900 and 5900 Series](#).

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 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(config-if)#show isis spbm unicast-fib
=====
                        SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION          BVLAN  SYSID          HOST-NAME          OUTGOING          COST
ADDRESS              ADDRESS              NAME              INTERFACE
-----
00:01:20:00:00:d1   1000   0001.2000.00d1  D1                 Port: 37          10
00:01:20:00:00:d1   1001   0001.2000.00d1  D1                 Port: 37          10
00:01:20:00:00:d2   1000   0001.2000.00d2  D2                 Port: 37          20
00:01:20:00:00:d2   1001   0001.2000.00d2  D2                 Port: 37          20
00:01:20:00:00:d3   1000   0001.2000.00d3  D3                 Port: 37          20
00:01:20:00:00:d3   1001   0001.2000.00d3  D3                 Port: 37          20
00:01:20:00:00:d4   1000   0001.2000.00d4  D4                 Port: 37          20
00:01:20:00:00:d4   1001   0001.2000.00d4  D4                 Port: 37          20
```

Multicast FIB

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

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



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(config)#show isis spbm multicast-fib
=====
                        SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA             ISID   BVLAN  SYSID          HOST-NAME          OUTGOING
-INTERFACES
-----
03:00:61:00:00:64   100    10     0080.2dc1.37ce  4000-1             4/7
03:00:61:00:00:c8   200    10     0080.2dc1.37ce  4000-1             4/2,4/1

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

SPBM Script

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

SPBM and SPBM reserved port must be enabled before running the SPBM script.

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 CLI 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 appears.

SPBM and EAP

The SPBM and EAP feature introduces EAP functionality for C-VLANs in SPBM environments.

SPBM with EAP MHMA-MV

The SPBM with EAP MHMA-MV feature introduces EAP MHMA-MV functionality for C-VLANs in SPBM environments.

For more information, see [Configuring Security on Ethernet Routing Switch 4900 and 5900 Series](#).

SPBM and DAI

The SPBM and DAI feature introduces DAI functionality for BEB devices in SPBM environments.

*** Note:**

You cannot use DAI over SPBM on the uplink port.

SPBM and IPSG

The SPBM and IPSG feature introduces IPSG functionality for BEB devices in SPBM environments.

*** Note:**

You cannot use IPSG over SPBM on the uplink port.

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). 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, such as Open Shortest Path First (OSPF) combined with PIM-SM. OSPF combined with PIM-SM 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.

With IP Multicast over Fabric Connect, extensions to the SPBM IS-IS control plane 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.

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.

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




Figure 3: 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 192.0.2.1.
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) 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 192.0.2.1.
4. The BEB (acting as the IGMP Querier) queries the IS-IS database to find all senders for group 192.0.2.1. 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 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.

* Note:

In the context of IP Multicast over Fabric Connect, scope is the I-SID value of the Layer 2 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.

Data I-SIDs are allocated to streams as they are learned. The multicast traffic is load balanced over primary and secondary B-VLANs depending on the data I-SID allocated (and not depending on the scope I-SID as with unicast streams). If the data I-SID is uneven, the stream is sent over primary VLAN. If the data I-SID is even, the stream is sent over secondary VLAN. The Data I-SID allocated to a multicast stream is local to the BEB and has local meaning.

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

IGMP

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

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

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

*** Note:**

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

BEB as IGMP Querier

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

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

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

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

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

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

For more concept and configuration information on IGMP, see [Configuring IP Routing and Multicast on Ethernet Routing Switch 4900 and 5900 Series](#).

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

The current implementation of 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 currently support a configuration option to accept streams with non-local sources.

You must configure static IGMP receivers on the PIM interface that face the 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 limitations

Review the following limitations 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 drops it or ignores it and logs a message in syslog to highlight the event..

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), both unicast and multicast..

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.

The switch supports IGMPv3.

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.

Data I-SID

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

IP address

IP Multicast over Fabric Connect only supports IPv4 multicast traffic.

Supported services

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.
- Ip Shortcuts multicast service—Multicast traffic scope is within GRT.

Removal of partial-default requirement when enabling SPBM

ERS 5900 unit or stack does not reset to partial-default when SPB is enabled if the loopback port is configured on reserved-port front-panel. Configuration settings are maintained after the unit restarts. To enable SPBM (either multicast is used or not) on ERS 5900, you must loopback configuration.

Software upgrade

The reserved front panel ports remain hidden after software upgrade to 7.2 until boot default or until disabling SPBM. After that, when re-enabling SPBM the boot partial-default is removed and the reserved ports are not hidden anymore.

E-Tree and Private VLAN

Private VLANs

Private VLANs provide isolation between ports within a Layer 2 service.

The primary and secondary VLAN make the Private VLAN. Standard VLAN configuration takes place on the primary VLAN. The secondary VLAN is virtual and inherits configuration from the primary VLAN.

Ports in the Private VLAN are configured as isolated, promiscuous, or trunk. There is no default value. If the port type is not defined, it defaults to none. The following are Private VLAN port types:

Port type	Description
Promiscuous (tagged or untagged ports)	Promiscuous ports communicate with all other ports within the Private VLAN. Uses the primary VLAN.
Isolated (tagged or untagged ports)	Isolated ports communicate with the promiscuous ports, but not with any other isolated port. Uses the secondary VLAN.
Trunk (tagged ports)	Trunk ports carry traffic between other port members within the Private VLANs. Accepts either primary or secondary VLAN.

Trunk ports are automatically set as tagged. A port can be a single port or belong to MLT.

For detailed information about Private VLAN interaction with other features, see [Configuring VLANs, Spanning Tree, and MultiLink Trunking on Ethernet Routing Switch 4900 and 5900 Series](#).

E-Tree

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

Transport within the SPBM network is achieved by associating the Private VLAN with an I-SID. Flooded traffic from both promiscuous and isolated devices are 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. This I-SID can be associated with a single VLAN only. The I-SID/Private VLAN binding must be unique throughout the network.
- 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 Private VLAN ID.

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

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.
- Untagged promiscuous ports cannot belong to more than one Private VLAN. Isolated ports tagged or untagged cannot belong to more than one Private VLAN.

- 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.
- The maximum number of Private VLANs is currently limited to 200.
- The secondary VLAN is virtual and inherits configuration from the primary VLAN. This means that the secondary VLAN becomes a reserved VLAN ID and it is not counted to the number of existing VLANs on the switch.
- Non-Private VLAN ports cannot have a Private VLAN set as PVID, but Private-VLAN ports can have a non-Private VLAN set as PVID.
- Due to hardware restrictions, four extra bytes are added internally to each frame and removed before leaving the switch. These four bytes are counted against the MTU, so frames that are 4 bytes or less smaller than MTU are dropped. Due to these four extra bytes the port may become oversubscribed with less than 100% load.
- The primary and secondary I-SID must be the same. The I-SID value used across SPBM cloud must be the same for the same Private VLAN pair. The primary and secondary VLAN values must be the same across boxes, meaning that there is no remapping.

IPv6 Management over SPB

IPv6 management enables configuring an IPv6 address on the management VLAN. The IPv6 Management over SPB feature adds support for IPv6 management over Fabric Connect.

IPv6 Management over SPB is operational in the following conditions:

- IPv6 global admin status is enabled
- the management VLAN is an SPBM C-VLAN
- an IPv6 interface exists on the management VLAN
- no duplicate addresses exists on the management VLAN

IPv6 management over SPB requires installation of UDF filters matching MAC-in-MAC encapsulated ICMPv6 and ICMPv6 with Hop-By-Hop option packets. In order to save resources, the switch installs these filters only when an IPv6 interface exists on the management VLAN and the management VLAN is also an SPBM C-VLAN.

The switch generates the following syslog messages that describe the installation status for the UDF filters:

- “IPv6 over SPBM filters successfully installed.”
- “IPv6 over SPBM filters could NOT be installed.”

Note:

When IPv6 forwarding is enabled and IPv6 is configured on the management VLAN over C-VLAN, forwarding over C-VLAN only allows management functionality. Forwarding IPv6 over SPB at Line Rate is not supported in this release.

Fabric Attach

The following sections provide conceptual information to help you understand and configure Fabric Attach on switch.

Fabric Attach Fundamentals

Fabric Attach (FA) extends the fabric edge to devices that do not support Shortest Path Bridging MAC (SPBM). With FA, non-SPBM devices can take advantage of full SPBM support, when support is available.

FA also decreases the configuration requirements on SPBM devices by off-loading some configuration to the attached non-SPBM devices and by automating certain configuration steps that occur most often.

FA Signaling

The FA elements communicate between themselves using FA Signaling. FA Signaling is an application level protocol that leverages standard network protocols, such as LLDP, to exchange messages and data between FA elements to orchestrate network automation.

FA Network Elements

The FA architecture involves the following FA elements:

- **FA Server**—An SPB capable network device connected to the fabric edge running the FA agent in FA Server mode. FA Servers receive requests to create services with specific I-SID/VLAN bindings.

In the SPBM architecture an FA Server is a BEB. FA servers process requests for service creation from FA Proxy and/or FA Clients. An FA Server can operate in SPBM or VLAN provisioning mode.

- **FA Proxy**—A device running the FA agent in FA Proxy mode.

An FA Proxy device may be capable of running SPB or not. SPB is always disabled on devices running FA Proxy. FA Proxy mode is enabled by default on devices supporting this mode.

FA Proxies support I-SID/VLAN assignment definition and have the ability to advertise these assignments for possible use by an FA Server, if connectivity permits.

- **FA Client**—A non-SPB network attached device running the FA agent in FA Client mode and able to advertise ISID/VLAN binding requests for service creation to an FA Proxy or FA Server. Non-FA clients without an FA agent are supported through the FA EAP support.
- **FA Standalone Proxy**—An FA device running the FA agent in FA Standalone Proxy mode. FA Standalone Proxy supports FA Proxy functionality in environments without an FA Server.

An FA Standalone Proxy can be used to automate the configuration of traditional VLANs for devices connected to it, such as WLAN Access Points.

The FA Standalone Proxy does not send provisioning requests upstream. An FA Standalone Proxy automatically accepts requests from FA clients and assumes that the upstream network has been provisioned appropriately.

FA Standalone Proxy can be used in environments where the devices upstream from the FA Standalone Proxy do not support Fabric Attach, but the devices downstream from it support Fabric Attach.

FA Server, FA Proxy and FA Standalone Proxy devices use FA signaling in conjunction with Extreme Networks Identity Engines in order to automate configuration of services.

FA Element Discovery

An FA agent which controls FA functionality resides on all FA-capable devices (FA Server, FA Proxy, FA Standalone Proxy or FA Client). No agent-specific configuration is necessary.

FA Proxy and FA Server elements control FA through a global FA service setting (global SPBM setting) and through per-port settings that control the transmission of FA information using FA Signaling.

The first stage of establishing FA connectivity involves element discovery. In order for FA discovery to function, FA service and per-port settings must be enabled. Once these settings are enabled, the FA agent advertises its capabilities (FA Server, FA Proxy or FA Client) through FA Signaling. Following discovery, an FA agent is aware of all FA services currently provided by the network elements to which it is directly connected. Based on this information, an FA Client or an FA Proxy agent can determine whether FA data (I-SID/VLAN assignments) should be exported to an FA Proxy that acts as an external client proxy or an FA Server.

The FA service is enabled by default on FA Servers and FA Proxies. It is disabled by default on FA Standalone Proxy-only devices. Per-port settings are, by default, enabled on FA Proxies and disabled on FA Servers.

Port VLAN tagging mode updates occur when an element is discovered, provided that base zero-touch functionality is enabled. When an element is deleted or expires, all updated settings are cleared and roll back to their previous values.

Note:

An FA Proxy can communicate with, at most, one FA Server at a time. If multiple server connections exist, the first discovered server is considered the primary server. Multiple links (trunked) to a single server are supported as long as they form a logical interface. Multiple non-trunked links are not supported and data received on non-primary ports is ignored by an FA Proxy. FA Proxies or FA Clients can connect through a LAG/MLT to two FA Servers which form a Split-LAG or SMLT pair. Connections which may create loops, to multiple servers that are not in Split-LAG or SMLT mode, are not supported.

An FA Server can communicate with multiple, different FA Proxies and FA Clients.

FA Agent Startup and Initialization

During the FA agent startup and initialization sequence, the following are restored from non-volatile memory:

- FA service status
- FA port-level settings
- external client proxy status

- message authentication status and keys for all ports
- previously configured I-SID/VLAN assignments
- Auto Provision status
- Zero Touch settings
- FA Standalone Proxy settings
- extended logging support

In a stack environment, FA agent startup and initialization occurs on every unit in the stack, using the data restored from non-volatile memory.

The initialization sequence can also include operations geared towards cleaning-up settings that were previously configured in support of FA I-SID/VLAN assignments that were active on an FA Proxy or an FA Server before a system reset.

*** Note:**

After a reboot the switch does not retain configurations (such as DHCP Snooping, ARP Inspection, IP Source Guard) you apply on a dynamic VLAN created by the FA agent.

FA Proxy I-SID-to-VLAN Assignment

Although administrators may configure I-SID-to-VLAN bindings on FA Proxies, I-SID-to-VLAN bindings are typically received by FA Proxies from FA Clients. If external client proxy support is enabled, standard processing requirements for bindings received from an FA Client are managed the same way that processing requirements for locally configured bindings are managed.

Each I-SID/VLAN association that is configured on an FA Proxy creates a Customer VLAN (C-VLAN) User-Network Interface (UNI), once the assignment becomes active following acceptance by an FA Server.

*** Note:**

FA Proxy devices only support C-VLAN UNIs and don't support switched UNIs.

If an I-SID-to-VLAN assignment is accepted by the FA Server, the assignment state is updated to *active*. If an I-SID-to-VLAN assignment is not accepted by the FA Server, the assignment state is updated to *rejected*.

The FA Proxy receives and displays assignment status information from the FA Server for each pending I-SID-to-VLAN assignment. Possible responses include:

- Assignment accepted (2)
- Rejection: generic (3)
- Rejection: Fabric Attach resources unavailable (4)
- Rejection: VLAN invalid (6)
- Rejection: VLAN resources unavailable (8)
- Rejection: application interaction issue (9)

*** Note:**

Data exchanges (I-SID/VLAN assignments) between an FA Proxy and an FA Server/FA Client are supported, as are exchanges between an FA Server and an FA Proxy/FA Client. FA Proxy to FA Proxy and FA Server to FA Server interactions are not supported.

If the FA Proxy or FA Client has access to an FA Server, these assignments are advertised for possible use by the FA Server, using FA signaling.

All I-SID/VLAN assignments defined on an FA Proxy, as well as those received from FA Clients when client proxy operation is enabled, start in the 'pending' state. The I-SID/VLAN assignment state is updated based on feedback received from the FA Server. If an assignment is accepted by the FA Server, its state is updated to 'active'. A server can also reject proposed I-SID/VLAN assignments. In this case, the assignment state is updated to 'rejected'. Data describing the reason for the rejection may also be available.

FA Data Processing

Following discovery, an FA Proxy or FA Client transmits locally-defined I-SID/VLAN assignments through FA Signaling to an FA Server, which accepts or rejects these assignments.

The I-SID/VLAN assignment acceptance by the server can require actions to be performed by the FA agent on both the FA Proxy and the FA Server, to appropriately configure the communication channel (uplink) between the FA Proxy or FA Client and FA Server. Most actions undertaken based on assignment acceptance are undone when the I-SID/VLAN assignment is no longer needed.

I-SID/VLAN assignment rejection by the FA Server requires the FA Proxy to clean up any settings that the FA agent made related to feature operation, as well as log the rejection and any associated error type information for later analysis by an administrator. The amount of clean-up required depends on whether the port VLAN membership was established by the FA Proxy agent or by the administrator outside of the FA feature operation. An uplink port that is associated with a VLAN because of an accepted FA Proxy I-SID/VLAN assignment, and not because of an explicit administrator port VLAN membership action, will have the port VLAN membership cleared when the related I-SID/VLAN assignment is rejected by the FA Server or deleted by the FA Proxy administrator.

VLANs that are automatically created on an FA Proxy due to I-SID/VLAN assignment acceptance are automatically deleted when bindings are rejected or deleted.

No more than a single log message is generated for a rejected I-SID/VLAN assignment, regardless of how many times the assignments have been requested and rejected. Assignments that are rejected, accepted, and later rejected result in a log message being generated for each "new" rejection (two I-SID/VLAN assignment rejection log messages are generated in this case).

FA Proxy I-SID/VLAN assignment addition actions:

- Create port-based VLAN corresponding to I-SID/VLAN assignment VLAN.
- Update port VLAN membership to include I-SID/VLAN assignment VLAN.

FA Server I-SID/VLAN assignment addition actions:

- Create SPBM switched UNI VLAN corresponding to I-SID/VLAN assignment VLAN.
 - C-VLAN join operation does not initiate VLAN creation (VLAN already exists and is associated with the I-SID/VLAN binding I-SID).
- Update I-SID/VLAN mapping data to ensure Shortest Path Bridging-MAC (SPBM)-switched UNI support is enabled for the I-SID/VLAN/port tuple (in other words, create switched UNI). Port VLAN membership is updated by this action.

Additional actions can be required for I-SID/VLAN binding state transitions involving FA Client-generated data. The communication channel (that is, the downlink) between the FA Client and FA Proxy must be appropriately configured. This can require actions to be performed on the switch.

FA Proxy external client proxy I-SID/VLAN assignment addition actions:

- Update downlink port VLAN membership to include I-SID/VLAN assignment VLAN.

Each of these actions is performed by the FA Proxy and FA Server for each I-SID/VLAN assignment, unless the required data/settings have already been configured by the administrator. The successful transition from 'pending' to 'active' is gated by the successful completion of these actions. The FA agent tracks which settings have been updated based on I-SID/VLAN assignment processing (comparing them with settings established by the administrator), and cleans-up or undoes the settings that are related to I-SID/VLAN assignment support as much as possible when an assignment is no longer needed.

I-SID/VLAN assignment state transitions from 'active' to 'rejected' require complementary actions be performed by the FA Proxy and the FA Server to eliminate assignment-related settings:

FA Proxy I-SID/VLAN assignment deletion actions:

- Update uplink port VLAN membership to exclude I-SID/VLAN assignment VLAN.
- Delete port-based VLAN corresponding to I-SID/VLAN assignment VLAN.

FA Server I-SID/VLAN assignment deletion actions:

- Delete I-SID/VLAN/port association data to disable SPBM-switched UNI support for the I-SID/VLAN/port tuple (to delete switched UNI). This action updates port VLAN membership.
- Delete SPBM-switched UNI VLAN corresponding to I-SID/VLAN assignment VLAN.
 - Previously joined C-VLANs are not deleted.

State transitions related to FA Client-generated bindings require additional complementary actions to be performed by the FA Proxy to eliminate assignment-related settings:

FA Proxy external client proxy I-SID/VLAN assignment deletion actions:

- Update downlink port VLAN membership to exclude I-SID/VLAN assignment VLAN.
- Delete port-based VLAN corresponding to I-SID/VLAN assignment VLAN.

Assignment status data returned by the FA Server for each pending I-SID/VLAN assignment drives the FA Proxy response processing. Assignment rejections can include information to indicate the reason for the rejection.

Rejection error codes include:

- FA resources unavailable(4)—the resources that are required for the FA agent to support additional I-SID/VLAN assignments are currently exhausted. The maximum number of assignments that can be supported has been reached.
- VLAN invalid(6)—the specified VLAN can't be used to create a switched UNI at this time. The VLAN already exists and is either inactive or has an incorrect type for this application. This error is also returned if an FA Client or FA Proxy exports an bindings with an I-SID value of 0 and SPBM provisioning is enabled.
- VLAN resources unavailable(8)—the maximum number of VLANs that can be supported by the device has been reached.
- Application interaction issue(9)—a failure has been detected during FA interactions with the VLAN and/or the SPBM applications. The VLAN operations to create the required SPBM switched UNI VLAN or enable port tagging may have failed or the SPBM operation to create the switched UNI may have failed.

As with the actions initiated to support an assignment addition, actions related to assignment deletion are performed only if the targeted data was created during the I-SID/VLAN assignment addition phase. Previously-existing configuration data is not changed. No artifacts are left behind to indicate that automated operations have taken place, following an addition or deletion sequence. This goal may not always be achievable but all attempts are made to satisfy this requirement.

In addition to explicit I-SID/VLAN assignment state transitions, several events can occur that initiate assignment deletion processing. These include:

- I-SID/VLAN assignment timeout—A “last updated” timestamp is associated with all active assignments on the FA Server. When this value is not updated for a predetermined amount of time, the I-SID/VLAN assignment is considered obsolete. Obsolete assignment data and related settings are removed by the FA server agent. The timeout duration value allows FA Server settings to be maintained if temporary connectivity issues are encountered.

I-SID/VLAN binding timeout is also performed by an FA Proxy when it is providing client proxy services and FA Client data is present. Processing similar to that performed by the FA Server related to data aging is supported.

- I-SID/VLAN assignment list updates—The current I-SID/VLAN assignment list is advertised by an FA Proxy at regular intervals (dictated by FA Signaling). During processing of this data, an FA Server must handle list updates and delete assignments from previous advertisements that are no longer present. Though these entries would be processed appropriately when they timeout, the FA agent attempts to update the data in real-time and initiates deletion immediately upon detection of this condition.
- FA Server inactivity timeout—If primary FA Server advertisements are not received for a predetermined amount of time, the I-SID/VLAN assignments accepted by the server are considered rejected. I-SID/VLAN assignment data is defaulted (reverts to the ‘pending’ state) and related settings are removed by the FA Proxy agent. The timeout duration value has been chosen to allow FA Proxy settings to be maintained if temporary connectivity issues are encountered.

You can configure the timeout value used for FA device or binding aging with the `fa timeout` command. The default value is 240 seconds.

FA Proxy and FA Server Connection Maintenance

An FA Proxy can only interact with one FA Server at a time. If multiple server connections exist, the first discovered server is considered the primary server. All other servers discovered after this point in time are considered alternates. Typically only a single FA Server is discovered. If multiple servers are discovered, an indication is logged to identify this situation in case it is not intended. I-SID/VLAN assignment data is only exchanged between the FA Proxy and the primary FA Server.

When using LACP for uplink/downlink trunk, ports should be aggregated into a trunk and the LACP key should explicitly be associated with a MLT ID through the LACP Key/MLT ID mapping table.

Primary server failure is detected using a capabilities advertisement timeout. Once a predefined period of time without an FA Server advertisement from the current primary server expires, the primary server becomes undefined. Any FA Proxy I-SID/VLAN assignments previously accepted by the server are defaulted (reset to the 'pending' state) and related settings are cleared. An informational message (primary server lost) is logged when this transition occurs. I-SID/VLAN assignment data is not advertised until a new primary FA Server is selected. The same algorithm used at startup to select an initial primary server is used to select a new primary server.

FA Proxy/FA Server connectivity using Multi-link Trunking (MLT), Distributed Multi-Link Trunking (DMLT) or Split Multi-Link Trunking (SMLT) connections is supported.

Multiple links associated with the same trunk are treated as a single logical connection. The FA agent reconciles any issues related to MLT, DMLT and SMLT server connectivity and recognizes server uniqueness in the presence of (potentially) multiple capabilities advertisements (that is, FA Signaling received on multiple ports generated by the same server).

In MLT, DMLT and SMLT environments, FA Signaling is generated and received on all links connecting the FA Proxy and FA Server. An FA Proxy receiving an FA Server advertisement determines if a primary FA Server has been selected. If not, the FA Element System ID associated with an advertising FA Server is saved and primary server selection is completed. Once a primary server has been selected, system ID data associated with FA Server advertisements received on other ports is compared against the primary server data. If the system ID values are not the same, an error indication is logged. In all cases, the FA Proxy only generates FA Signaling containing I-SID/VLAN assignment data on the interfaces associated with the primary FA Server.

Note:

The FA Element System ID is structured such that the same system ID is generated on all links associated with a trunk connection between an FA Proxy and an FA Server even in an SMLT scenario where different physical devices are acting as a single logical entity.

In an SMLT environment, an FA Server takes additional actions to ensure that data is synchronized on both SMLT aggregation peers. In this configuration, the FA Server that receives and accepts advertised FA I-SID/VLAN assignments is responsible for generating messages that are sent across the Inter-Switch Trunk (IST) to inform the partner aggregation switch about FA settings that have been configured (for example, SPBM switched UNI VLAN). Similar actions are required when I-SID/VLAN assignments are deactivated.

Agent Stacking functionality

The FA agent is able to function in both standalone and stacked configurations. In a stack, the base unit FA agent acts as the master and pushes its configuration settings to all non-base units (NBUs), to synchronize data across all units. FA agents are active on all units and are able to process stack events as well as data distribution messages.

On an FA Proxy, connections to the primary FA Server can exist on any unit in the stack. When the unit with the active FA Proxy-to-FA primary server interface leaves the stack, any I-SID-to-VLAN assignments accepted by the server are aged-out. I-SID-to-VLAN assignment data is restored to the default *pending* state and related settings are removed by the FA Proxy agent.

The presence of multiple FA Server connections (for example, DMLT FA Proxy - Server connection) is taken into account when determining if FA Server connectivity has been lost.

Note:

For a stack unit failover, low convergence time for LACP or MLT over FA can be achieved by using a stack of three or more units. Starting with 7.8.1, convergence times for LACP over FA are comparable to MLT over FA (a few seconds).

In a stack of two units, a stack unit failover will trigger a stack to standalone transition. This transition requires FA and trunk re-initialization as the standalone switch reconfigures ports to reflect the new state. The convergence times in this case are comparable to the time required for a switch reboot.

FA Message Authentication and Integrity Protection

In order to secure the FA communication in terms of data integrity and authenticity, a keyed-hash message authentication code transmitted with FA TLV data is used to protect all FA signaling exchanges. The standard HMAC-SHA256 algorithm is used to calculate the message authentication code (digest) involving a cryptographic hash function (SHA-256) in combination with a shared secret key. The key is symmetric (known by both source and destination parties). By default, FA message authentication is enabled and a default key is defined to provide secure communication out-of-the-box.

You can configure message authentication status and authentication keys on a per-port basis.

When FA message authentication is enabled, the FA key (default or configured) is used to generate a Hash-based Message Authentication Code (HMAC) digest that is included in all FA TLVs (the FA Element TLV and the FA I-SID/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 not, 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.

Multiple authentication key support provides support for authentication using multiple keys, a user-defined key and a default key. Key usage can be restricted. Only the user-defined key (strict key-mode) or both the user-defined key followed if necessary by the default key (standard key-mode) can be used for authenticating messages. By default, only the user-defined key (strict key-mode) is used for authentication.

Message authentication status, authentication key and key-mode settings are maintained on a per-port basis.

Information related to authentication failures is passed to the EAP/NEAP agent for forwarding to a FA policy server for potential processing when the following criteria are met:

- the interface on which the FA Client is discovered is EAP/NEAP enabled
- the automated FA Client Port Mode Zero Touch option is enabled for FA Client element type

FA Client ingress interface, element type, authentication status, and related key information can be provided for additional upstream client processing.

FA Clients

FA Clients connect to an FA Proxy through standard, non MAC-in-MAC access ports, advertising configured I-SID/VLAN requests to the FA Server. In this scenario, the FA Proxy acts as a client proxy for the FA Client by passing I-SID/VLAN binding requests to a discovered FA Server and returning assignment status information to the FA Client. FA Clients can connect directly to an FA Server, as well.

* Note:

External client proxy support must be enabled on an FA Proxy switch before FA client data is accepted by the FA Proxy. By default, external client proxy support is enabled on an FA Proxy.

I-SID/VLAN bindings received from an FA Client by an FA Proxy acting as a proxy for external clients are processed in much the same way locally administered assignments are processed. FA Proxy response processing takes care of VLAN creation and updates VLAN membership.

If the I-SID/VLAN client assignment is rejected by the FA Server, the FA Proxy performs any required clean-up tasks and also logs the rejection and any associated error type information for later analysis by an administrator.

* Note:

A user assigned to Fail Open VLAN is not removed from I-SID/VLAN bindings using MHSA mode when the RADIUS server becomes unreachable.

FA Auto Provision

You can use Auto Provision with an FA Server-capable device to take advantage of Fabric Attach functionality in non-SPB environments. Auto Provision allows an FA Proxy device (that is also FA Server-capable) to function as an FA Server when SPBM is disabled. With Auto Provision you can designate the device as an FA Proxy or FA Server.

FA VLAN definitions, configured locally on an FA Proxy or through client processing, transparently replace I-SID/VLAN binding definitions in this scenario and allow all of the automated FA processing, with the exception of switched UNI-related operations, to be performed in the absence of SPBM operations. All existing FA default settings remain unchanged.

The Auto Provision support is set to *proxy* by default on an FA Server. The global SPBM setting always overrides the Auto Provision setting, therefore FA operation in an SPBM environment is not impacted at all by Auto Provision.

An FA Server can operate in SPBM or VLAN provisioning mode. In an SPB environment, when SPBM provisioning is operational, for each VLAN associated with an accepted I-SID/VLAN

assignment, the FA Server creates an SPBM switched UNI VLAN , if the VLAN does not already exist. In a non-SPB environment, when VLAN auto-provisioning is operational, the FA Server creates port-based VLANs instead of SPBM switched UNI VLANs.

Once the FA Proxy selects a primary FA Server, the FA Proxy provision mode transitions to the provisioning mode operational on the FA Server.

The current provisioning mode on an FA Server determines the range of I-SID values that are acceptable in the proposed I-SID/VLAN assignment list. When SPBM is enabled, the acceptable I-SID range is 0-16777214. When SPBM Multicast is enabled, the acceptable I-SID range is 0-15999999. When SPBM is disabled and the auto provision mode is set to *server*, the FA Server only accepts bindings with an I-SID value of 0.

FA Zero Touch

FA Zero Touch eases the configuration process on FA-capable devices by automating specific configuration tasks required for FA functionality. For situations when you prefer or require manual configuration of the settings affected by Zero Touch, feature control is provided.

Fabric Attach must be enabled in order for Zero Touch to function. You must manually configure which FA Clients to associate with a Zero Touch option that automates tasks based on FA Clients discovery.

When base Zero Touch functionality is enabled, FA Proxy and FA Client devices can acquire management VLAN data from the connected FA Server or FA Proxy and use it to facilitate manageability and network configuration. When the feature is enabled, base Zero Touch auto-attach operation extracts management VLAN data from the primary FA Server advertisements and potentially uses this data to update the in-use management VLAN. This information can be cascaded to FA Clients, as well.

If the management VLAN being replaced was originally learned by the FA Proxy from FA Element TLV data pushed by the FA Server, the port membership of the now obsolete management VLAN is migrated to the new management VLAN automatically. If there is any user intervention during this automated process (for example, the Zero Touch auto-attach status is modified or the device management VLAN is manually updated) the obsolete management VLAN data remains as is.

Base ZT auto-attach functionality must be enabled in order to support port VLAN tagging mode updates.

Base ZT auto-attach support also transitions the connection between an FA Proxy and an FA Server to *Trusted* if it doesn't already support trusted QoS traffic processing. If the uplink (FA Proxy) or downlink (FA Server) interface is not already associated with a QoS Trusted interface group, a new QoS Trusted interface group ('FaTrustedIfcs') is created if necessary, and the interface is assigned to the interface group. FA Proxy or FA Server connection termination causes the QoS interface group associations to revert to their previous setting, or the default setting if prior setting data is not available.

By default, base Zero Touch support is enabled.

In addition to base Zero Touch functionality, you can configure the following Zero Touch options on an FA device:

IP Address Source Mode Update

When this option is enabled, IP address source mode is updated on the FA Proxy device (receiver) to *DHCP-When-Needed* and initiates DHCP-based IP address acquisition if an IP address is not manually configured.

IP address source mode update only occurs during base Zero Touch processing when a new management VLAN is processed if this option is enabled.

Automated trusted FA Client connection

When this option is enabled, the FA agent examines the list of discovered FA Clients and updates the QoS interface class assignment to 'Trusted' for certain client types, if the interface is not already associated with a 'Trusted' interface class.

QoS interface class assignment data that is updated by FA when an FA Client is discovered resets to the previous QoS interface class assignment when the FA Client information expires. FA Proxy or FA Server connection termination causes the QoS interface group associations to revert to their previous settings, or to default setting if prior setting data is not available. A system reset also causes the QoS interface class assignment for FA-updated interfaces to revert to the previous setting.

All FA updates to QoS settings are dynamic, with the exception of the creation of an FA Trusted QoS interface group ('FaTrustedIfcs').

QoS interface class data is updated based on the discovery and deletion (based on aging) of the following FA Client types:

- Wireless Access Point Type 1
- Wireless Access Point Type 2
- Switch
- Router
- IP Phone
- IP Camera
- IP Video
- Security Device
- Virtual Switch
- Server Endpoint
- ONA SDN
- ONA SPB-over-IP

Automated configuration only applies to FA-enabled ports.

Automated FA Client Port Mode

When this option is enabled and FA Clients are present, the EAP settings for the interface on which the client is discovered, are automatically updated based on the FA Client type. If the FA Clients of the appropriate type are deemed no longer valid (when element aging causes the FA Client to be deleted from the discovered elements list), the EAP port settings revert to the previous state. This is applicable for FA Proxy and FA Server devices.

When EAP port mode configuration is automated using the auto-port-mode-fa-client option, I-SID/VLAN binding data advertised by the FA Client is ignored by the FA agent and is considered

'Untrusted'. All I-SID/VLAN binding installation/deletion is controlled by EAP based on information received from IDE. With the Trusted FA Client feature, certain FA clients can be designated as 'trusted' based on the FA client MAC address and client type. This client designation is passed from EAP/IDE to FA to allow I-SID/VLAN binding data specified by the client to be processed by FA.

When an FA Client is detected on the port, if the auto-port-mode option is enabled in conjunction with a Zero Touch configuration, the port is configured with MHSA. The following settings are applied automatically:

- Global
 - radius-non-eap-enable
 - auto-non-eap-mhlsa-enable
 - eapol enable
- Interface
 - radius-non-eap-enable
 - auto-non-eap-mhlsa-enable
 - mhlsa-no-limit
 - eapol status auto
 - Tagging is set to either UntagPvidOnly or TagAll depending on the FA Client request.

 **Note:**

If the required EAP settings differ from the dynamic ones, it is recommended that a static EAP configuration be applied instead. The port tagging configuration must also be changed manually.

Automated configuration only applies to FA-enabled ports.

The following FA Client types are supported:

- Wireless Access Point Type 1
- Wireless Access Point Type 2
- Switch
- Router
- IP Phone
- IP Camera
- IP Video
- Security Device
- Virtual Switch
- Server Endpoint
- ONA SDN
- ONA SPB-over-IP

*** Note:**

The `auto-port-mode-fa-client` option is incompatible with the Zero Touch Client `auto-client-attach` and `auto-pvid-mode-fa-client` options.

Zero Touch Client installation

Zero Touch Client (ZTC) functionality supports automatically updating port VLAN membership, the port PVID, and possibly the default port priority, based on the presence and type of discovered FA Clients. An I-SID/VLAN binding can be installed, as well.

*** Note:**

The `auto-client-attach` option must be enabled before Zero Touch Client specifications can be applied (either during discovery or retroactively).

*** Note:**

The `auto-client-attach` option is incompatible with both the `auto-port-mode-fa-client` and `auto-pvid-mode-fa-client` options. You cannot enable these Zero Touch options for a client type at the same time.

The following FA Client types are supported:

- Wireless Access Point Type 1
- Wireless Access Point Type 2
- Switch
- Router
- IP Phone
- IP Camera
- IP Video
- Security Device
- Virtual Switch
- Server Endpoint
- ONA SDN
- ONA SPB-over-IP

Automated PVID FA Client Port Mode

Enabling this option initiates automatic port PVID and port management VLAN membership, based on the type of discovered FA Clients. This is applicable for FA Proxy and FA Server devices. Automated configuration is only applied to FA-enabled ports.

Data updated by Automated PVID FA Client Port Mode when an FA Client is discovered resets to the previous value when the FA Client information expires.

PVID and port VLAN data are updated based on the discovery and deletion (based on aging and port events) of the following FA Client types:

- Wireless Access Point Type 1

- Wireless Access Point Type 2
- Switch
- Router
- IP Phone
- IP Camera
- IP Video
- Security Device
- Virtual Switch
- Server Endpoint
- ONA SDN
- ONA SPB-over-IP

The `auto-pvid-mode-fa-client` option does not function over EAPOL.

*** Note:**

The `auto-port-mode-fa-client` option is incompatible with the `auto-pvid-mode-fa-client` option. You cannot enable these Zero Touch options for a client type at the same time.

*** Note:**

The `auto-pvid-mode-fa-client` option is incompatible with with the Zero Touch Client `auto-client-attach` option. You cannot enable these Zero Touch options for a client type at the same time.

Management VLAN Advertisement Blocking

The `fa zero-touch auto-attach` command is augmented with the optional parameter `disable-mgmt-vlan-distribution`. When this parameter is not specified, management VLAN data in the FA Element TLV is included, by default. This parameter causes the management VLAN data in the FA Element TLV to be zeroed indicating to the downstream FA devices that management VLAN data is not being advertised.

*** Note:**

The management VLAN distribution setting does not impact FA agent management VLAN processing or usage in any other way. A management VLAN can still be learned or updated based on FA Server advertisements. Management VLAN port associations can be updated through all current mechanisms, including various zero-touch operations. The Zero Touch Auto-Attach setting overrides the management VLAN distribution setting. When Zero Touch Auto-Attach is disabled, management VLAN advertisement stops regardless of the management VLAN distribution setting.

For more information, see [Disabling Management VLAN Distribution](#) on page 96.

Automatic Management VLAN Assignment

The current `auto-pvid-mode-fa-client` ZT option allows the port PVID and port VLAN membership to be updated following FA Client discovery on a per-client basis using the management VLAN. This

works well for FA Clients generating a mix of tagged and untagged data with management traffic being untagged (FA Proxy port tagging mode on the FA Client link set to **untagPvidOnly**).

For FA Clients that send management traffic tagged (possibly after learning the management VLAN from the FA Proxy), a slightly different option is now supported. The new per-client ZT option **auto-mgmt-vlan-fa-client** updates the port VLAN membership with the switch management VLAN but does not update the port PVID. If **auto-mgmt-vlan-fa-client** is enabled for a specific client type and if the specified FA Client type is discovered, the management VLAN is automatically added to the FA Client port on the switch. The dynamically added management VLAN ID is automatically removed from the port when the FA Client disconnects.

*** Note:**

The **auto-mgmt-vlan-fa-client** option is incompatible with the **auto-pvid-mode-fa-client** and the **auto-port-mode-fa-client** options, as well as with the Zero Touch Client (ZTC) auto-attach support. Attempts to enable these Zero Touch options for a specific client type at the same time are rejected.

To configure Zero Touch options, see [Configuring Fabric Attach Zero Touch options](#) on page 99.

EAP and FA

With EAP and FA, FA-capable switches or stacks can forward traffic from EAP/NEAP clients over the SPB cloud. The traffic for authenticated clients is mapped to I-SIDs received from the Extreme Networks Identity Engines RADIUS server.

You must configure the desired bindings for EAP/NEAP clients on the RADIUS server. When confirming the authentication request, the RADIUS server also sends the corresponding binding for the EAP/NEAP client.

In MHSA and MHMV modes, the VLAN from the I-SID/VLAN binding received from the RADIUS server is automatically created if it is not already present.

The following VLAN types are automatically created:

- port-based VLANs, if the EAP/NEAP client is connected via an FA Proxy
- Switched UNI VLANs, if the EAP/NEAP client is connected via an FA Server

After an EAP/NEAP client is disconnected, the switch cleans-up the binding associated with the client, if no other EAP/NEAP client on that port uses it.

When an EAP/NEAP client successfully authenticates on an FA Proxy, the client port becomes a member of the VLAN from the I-SID/VLAN pair. The FA Proxy sends to the FA Server the binding received from the RADIUS server. If the FA Server rejects all the bindings, the client is disconnected. EAP clients are moved from AUTHENTICATED state to HELD state.

*** Note:**

In case of a rejected binding, a delay of up to 30 seconds may exist from the time the client authenticates on the FA Proxy until the FA Server rejection response is received by the FA Proxy. Therefore, EAP client traffic may flow for up to 30 seconds until dropped.

On an FA Server, when an EAP/NEAP device is authenticated and an FA binding is received from the RADIUS server, a switched-UNI is created. This is automatically cleaned-up when the client is disconnected.

Access Points authentication

In MHSa mode, the switch also supports NEAP authentication for Access Points. Because Access Points cannot authenticate via EAP, the MHSa mode was improved as follows:

- MHSa now allows the first connected client to be a NEAP client. For each MAC seen on the port, the switch sends an Access Request to the RADIUS Server. After the first successful authentication, a configured number of auto-learned clients are granted access, as in previous MHSa behavior.
- a new option, 'no-limit', is available for configuring the switch to support an unlimited number of NEAP auto-learned clients. You can use this option when an Access Point connected to the switch supports an indeterminate number of devices.

Previously, after the first successful EAP authentication, the switch allowed only a limited number of auto-learned NEAP clients.

When the 'no-limit' option is enabled, the port forwards the traffic from all the devices on that port, without limiting their number. When the Access Point disconnects, the switch clears the mac-address-table for that port and blocks again all traffic. By default, the 'no-limit' option is disabled.

* Note:

In FA Standalone mode, the uplink port is automatically added to the Guest VLAN or Fail Open VLAN only when these VLANs are created using the `fa vlan` command.

* Note:

EAP ports configured in MHSa mode with AP detected as an FA client will not be added to the Fail Open VLAN.

* Note:

If SPBM is enabled, the ISID/VLAN bindings does not appear in the `show fa assignment` command. They can be checked for successful auto-creation using `show i-sid`, `show vlan dynamic`, and other VLAN show commands. In the LLDP TLV show commands they might also appear as *Pending* because only EAP is in charge of auto-creating bindings, if passed from RADIUS.

CoA Support for FA

This feature provides support for processing FA binding VSA attributes present in a Change of Authorization (CoA) request.

These FA bindings follow the same rules that apply to FA bindings received from the RADIUS server. The switch supports FA bindings in CoA requests for EAP and NEAP authenticated clients.

* Note:

In this release, the switch supports only the FA binding attributes. If other FA VSA attributes are present in a CoA request, the switch sends a CoA -NAK response.

VSAs

The following is a list of VSAs added to support EAP FA functionality:

VSAs sent from RADIUS server to switch:

- Extreme-Fabric-Attach-VLAN-ISID

This VSA consists of a (VLAN, I-SID) pair.

Multiple (VLAN, I-SID) pairs are processed only in MHSA mode.

- Extreme-Auto-VLAN-Create

If this VSA is set to TRUE, the VLANs received in all (VLAN, I-SID) pairs will be automatically created if they do not exist. This VSA is processed only in MHSA and MHMV modes.

- Extreme-Fabric-Attach-VLAN-PVID

This VSA contains the value of the PVID that should be set on the port with the authenticated client. The Extreme-Fabric-Attach-VLAN-PVID VSA is processed only in MHSA mode.

Trusted FA Client

With Trusted FA Client, binding requests from FA Clients are processed by the FA agent even when the EAP port mode configuration is automated using the auto-port-mode-fa-client option or the port has been manually EAP-enabled.

The following limitations apply to the Trusted FA Client feature:

- Trusted FA Client functions in MHMV and MHSA modes.
- In this release the feature supports only NEAP clients.
- You must allow the use of VLAN IDs assigned by RADIUS for non-EAP clients in order for the feature to function in MHMV mode.
- If the Auto-Create-Vlan VSA is not sent by the RADIUS server, the FA Client binding is processed only when the VLAN already exists on switch.
- The bindings received from a trusted FA Client must meet the same restrictions as bindings received from the RADIUS server.

VSAs sent from IDE to the ERS switch:

- Extreme-Fabric-Attach-Client-Trust

This VSA tells the switch to trust or not to trust the VLAN:ISID binding requests coming from an FA client when the port is in EAP/NEAP mode.

This VSA can have the following values:

- Extreme-Fabric-Attach-Client-Trust = 0 or Not sent: No FA Client VLAN:ISID binding request is trusted.
 - Extreme-Fabric-Attach-Client-Trust = 1: All FA Client VLAN:ISID binding requests are trusted.
 - Extreme-Fabric-Attach-Client-Trust = 2: Some FA Client VLAN:ISID binding requests are trusted (based on the VSA Fabric-Attach-Client-Trusted-Binding).
- Extreme-Fabric-Attach-Client-Trusted-Binding

This VSA enables the administrator to control which VLANs and ISIDs the FA client can request. A maximum of 10 VSAs per client can be sent from Radius Server to the Switch. This VSA has a string maximum length of 64.

This VSA has the following values:

- For Extreme-Fabric-Attach-Client-Trust = 1:

All FA Client VLAN:ISID binding requests are trusted and applied upon RADIUS authentication. You do not need to specify a binding range.

- For Extreme-Fabric-Attach-Client-Trust = 2:

Some FA Client VLAN:ISID binding requests are trusted. You must specify all VLAN:ISID bindings, including the individual VLAN:ISID bindings, as a range.

The format of the VLAN:ISID binding is:

minVlanValue-maxVlanValue:minIsidValue-maxIsidValue, where:

- *minVlanValue* = minimum value of VLAN range
- *maxVlanValue* = maximum value of VLAN range
- *minIsidValue* = minimum value of ISID range
- *maxIsidValue* = maximum value of ISID range

Example:

If an FA client sends VLANs 101 through 110 with ISIDs 10010001 through 10010010, specify the VLAN:ISID bindings as a range, including the individual VLAN:ISID bindings (for example, for VLANs 101 and 108), as follows:

- Extreme-Fabric-Attach-Client-Trusted-Binding += 101-101:10010001-10010001,
- Extreme-Fabric-Attach-Client-Trusted-Binding += 103-105:10010003-10010005,
- Extreme-Fabric-Attach-Client-Trusted-Binding += 108-108:10010008-10010008

VSAs sent from switch to RADIUS server:

- Extreme-Fabric-Attach-Mode

This VSA can have the following values:

- 0 or not sent, when Switch is assumed to have no concept of SPB/AutoProv
- 1, when the switch is an FA Server in VLAN provision mode
- 2, when the switch is an FA Server in SPBM mode
- 3, when the switch is an FA Proxy with the connected FA Server in VLAN provision mode
- 4, when the switch is an FA Proxy with the connected FA Server in SPBM mode
- 5, when the switch is a FA Standalone Proxy

- Extreme-Fabric-Attach-Client-Type

This VSA can have the following values:

- 1, FA Element Type Other
- 2, FA Server

- 3, FA Proxy
- 4, FA Server No Authentication
- 5, FA Proxy No Authentication
- 6, FA Client – Wireless AP Type 1 [clients direct network attachment]
- 7, FA Client – Wireless Ap Type 2 [clients tunneled to controller]
- Fabric-Attach-Client-PSK
This VSA can have the following values:
 - Not sent when PSK used unknown
 - 0, When Dual-key authentication is disabled
 - 10, When FA Client Failed FA TLV authentication using Default PSK
 - 11, When FA Client Passed FA TLV authentication using Default PSK
 - 100, When FA Client Failed FA TLV authentication using User Defined PSK
 - 101, When FA Client Passed FA TLV authentication using User Defined PSK
- Extreme-Fabric-Attach-Client-Id
This VSA contains the MAC address of the FA client, exported via FA Signaling.

FA Standalone Proxy

FA Standalone Proxy introduces FA Proxy functionality in environments without an FA Server. Regardless of whether the FA Standalone Proxy upstream device is a non-Extreme Networks component or an Extreme Networks device on which FA Server functionality is not available, FA Standalone Proxy operation supports standard FA Proxy processing as if an FA Server has been discovered.

You can enable or disable FA Standalone Proxy support. By default, it is disabled. Enabling the FA Standalone Proxy mode enables immediate processing of pending I-SID/VLAN bindings, if other configuration data such as static uplink data allows the processing. Previously established settings based on FA Proxy operation, if present, are reset when FA Standalone Proxy operation is enabled.

Note:

In FA Standalone Proxy mode, I-SID values are not specified and are implicitly 0. Only bindings with an I-SID value equal to 0 are accepted for processing.

Disabling the FA Standalone Proxy mode resets configured I-SID/VLAN binding data to its default state and enables full FA Proxy operation.

In FA Standalone Proxy mode you must provide the FA Server uplink information, which is typically gathered through FA Server discovery. Once you provide this information, FA Standalone Proxy mode operates as if an FA Server has been discovered and is accepting I-SID/VLAN binding requests. The binding clean-up is similar to an FA Server timeout event, and occurs when the static uplink is deleted and when FA Standalone Proxy operation is disabled.

*** Note:**

FA Standalone Proxy and SPBM are mutually exclusive. Enabling SPBM transitions the device to FA Server mode from FA Standalone Proxy mode. The FA service remains enabled. Settings configured by FA Standalone Proxy operation, that is I-SID/VLAN binding data, revert to default values.

*** Note:**

No interactions with an FA Server are supported in FA Standalone Proxy mode.

*** Note:**

Before creating static uplink over a LAG, it's highly recommended to manually bind LACP-key to an MLT-ID.

When using LACP for uplink trunk, ports should be aggregated into trunk.

Edge Automation Enhancements

Edge Automation Enhancements enable dynamic configuration of ports and VLANs that have users or devices connected to them, such as IP cameras, or Access Points.

RADIUS service requests are specified using the Fabric-Attach-Service-Request VSA.

Dynamic Configuration of Port-only Settings

Dynamic configuration can be applied to the following port-only settings:

- ability to configure port speed and duplex
- ability to enable BPDU filtering
- ability to enable SLPP Guard
- ability to enable IP Source Guard
- ability to set traffic-control for Wake on LAN (WoL) capable devices

The RADIUS user configuration attributes, which request the settings are as follows:

- Fabric-Attach-Service-Request = "BPDU"
- Fabric-Attach-Service-Request = "SLPPGUARD"
- Fabric-Attach-Service-Request = "SPEED:<speed>"
- Fabric-Attach-Service-Request = "DUPLEX:<duplex>"

*** Note:**

Values for the DUPLEX attributes require the uppercase characters (HALF/FULL)

- Fabric-Attach-Service-Request = "IPSG"
- Fabric-Attach-Service-Request = "DHCP Snooping82SUBID:<subscriber_id>"
- Fabric-Attach-Service-Request = "WOL"

*** Note:**

IP Source Guard can be enabled only if the port is a member of a Dynamic ARP Inspection and DHCP Snooping enabled VLAN. DHCP must also be enabled globally.

*** Note:**

As a best practice, speed and duplex must be sent from RADIUS if they are to be applied. This ensures that the link has the exact parameters desired for the device in question.

Port settings are applied when the client is the first and the only client authenticated on the port. The configuration settings are ignored if there is at least one authenticated client that pushed a set of port settings. It is expected that all clients connecting to a specific port require the same port settings. When all users on a port disconnect, all settings return to the values configured before the dynamic port settings were applied, with the exception of the traffic-control setting for WoL. For speed and duplex, port auto-negotiation returns to the original state.

*** Note:**

Speed and duplex settings cause a port link-down link-up bounce, which removes the authenticated clients. It is necessary for the clients to reauthenticate immediately if these settings are pushed.

*** Note:**

It is recommended to set MSTP Edge Port to True (or Spanning Tree Fast Learning if in STPG mode) on EAP-enabled ports. This prevents topology change notifications from being sent on that port and MAC addresses will not be cleared on outside topology changes, which prevents EAP clients from re-authenticating. This also speeds up reauthentication after a port bounce caused by changing speed and duplex.

Dynamic Configuration of VLAN Settings

Dynamic configuration can be applied to the following VLAN settings:

- ability to enable Dynamic ARP Inspection
- ability to enable DHCP Snooping
- ability to enable DHCP Snooping Option 82
- ability to enable IGMP snooping

*** Note:**

DHCP Snooping and DHCP Snooping Option 82 have a global setting, which must be enabled statically in order for the feature to function properly. DHCP Snooping and Dynamic ARP Inspection trusted ports should also be configured statically.

The RADIUS user configuration attributes for VLAN settings can specify a single VLAN or a range of VLANs for each setting request.

The RADIUS user configuration attributes, which request the settings are as follows:

- Fabric-Attach-Service-Request += "DAI:<vid>[-<vid>]"

- Fabric-Attach-Service-Request += "DHCP Snooping: <vid>[-<vid>]"
- Fabric-Attach-Service-Request += "DHCP Snooping82: <vid>[-<vid>]"
- Fabric-Attach-Service-Request += "IGMP Snooping: <vid>[-<vid>]"

*** Note:**

IGMP Snooping cannot be enabled on SPBM switched UNI VLANs.

VLAN settings are applied on auto-created VLANs only. VLAN settings are applied if the client is the first and the only client authenticated that pushed the settings. It is assumed that all clients connecting to a specific VLAN require the same VLAN settings. If two or more clients join a VLAN, it is assumed that on that specific VLAN there is a set of enabled features wanted by all the clients joining the VLAN. If a client requests additional settings than those pushed by the first client, those requests are ignored.

VLAN settings persist until the auto-created VLAN is removed. If EAP GuestVLAN or FailOpenVLAN are manually configured with an auto-created VLAN ID, which is not recommended, the setting will be commented in ASCII Command Generator (ACG), and lost after reboot. However, the dynamic VLAN will not be deleted once it runs out of port members, so the Guest or FailOpen setting will persist until reboot.

*** Note:**

The new settings are not applied if the user authentication fails, if the new session is not valid in the FA context, or if Private VLAN context or the bindings are not consistent with the current configuration.

*** Note:**

Both port and VLAN settings are saved in the NVRAM. They are deleted when the clients that request them deauthenticate. Likewise, a soft reboot of the switch deletes these settings.

A hard reboot preserves the port settings, but deletes the settings on the dynamically created VLANs since those VLANs themselves will no longer exist.

SPBM and IS-IS infrastructure configuration using CLI

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

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 the SPBM and SPBM reserved-port before running the SPBM script.

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.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. To run the SPBM script, enter the following command at the command prompt:

```
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>run spbm
```

```
SPBM Secondary Backbone VLAN <2-4059>[4052]:
SPBM Nickname <x.xx.xx>[e.ed.00]: derived from chassis mac
Manual-area <xx.xxxx.xxxx...xxx>[49.0000]:
ISIS Ethernet Interface <List of ports>[:8
ISIS System Name [Switch]:ISIS
System ID <xxxx.xxxx.xxxx>[7030.183e.efdf]: derived from chassis mac %%
Applying SPBM parameters...
%% SPBM Ethertype is set to 0x8100
%% SPBM Backbone VLAN is set to 4051
%% SPBM Secondary Backbone VLAN is set to 4052
%% VLANs 4051,4052 are associated with SPBM instance
%% SPBM Nickname is set to e.ed.00
%% Manual-area is set to 49.0000
%% ISIS is enabled on Ethernet 8
%% ISIS System Name is set to 7024XLS
%% ISIS System ID is set to 7030.183e.efdf
```

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.

Before you begin

Configure the loopback port.

Procedure

1. Enter Global Configuration mode:

```
enable
```


- ```
configure terminal
```
2. Enable SPBM globally:
 

```
spbm
```
  3. Log on to the IS-IS Router Configuration mode:
 

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

```
spbm <1-100>
```
  5. Exit IS-IS Router Configuration mode to Global Configuration mode:
 

```
exit
```
  6. Create the primary SPBM backbone VLAN (B-VLAN):
 

```
vlan create <2-4094> type spbm-bvlan
```
  7. Create the secondary SPBM backbone VLAN (B-VLAN):
 

```
vlan create <2-4094> type spbm-bvlan
```
  8. Log on to the IS-IS Router Configuration mode:
 

```
router isis
```
  9. Add the SPBM B-VLANs to the SPBM instance:
 

```
spbm <1-100> b-vid {<vlan-id [-vlan-id][, ...]} [primary <1-4094>]
```

To remove the specified B-VLAN from the SPBM instance, use the following command:

```
no spbm <1-100> b-vid {<vlan-id [-vlan-id] [, ...]}
```
  10. Configure the system nickname (2.5 bytes in the format <x.xx.xx>):
 

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

To delete the configured nickname, use one of the following commands:

```
no spbm <1-100> nick-name
```

OR

```
default spbm <1-100> nick-name
```
- \* 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 <xxxxx.xxxx.xxxx>` command). This helps to recognize source and destination addresses for troubleshooting purposes.
11. Configure an IS-IS manual area (1-13 bytes in the format <xx.xxxx.xxxx...xxxx>. Only one manual area is supported.):

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

To delete the manual area, use one of the following commands:

```
no manual-area
```

OR

```
default manual-area
```

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

```
exit
```

13. Log on to Interface Configuration mode, by specifying the ports that are going to link to the SPBM network:

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

14. Create an IS-IS circuit and interface on the selected ports:

```
isis
```

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

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

16. Enable the IS-IS circuit/interface on the selected ports:

```
isis enable
```

To disable IS-IS on the specified interface, use the following command:

```
no isis enable
```

17. Exit Interface Configuration mode:

```
exit
```

18. Remove the selected port for IS-IS from the default VLAN.

```
vlan member remove [vlan-id] [port]
```

 **Note:**

By default, all ports are enabled in VLAN 1. Ensure the port for the IS-IS interface is removed from VLAN 1 and all other normal VLANs. Also, disable Spanning Tree participation.

19. Enable IS-IS globally:

```
router isis enable
```

20. Display the SPBM configurations:

```
show isis spbm
```

21. Display the global IS-IS configuration:

```
show isis
```

## 22. 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)# exit
Switch(config)# vlan create 1000 type spbm-bvlan
Switch(config)# vlan create 2000 type spbm-bvlan
Switch(config)# router isis
Switch(config-isis)# spbm 1 b-vid 1000,2000 primary 1000
Switch(config-isis)# spbm 1 nick-name 1.11.16
Switch(config-isis)# manual-area c0.2000.0000.0000
Switch(config-isis)# exit
Switch(config)# interface Ethernet 3
Switch(config-if)# isis
Switch(config-if)# isis spbm 1
Switch(config-if)# isis enable
Switch(config-if)# exit
Switch(config)# vlan member remove 1 3
Switch(config)# router isis enable
Switch(config)# show isis spbm
```

```
=====
 ISIS SPBM Info
=====
SPBM B-VID PRIMARY NICK LSDB
INSTANCE B-VID VLAN NAME TRAP

1 1000,2000 1000 1.11.16 disable
=====
```

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

```
=====
 ISIS General Info
=====
```

```

AdminState : enabled
RouterType : Level 1
System ID :0014.c7e1.33df
Max LSP Gen Interval : 900
Min LSP Gen Interval : 30
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name :
Num of Interfaces : 2
Num of Area Addresses : 1

```

Switch(config)# show isis interface

```

=====
ISIS Interfaces
=====
IFIDX TYPE LEVEL OP-STATE ADM-STATE ADJ UP-ADJ SPBM-L1-METRIC

Mlt2 pt-pt Level 1 UP UP 1 1 10
Port3 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.<br><br>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.xxxx.xxxx...xxxx> | Specifies the IS-IS manual-area in hexadecimal format (1-13 bytes in the format <xx.xxxx.xxxx...xxxx>). In this release, only one manual area is supported. For IS-IS to operate, you must configure at least one area.<br><br>Use the no option to delete the manual area. |

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

| Variable | Value                                                                            |
|----------|----------------------------------------------------------------------------------|
| <1-100>  | Creates the SPBM instance. In this release, only one SPBM instance is supported. |

*Table continues...*

| Variable                           | Value                                                                                                                                                                    |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| b-vid {<vlan-id [-vlan-id] [,...]} | Sets the ISIS SPBM instance data VLANs.<br>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.<br>The value is 2.5 bytes in the format <x.xx.xx>. Use the no or default options to delete the configured nickname. |
| primary <1–4094>                   | Sets the IS-IS instance primary data VLAN.                                                                                                                               |

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

| Variable                                                                                                                                                                                                                                                                                                                                       | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <2–4094>                                                                                                                                                                                                                                                                                                                                       | Specifies the VLAN ID. Creates an SPBM Backbone VLAN (B-VLAN). You can optionally specify a name for the SPBM B-VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| type {port protocol-decEther2 protocol-decOtherEther2 protocol-ipEther2 protocol-ipv6Ether2 protocol-ipx802.2 protocol-ipx802.3 protocol-ipxEther2 protocol-ipxSnap protocol-Netbios protocol-RarpEther2 protocol-sna802.2 protocol-snaEther2 protocol-Userdef protocol-vinesEther2 protocol-xnsEther2 spbm-bvlan spbm-switchedUni voice-vlan} | Specifies the type of VLAN created. <ul style="list-style-type: none"> <li>• port — port-based</li> <li>• protocol-decEther2 — Specify a decEther2 protocol-based VLAN.</li> <li>• protocol-decEther2 — Specify a OtherdecEther2 protocol-based VLAN.</li> <li>• protocol-ipEther2 — Specify an ipEther2 protocol-based VLAN.</li> <li>• protocol-ipv6Ether2 — Specify an ipv6Ether2 protocol-based VLAN.</li> <li>• protocol-ipx802.2 — Specify an ipx802.2 protocol-based VLAN.</li> <li>• protocol-ipx802.3 — Specify an ipx802.3 protocol-based VLAN.</li> <li>• protocol-ipxEther2 — Specify an ipxEther2 protocol-based VLAN.</li> <li>• protocol-ipxSnap — Specify an ipxSnap protocol-based VLAN.</li> <li>• protocol-Netbios — Specify a NetBIOS protocol-based VLAN.</li> <li>• protocol-RarpEther2 — Specify a RarpEther2 protocol-based VLAN.</li> <li>• protocol-sna802.2 — Specify a sna802.2 VLAN.</li> <li>• protocol-snaEther2 — Specify an snaEther2 protocol-based VLAN.</li> </ul> |

*Table continues...*

| Variable | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | <ul style="list-style-type: none"> <li>• protocol-Userdef — Specify a user-defined protocol-based VLAN. Enter optional parameters.                             <ul style="list-style-type: none"> <li>- all – display all Userdef VLANs</li> <li>- ether – display Ethernet II Userdef VLANs</li> <li>- llc – display LLC Userdef VLANs</li> </ul> </li> <li>• protocol-vinesEther2 — Specify a vinesEther2 protocol-based VLAN.</li> <li>• protocol-xnsEther2 — Specify an xnsEther2 protocol-based VLAN.</li> <li>• spbm-bvlan — Specify an SPBM-BVLAN.</li> <li>• spbm-switchedUni — Specify an SPBM-switchedUni</li> <li>• voice-vlan — Specify voice VLAN information</li> </ul> |

## Job aid

### Important:

After you configure the SPBM nickname and enable 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 might not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps:

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

---

## Displaying global SPBM parameters

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

### Procedure

1. Log on to CLI to enter User EXEC mode.
2. At the command prompt, enter the following command to check if SPBM is enabled:

```
show spbm
```

- At the command prompt, enter the following command:

```
show isis spbm
```

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

```
show vlan
```

### Example

```
Switch(config)#show spbm
SPBM Global: Disabled
SPBM Ethertype: 0x8100
```

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

```
=====
 ISIS SPBM Info
=====
SPBM B-VID PRIMARY NICK LSDB IP MULTICAST
INSTANCE VLAN VLAN NAME TRAP
=====
4 400-401 400 1.10.01 disable enable enable
```

```
Switch# show vlan
```

```
Id Name Type Protocol PID Active IVL/SVL Mgmt

1 VLAN #1 Port None 0x0000 Yes IVL Yes
 Port Members: ALL
2 VLAN #2 Port None 0x0000 Yes IVL No
 Port Members: 15-18
3 VLAN #3 Port None 0x0000 Yes IVL No
 Port Members: NONE
4 VLAN #4 B-VLAN None 0x0000 Yes IVL No
 Port Members: NONE
```

## Variable definitions

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

| Parameter      | Description                               |
|----------------|-------------------------------------------|
| SPBM Global    | Indicates if SPBM is enabled or disabled. |
| SPBM Ethertype | Indicates the SPB EtherType value.        |

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

| Parameter     | Description                                                                                            |
|---------------|--------------------------------------------------------------------------------------------------------|
| SPBM INSTANCE | Indicates the SPBM instance identifier. You can only create one SPBM instance.                         |
| B-VID         | Indicates the SPBM B-VLAN associated with the SPBM instance.                                           |
| PRIMARY VLAN  | Indicates the primary SPBM B-VLAN.                                                                     |
| NICK NAME     | Indicates the SPBM node nickname. The nickname is used to calculate the I-SID multicast MAC address.   |
| LSDB TRAP     | Indicates the status of the IS-IS SPBM LSDB update trap on this SPBM instance. The default is disable. |

*Table continues...*

| Parameter | Description                              |
|-----------|------------------------------------------|
| IP        | Indicates the state of SPBM IP Shortcut. |
| MULTICAST | Indicates the state of multicast.        |

## Displaying global IS-IS parameters

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

### Procedure

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

```
show isis
```
3. Display the IS-IS system-id:  

```
show isis system-id
```
4. Display IS-IS net info:  

```
show isis net
```

### Example

```
Switch#show isis
=====
 ISIS General Info
=====
AdminState : enabled
RouterType : Level 1
System ID : 0000.0010.1010
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : true
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
Spf-delay : 100
Router Name : A1
ip-source-address : 192.0.2.1
Num of Interfaces : 2
Num of Area Addresses : 1
Switch(config)#
```

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

```
Switch#show isis net
=====
 ISIS Network Entity Title Info
=====
NET
```



```
=====
c0.2000.0000.0000.14c7.e133.df00
```

## Variable definitions

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. |
| Min LSP Gen Interval  | Indicates the minimum time between LSP updates in seconds. |
| Metric                | Indicates if the metric is narrow or wide.                 |
| Overload-on-startup   | Indicates the overload-on-startup value.                   |
| Overload              | Indicates if there is an overload condition.               |
| Csnp Interval         | Indicates the interval between CSNP updates in seconds.    |
| PSNP Interval         | Indicates the interval between PSNP updates in seconds.    |
| Rxmt LSP Interval     | Indicates the received LSP time interval.                  |
| spf-delay             | Indicates the Shortest Path First delay in milliseconds.   |
| Router Name           | Indicates the IS-IS name of the router.                    |
| ip-source-address     | Indicates the source IP address.                           |
| Num of Interfaces     | Indicates the number of interfaces on the router.          |
| Num of Area Addresses | Indicates the number of area addresses on the router.      |

### show isis system-id

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

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

### show isis net

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

| Parameter | Description                                                                                                                                       |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| NET       | Shows the NET address. Output from this command is from the global IS-IS configuration of the manual area and the configuration of the system ID. |

| Parameter | Description                                                                                                                                     |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------|
|           | There is only one manual area 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. These I-SIDs are reserved for multicast and cannot be manually configured on a C-VLAN.

### Before you begin

- You must configure a loopback port. For more information, see [Configuring the loopback port](#) on page 166.

### \* Note:

Configuring the loopback port requires a reset.

- 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).
- You must add IST to the C-VLAN for an SMLT topology.

### \* Note:

You should use the `spbm reserved-port {front-panel | stack}` command to enable IP Multicast over Fabric Connect and to configure the loopback port simultaneously to avoid multiple device reboots.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Verify no I-SIDs exist in the default reserved range:

For Layer 2 use the following command:

```
show vlan i-sid
```

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:**

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

The device must be reset in order for the configuration change to become effective.

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(config)#show vlan i-sid
=====
Vlan I-SID
=====
VLAN_ID I-SID

1
50 200
51
52
53
54
55
56
57
9 out of 9 Total Num of Vlans displayed
switch:1>enable
switch:1#configure terminal
switch:1(config)#router isis
switch:1(config-isis)#spbm 1 multicast enable
```

**Variable definitions**

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

| Variable | Value                                                                                             |
|----------|---------------------------------------------------------------------------------------------------|
| <1-100>  | Enables IP Multicast over SPBM globally. The default is disabled.<br>Specifies the SPBM instance. |

## Displaying IP Multicast over Fabric Connect information

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

### Procedure

1. To enter User EXEC mode, log on to the switch.
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 with the corresponding Scope I-SID, Data I-SID, and the host name of the source:

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

### Example

The following example displays the output for the `show isis spbm multicast` command.

```
Switch:1>show isis spbm multicast
multicast : enable
fwd-cache-timeout : 210
```

The following example displays the output for the `show isis spb-mcast-summary` command.

```
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

80 192.0.2.1 203.0.113.3 16300014 1000 0x0 MERS2-8606
80 192.0.2.1 203.0.113.4 16300015 1000 0x0 MERS2-8606
80 192.0.2.3 203.0.113.3 16300001 1001 0x0 MERS4-8606
80 192.0.2.3 203.0.113.4 16300002 1001 0x0 MERS4-8606
200 192.0.2.4 203.0.113.2 16000001 1000 0x1 4826GTS
80 192.0.2.5 203.0.113.2 16000003 1000 0x1 4826GTS

6 out of 6 Total Num of Entries
```

## Variable definitions

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

| Variable                               | Value                                                                                                                                 |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| host-name <i>WORD&lt;0-255&gt;</i>     | Displays the IP Multicast over SPBM summary information for a specific host-name.                                                     |
| lspid <i>&lt;xxx.xxx.xxx.xx-xx&gt;</i> | Displays the IP Multicast over SPBM summary information for the specified LSP ID that you enter in xxx.xxx.xxx.xx-xx — 8 byte format. |

## Displaying IS-IS areas

Use the following procedure to display IS-IS areas.

### Procedure

1. Log on to CLI to enter User EXEC mode.
2. At the command prompt, enter the following command:

```
show isis manual-area
```

### Example

```
Switch#show isis manual-area
```

```
=====
 ISIS Manual Area Address
=====
AREA ADDRESS

c0.2000.0000.00
```

## Variable definitions

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. The manual area can be from 1-13 bytes in length. |

## 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 stack operation mode:

```
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. Log on to the IS-IS Router Configuration mode:  

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

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

To disable LSDB traps, use the following command:

```
no 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 Ethernet <port>
```
  - 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 ptpt
```
  - d. Configure the SPBM instance level 1 metric on the IS-IS interface:  

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

To set the l1-metric to the default value of 10, use one of the following commands:

```
no isis spbm <1-100> l1-metric
```

OR

```
default isis spbm <1-100> l1-metric
```
  - 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 ethernet 3
Switch(config-if)# no isis enable
Switch(config-if)# isis spbm 1 interface-type ptpt
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} | <p>Specifies the global Ethertype value as 0x8100 or 0x88a8. The default value is 0x8100.</p> <p>This value allows SPB to be transported across non-SPB networks, that is, transparent VLAN service or a traditional Ethernet network. For SPB interoperability between different vendors, you must change this value to the STP standard EtherType value of 0x88a8 unless this vendor also supports an SPB EtherType value of 0x8100.</p> |
| <1-100> lsdb-trap enable    | <p>Configures whether to enable or disable a trap when the SPBM LSDB changes.</p> <p>The default is disabled. Use the no or default options to disable LSDB traps.</p>                                                                                                                                                                                                                                                                     |

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

| Variable                       | Value                                                                                                                                                                                                                                                                                         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <1-100> interface-type ptpt    | <p>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.</p> <p>The default is pt-pt. Use the no or default options to set this parameter to the default value of pt-pt.</p> |
| <1-100> l1-metric <1-16777215> | <p>Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10.</p> <p>Use the no or default options to set this parameter to the default.</p>                                                                                |

---

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

```
max-lsp-gen-interval <30-900>
```
  - d. Configure the IS-IS metric type:

```
metric {wide}
```
  - e. Set or clear the overload condition:

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

```
overload-on-startup <15-3600>
```
  - g. Configure the Partial Sequence Number Packet (PSNP) in seconds:

```
psnp-interval <1-120>
```
  - h. Configure the minimum time between retransmission of an LSP:

```
retransmit-lsp-interval <1-300>
```
  - i. Configure the SPF delay in milliseconds:

```
spf-delay <0-5000>
```
  - j. Configure the name for the system:

```
sys-name WORD <1-255>
```
  - 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)# snp-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 20
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.<br><br>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                                                                                                                                                                                                               |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| {1}      | Sets the router type globally: <ul style="list-style-type: none"> <li>• l1: Level-1 router type</li> </ul> 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.<br><br>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                                                                    |
|----------|--------------------------------------------------------------------------|
| {wide}   | Specifies the IS-IS metric type. Only wide is supported in this release. |

| Variable | Value                                                                                                        |
|----------|--------------------------------------------------------------------------------------------------------------|
|          | 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                                                                                                                                                              |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>overload</i> | Sets or clears the overload condition.<br><br>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                                                              |
|----------------------------|--------------------------------------------------------------------|
| <i>overload-on-startup</i> | Specifies the overload-on-startup value.<br><br>The default is 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.<br><br>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.<br><br>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 |

| Variable | Value                                                                                                                                                                                                                                     |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | <p>scheduled back-to-back. The mechanism for pacing SPF allows two back-to-back SPF runs.</p> <p>The default value is 100 milliseconds. Use the no or default options to set this parameter to the default value of 100 milliseconds.</p> |

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

| Variable            | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>WORD</i> <1–255> | <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 no or default options to set this parameter to the default value (host name).</p> <p> <b>Note:</b></p> <p>In this release, no consistency checks appear when you edit sys-name on the switch.</p> |

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

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

## Job aid

### Important:

After you configure the SPBM nickname and enable 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 might not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps:

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

---

## Configuring optional IS-IS interface parameters

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

### Procedure

1. Enter Ethernet Interface Configuration mode:

```
enable
configure terminal
interface Ethernet <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> enable
Switch# configure terminal
Switch(config)# interface ethernet 3
Switch(config-if)# isis
Switch(config-if)# isis hello-auth type hmac-md5 key test
Switch(config-if)# isis l1-dr-priority 100
```

```
Switch(config-if)# isis l1-hello-interval 20
Switch(config-if)# isis l1-hello-multiplier 10
```

## Variable definitions

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

| Variable                                                                        | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hello-auth type {none simple hmac-md5} [key [key WORD<1–16> ] [key-id <1–255> ] | <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"> <li>• none</li> <li>• 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.</li> <li>• 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.</li> </ul> <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> <b>Note:</b></p> <p>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.</p>                                                                                                                                                                                                                                                                                                                                        |
| l1-hello-interval <1–600>                                                       | <p>Configures the level 1 hello interval. The default value is 9 seconds.</p> <p>Use the no or default options to set this parameter to the default value of 9 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| l1-hello-multiplier <1–600>                                                     | <p>Configures the level 1 hello multiplier. The default value is 3 seconds.</p> <p>Use the no or default options to set this parameter to the default value of 3 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Displaying IS-IS interface parameters

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

### Procedure

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

```
show isis interface [11]
```

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

The following example displays the output of the **show isis interface** command.

```
Switch>show isis interface
=====
 ISIS Interfaces
=====
IFIDX TYPE LEVEL OP-STATE ADM-STATE ADJ UP-ADJ SPBM-L1-METRIC

Trunk: 2 pt-pt Level 1 UP UP 1 1 10
Port: 21 pt-pt Level 1 UP UP 1 1 10
=====
```

The following example displays the output of the **show isis int-auth** command.

```
Switch>show isis int-auth
=====
 ISIS Interface Auth
=====
IFIDX AUTH-TYPE AUTH-KEYID AUTH-KEY

Trunk: 3 none 0
Port: 21 none 0
=====
```

The following example displays the output of the **show isis int-timers** command.

```
Switch>show isis int-timers
=====
 ISIS Interface Timers
=====
IFIDX LEVEL HELLO HELLO HELLO
 LEVEL INTERVAL MULTIPLIER DR

Trunk: 2 Level 1 9 3 3
Port: 21 Level 1 9 3 3
=====
```

The following example displays the output of the `show isis int-ckt-level` command.

```
Switch>show isis int-ckt-level
=====
ISIS Circuit level parameters
=====
IFIDX LEVEL DIS CKTID

Trunk: 2 Level 1
Port: 21 Level 1 1
 2
```

## Variable definitions

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

| Variable | Value                                                           |
|----------|-----------------------------------------------------------------|
| [1]      | Displays the interface information for the specified level: l1. |

## Displaying the 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 allows for Global Routing Table (GRT) IP networks to be transported across IS-IS.

### Procedure

1. Log on to CLI to enter User EXEC mode.
2. Display the SPBM multicast FIB:

```
show isis spbm multicast-fib [vlan <0-4094>] [i-sid <1-16777215>]
[nick-name <x.xx.xx>] [summary]
```

3. Display the SPBM unicast FIB:

```
show isis spbm unicast-fib [b-mac <0x00:0x00:0x00:0x00:0x00:0x00>]
[vlan <0-4094>] [summary]
```

4. Display the SPBM unicast tree:

```
show isis spbm unicast-tree <1-4094> [destination <xxxx.xxxx.xxxx>]
```

**Example**

The following example displays the output for the **show isis spbm multicast-fib** command.

```
Switch>show isis spbm multicast-fib
=====
 SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA ISID BVLAN SYSID HOST-NAME OUTGOING-INTERFACES

13:11:16:00:00:c8 200 1000 0014.c7e1.33df SPBM-1 MLT-2,3/21,3/37
13:11:16:00:01:2c 300 1000 0014.c7e1.33df SPBM-1 MLT-2,4/21
13:11:16:00:01:90 400 1000 0014.c7e1.33df SPBM-1 MLT-2,3/21
13:11:16:00:00:c8 200 2000 0014.c7e1.33df SPBM-1 MLT-2,3/21,3/31,3/37

Total number of SPBM MULTICAST FIB entries 4
=====
```

The following example displays the output for the **show isis spbm unicast-fib** command.

```
Switch>show isis spbm unicast-fib
=====
 SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION BVLAN SYSID HOST-NAME OUTGOING COST
ADDRESS INTERFACE

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

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

The following example displays the output for the **show isis spbm unicast-tree** command.

```
Switch>show isis spbm unicast-tree 1000
Node:0018.b0bb.b3df.00 (Switch) -> ROOT
Node:0016.ca23.73df.00 (Switch) -> ROOT
```

**Variable definitions**

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

| Variable            | Value                                         |
|---------------------|-----------------------------------------------|
| vlan <0-4094>       | Displays the FIB for the specified SPBM VLAN. |
| i-sid <1-1677215>   | 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 BMAC.      |
| vlan <0–4094>                         | 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                                                    |
|------------------------------|----------------------------------------------------------|
| <1–4094>                     | 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 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.                         |
| OUTGOING INTERFACES | Indicates the outgoing interface of the multicast FIB entry.                |

### show isis spbm unicast-fib

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

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

## Displaying IS-IS LSDB and adjacencies

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

### Procedure

1. Log on to CLI to enter User EXEC mode.

2. 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]
```

3. Display IS-IS adjacencies:

```
show isis adjacencies
```

4. Clear IS-IS LSDB:

```
clear isis lsdb
```

### Example

The following example displays the output of the **show isis lsdb** command.

```
Switch>show isis lsdb
=====
 ISIS LSDB
=====
LSP ID LEVEL LIFETIME SEQNUM CHKSUM HOST-NAME

0014.c7e1.33df.00-00 1 545 0xb1 0xed28 NewYork
0016.ca23.73df.00-00 1 1119 0x9f 0x9c9d VSP-
Lab2
0018.b0bb.b3df.00-00 1 708 0xb9 0xcb1a VSP-Lab1

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

The following example displays the output of the **show isis lsdb** command using the **detail** modifier.

```
Switch>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 (ERS0) Metric:10
 SPBM Sub TLV:
 port id: 640 num_port 1
 Metric: 10
0000.beb1.00b1.01 (VSP1) Metric:10
 SPBM Sub TLV:
 port id: 643 num_port 1
 Metric: 10
0000.bcb1.0004.01 (C1) Metric:10
 SPBM Sub TLV:
 port id: 6144 num_port 1
 Metric: 10
0000.beb1.00ca.01 (VSP2) Metric:10
 SPBM Sub TLV:
 port id: 6156 num_port 1
 Metric: 10
0000.beb1.00a5.01 (VSS0) Metric:10
 SPBM Sub TLV:
 port id: 651 num_port 1
 Metric: 10
0000.beb1.00b2.01 (VSS1) Metric:10
 SPBM Sub TLV:
 port id: 645 num_port 1
 Metric: 10
0000.beb1.0008.01 (VSP1) Metric:10
 SPBM Sub TLV:
 port id: 652 num_port 1
 Metric: 10

TLV:129 Protocol Supported: SPBM

TLV:137 Host_name: C0#

TLV:144 SUB-TLV 1 SPBM INSTANCE:
 Instance: 0
 bridge_pri: 0
 OUI: 00-33-33
 num of trees: 2
 vid tuple : u-bit 1 m-bit 1 ect-alg 0x80c201 base vid 1000
 vid tuple : u-bit 1 m-bit 1 ect-alg 0x80c202 base vid 1001
TLV:144 SUB-TLV 3 ISID:
 Instance: 0

```

```

Metric: 0
B-MAC: 00-00-bc-b1-00-03
BVID:1000
Number of ISID's:8
 3001 (Both), 3002 (Rx), 3003 (Both), 3004 (Rx), 4001 (Both), 4002 (
Rx), 4003 (Both), 4004 (Rx)

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

--More-- (q = quit)

```

The following example displays the output of the **show isis adjacencies** command.

```

Switch>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 ERS-Lab1
Port3/21 1 UP 1d 03:57:16 127 27 0016.ca23.73df ERS-Lab2

 2 out of 2 Total Num of Adjacencies
=====

```

## Variable definitions

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

| Variable                                  | Value                                                                                                                                                                                                       |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| level <i>{1 2 12}</i>                     | Displays the LSDB for the specified level: l1, l2, or l12.<br><br> <b>Note:</b><br>Level 1 is supported in this release. |
| sysid <i>&lt;xxxx.xxxx.xxxx&gt;</i>       | Displays the LSDB for the specified system ID.                                                                                                                                                              |
| lspid <i>&lt;xxxx.xxxx.xxxx.xx-xx&gt;</i> | Displays the LSDB for the specified LSP ID.                                                                                                                                                                 |
| tlv <i>&lt;1-184&gt;</i>                  | Displays the LSDB by TLV type.                                                                                                                                                                              |
| detail                                    | Displays detailed information.                                                                                                                                                                              |

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. |

## Displaying IS-IS statistics and counters

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

### Procedure

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

```
show isis statistics
```
3. Display IS-IS interface counters:  

```
show isis int-counters
```
4. Display IS-IS level 1 control packet counters:  

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

#### \* Note:

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

5. Clear IS-IS statistics:

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

### Example

The following example displays the output of the `show isis statistics` command.

```
Switch>show isis adjacencies
=====
 ISIS Adjacencies
=====
INTERFACE L STATE UPTIME PRI HOLDTIME SYSID HOST-NAME

Mlt2 1 UP 1d 03:57:25 127 20 0018.b0bb.b3df ERS-Lab1
Port3/21 1 UP 1d 03:57:16 127 27 0016.ca23.73df ERS-Lab2

 2 out of 2 Total Num of Adjacencies
=====
```

The following example displays the output of the `show isis int-counters` command.

```
Switch>show isis int-counters
=====
 ISIS Interface Counters
=====
IFIDX LEVEL AUTH ADJ INIT REJ ID LEN MAX AREA LAN DIS

 FAILS CHANGES FAILS ADJ CHANGES

Mlt2 Level 1 0 1 0 0 0 0 0
```

```
Port3/21 Level 1 0 1 0 0 0 0 0
```

The following example displays the output of the `show isis int-l1-ctrl-pkts` command.

```
Switch>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
Port3/21 Transmitted 13340 227 2 226
Port3/21 Received 13335 226 1 227
```

## Variable definitions

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

| Variable        | Value                                   |
|-----------------|-----------------------------------------|
| error-counters  | Clears IS-IS stats for error-counters.  |
| packet-counters | Clears IS-IS stats for packet-counters. |

## 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 MultiLink Trunking on Ethernet Routing Switch 4900 and 5900 Series](#).

### 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 <1-4094> <0-16777214>

OR

i-sid <0-16777214> vlan <1-4094>
```
3. **(Optional)** Verify the Private VLAN configuration:
 

```
show vlan private-vlan
```

## Example

The following example displays sample output for the `show vlan private-vlan` command.

```
Switch(config)#show vlan private-vlan
Primary VLAN Primary I-SID Secondary VLAN Secondary I-SID

2 200 10 200
3 300 30 300
Total Private VLANs: 2
```

## Variable definitions

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

| Variable     | Value                                                                                                                                                                                                                                     |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1–4094       | Primary VLAN ID.<br>Specifies the VLAN ID in the range of 1 to 4094. VLAN IDs 1 to 4094 are configurable. The system reserves VLAN IDs 4060 to 4094 for internal use. VLAN ID 1 is the default and you cannot create or delete VLAN ID 1. |
| <0-16777215> | Specifies the I-SID value. This value is same for primary and secondary VLANs.                                                                                                                                                            |

## Fabric Attach configuration using Command Line Interface

This section provides procedural information you can use to configure Fabric Attach (FA) using Command Line Interface (CLI).

### Displaying FA-specific settings

Use this procedure to display the FA configuration status.

#### Procedure

1. Log on to CLI to enter User EXEC mode.
2. To display the FA configuration status, enter the following command:

```
show fa agent
```

#### Example

This example shows sample output for the `show fa agent` command default settings.

```
Switch(config)#show fa agent
Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Zero Touch Status: Enabled
Fabric Attach Mgmt VLAN Distribution: Enabled
Fabric Attach Auto Provision Setting: Proxy
Fabric Attach Provision Mode: Disabled
Fabric Attach Client Proxy Status: Enabled
Fabric Attach Standalone Proxy Status: Disabled
Fabric Attach Agent Timeout: 240 seconds
```

```
Fabric Attach Extended Logging Status: Disabled
Fabric Attach Primary Server Id: <none>
Fabric Attach Primary Server Descr: <none>
```

This example shows sample output for the **show fa agent** command when management VLAN data in the FA Element TLV is disabled.

```
Switch(config)#show fa agent

Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Zero Touch Status: Enabled
Fabric Attach Mgmt VLAN Distribution: Disabled
Fabric Attach Auto Provision Setting: Proxy
Fabric Attach Provision Mode: Disabled
Fabric Attach Client Proxy Status: Enabled
Fabric Attach Standalone Proxy Status: Disabled
Fabric Attach Agent Timeout: 240 seconds
Fabric Attach Extended Logging Status: Disabled
Fabric Attach Primary Server Id: <none>
Fabric Attach Primary Server Descr: <none>
```

This example shows sample output for the **show fa agent** command in FA Proxy mode.

## Displaying Fabric Attach elements

Use this procedure to display discovered Fabric Attach elements.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. To display the discovered FA elements, enter the following:

```
show fa elements [<portlist> | trunk <trunknumber> | element-type
{server | proxy | client} | auth-status {auth-pass | auth-fail |
not-auth} | client-type <6-17>]
```

### Example

The following example displays sample output for the **show fa elements** command.

```
=====
Fabric Attach Discovered Elements
=====
UNIT/ MGMT ELEM ASGN
PORT TYPE VLAN STATE SYSTEM ID AUTH AUTH

1/5 Server 1234 T / S 6c:fa:58:dc:fc:00:00:00:01:05 AP AP
1/36 Client 1 U / D fc:a8:41:fa:f8:00:00:00:00:24 AP N
=====

Fabric Attach Authentication Detail
=====
UNIT/ ELEM OPER ASGN OPER
PORT EXPANDED TYPE AUTH STATUS AUTH STATUS

1/5 Server (Auth) successAuth successAuth
1/36 Switch successAuth none
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 number of Fabric Attach discovered elements displayed

```

| Field                 | Definition                                                     |
|-----------------------|----------------------------------------------------------------|
| State                 | FA Element TLV state field data                                |
| Elem Auth             | FA Element TLV authentication status                           |
| Asgn Auth             | FA I-SID/VLAN Assignment TLV authentication status             |
| Elem Oper Auth Status | FA Element TLV authentication status detail data               |
| Asgn Oper Auth Status | FA I-SID/VLAN Assignment TLV authentication status detail data |

## Variable Definitions

The following table describes the parameters for the `show fa elements` command.

| Variable                                       | Value                                                                            |
|------------------------------------------------|----------------------------------------------------------------------------------|
| <portlist>                                     | Specifies a port or a list of ports for which to display discovered FA elements. |
| trunk <trunknumber>                            | Specifies a trunk number for which to display discovered FA elements.            |
| auth-status {auth-pass   auth-fail   not-auth} | Displays only specified authorized status FA elements.                           |
| element-type {server   proxy   client}         | Displays only specified element type.                                            |
| client-type <6-17>                             | Displays only specified client type.                                             |

## Activating FA Server mode

Use the following procedure to activate FA Server mode and enable the FA service.

### Procedure

1. Enter Global Configuration mode:
 

```
enable
configure terminal
```
2. To activate FA Server mode and enable the FA service, enter the following command:
 

```
spbm
```
3. To activate FA Server mode in non-SPB environments, enter the following command:
 

```
fa auto-provision server
```

## Displaying I-SID-to-VLAN assignment information

Use this procedure to display information about I-SID-to-VLAN assignments.

**Procedure**

1. Log on to CLI to enter User EXEC mode.
2. To display I-SID-to-VLAN assignment information on an FA Proxy, enter the following commands:

```
show fa i-sid [<1-16777214>]
```

```
show i-sid [<1-16777214>]
```

OR

```
show fa assignment [<1-16777214>]
```

```
show i-sid [<1-16777214>]
```

3. To display I-SID-to-VLAN assignment information on an FA Server (SPBM enabled), enter the following command:

```
show fa i-sid [<1-16777214>]
```

```
show i-sid [<1-16777214>]
```

OR

```
show fa assignment [<1-16777214>]
```

```
show i-sid [<1-16777214>]
```

**Example**

The following example displays sample output for the `show fa i-sid` command.

```
Switch(config)#show fa i-sid
```

| I-SID    | VLAN | Source        | Status                                   |
|----------|------|---------------|------------------------------------------|
| 500      | 5    | Proxy         | Active                                   |
| 501      | 25   | Client        | Active                                   |
| 600      | 6    | Proxy, Client | Active                                   |
| 13849    | 138  | Proxy         | Rejected (VLAN invalid)                  |
| 16000000 | 1000 | Proxy         | Rejected (application interaction issue) |

**Creating an I-SID-to-VLAN assignment on an FA proxy****About this task**

Use this procedure to create an association between an I-SID and a VLAN on an FA Proxy, when SPBM is disabled on switch.

**Procedure**

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. To create an I-SID-to-VLAN assignment, enter the following command:

```
i-sid <1-16777214> vlan <1-4094>
```

## Result

Each FA Proxy I-SID-to-VLAN assignment creates a C-VLAN User Network Interface (UNI) when the assignment is active and accepted by an FA server.

## Example

The following example creates an association between I-SID 600 and VLAN 3:

```
Switch(config)#i-sid 600 vlan 3
Switch(config)#
```

## Variable definitions

The following table describes the parameters for the `i-sid <1-16777214> vlan <1-4094>` command

| Variable                              | Value                                                                                     |
|---------------------------------------|-------------------------------------------------------------------------------------------|
| <code>i-sid &lt;1-16777214&gt;</code> | Specifies the I-SID to associate with the selected VLAN. Values range from 1 to 16777214. |
| <code>vlan &lt;1-4094&gt;</code>      | Specifies the VLAN to associate with the selected I-SID. Values range from 1 to 4094.     |

## Deleting an I-SID-to-VLAN assignment on an FA Proxy

Use this procedure to remove the association between an I-SID and a VLAN on an FA Proxy.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. To remove a specific I-SID-to-VLAN assignment, enter the following command:

```
no i-sid <I-SID> vlan <VLAN>
```

3. To remove all configured I-SID-to-VLAN assignments, enter the following command:

```
default i-sid
```

## Variable definitions

The following table describes the parameters for the `no i-sid <I-SID> vlan <VLAN>` command

| Variable                              | Value                                                                                                    |
|---------------------------------------|----------------------------------------------------------------------------------------------------------|
| <code>i-sid &lt;1-16777214&gt;</code> | Specifies the I-SID of the specific I-SID-to-VLAN assignment to remove. Values range from 1 to 16777214. |
| <code>vlan &lt;1-4094&gt;</code>      | Specifies the VLAN of the specific I-SID-to-VLAN assignment to remove. Values range from 1 to 4094.      |

## Configuring external client proxy support

Use this procedure to enable or disable external client proxy support.

### Before you begin

Disable SPBM globally on switch.

### About this task

This operation enables or disables external client proxy support. It does not impact communication with an FA Server.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. To enable external client proxy support, enter either of the following commands:

```
fa proxy
```

### OR

```
default fa proxy
```

3. To disable external client proxy support, enter the following command:

```
no fa proxy
```

## Configuring FA on switch ports

Use this procedure to enable or disable the FA operation on one or more switch ports.

### \* Note:

Disabling FA at the port level causes FA Client data associated with the port, such as I-SID/VLAN binding data or discovered element data, to be immediately flushed. A link-down event also causes port-specific FA Client data to be cleared. I-SID/VLAN binding data that is associated with FA-enabled ports that become EAP-enabled is deleted as well.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. To configure the FA operation on switch ports, enter the following command:

```
[no][default] fa port-enable [<portlist>]
```

### Variable definitions

The following table describes the parameters for the `[no][default] fa port-enable [<portlist>]` command.

| Variable     | Value                                                                                                                                                        |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [<portlist>] | Enables the FA operation on the specified switch port or ports.<br>If you do not specify a port, the FA operation is enabled on all switch ports.            |
| [no]         | Disables the FA operation on the specified switch port or ports.<br>If you do not specify a port or ports, the FA operation is disabled on all switch ports. |
| [default]    | Restores the FA operation on all switch ports to default.                                                                                                    |

## Displaying switch port FA operation status

Use this procedure to display per-port FA operation status.

### Procedure

1. Log on to CLI to enter User EXEC mode.
2. To display FA configuration information, enter one of the following commands:

```
show fa port-enable [<portlist> | enabled-port | disabled-port |
enabled-auth | disabled-auth]
```

OR

```
show fa interface [<portlist> | enabled-port | disabled-port |
enabled-auth | disabled-auth]
```

### Example

The following example displays sample output for the `show fa port-enable` command.

```
Switch(config)#show fa port-enable
Unit Port IfIndex Trunk Advertisement Authentication Keymode

1 1 1 Enabled Enabled Strict
1 2 2 Enabled Enabled Strict
1 3 3 Enabled Enabled Strict
1 4 4 2 Enabled Enabled Standard
1 5 5 2 Enabled Enabled Standard
1 6 6 Enabled Enabled Strict
```

## Variable Definitions

The following table describes the parameters for the `show fa port-enable` or `show fa interface` command.

| Variable     | Value                                                                                                                                                                              |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <portlist>   | Specifies a port or a list of ports for which to display FA operation status. If you do not specify a port or ports, the switch displays FA operation status for all switch ports. |
| enabled-port | Displays only FA enabled ports.                                                                                                                                                    |

*Table continues...*

| Variable      | Value                                        |
|---------------|----------------------------------------------|
| disabled-port | Displays only FA disabled ports.             |
| enabled-auth  | Displays only authentication enabled ports.  |
| disabled-auth | Displays only authentication disabled ports. |

## Configuring the FA authentication key

Use the following command to configure the FA authentication key on specified ports.

### \* Note:

You can configure the FA authentication key only on secure images.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the FA authentication key:

```
[default] fa authentication-key <portlist>
```

Enter the authentication key, and then re-enter the key for confirmation. For security purposes, key data is hidden.

### Variable Definitions

The following table describes the parameters for the `fa authentication-key` command.

| Variable   | Value                                                                           |
|------------|---------------------------------------------------------------------------------|
| <portlist> | Specifies a port or a list of ports for which to define the authentication key. |

## Configuring FA message authentication support

Use the following procedure to configure the FA message authentication support on specified ports.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable the FA message authentication support:

```
fa message-authentication [<PortList>] [key-mode <strict |
standard>]
```

3. **(Optional)** Reset the FA message authentication support to default:

```
default fa message-authentication
```

**\* Note:**

The default setting is *enabled*.

4. **(Optional)** Disable the FA message authentication support:

```
no fa message-authentication [<PortList>]
```

## Variable Definitions

The following table describes the parameters for the `fa message-authentication` command.

| Variable                     | Value                                                                                                                                                                                                                                         |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <portlist>                   | Specifies a port or a list of ports for which to enable the FA message authentication support.                                                                                                                                                |
| key-mode <strict   standard> | Specifies the Authentication key usage setting — the user-defined authentication key (strict) or both the user-defined and default authentication keys (standard) are used for FA TLV data authentication.<br><br>Default key-mode is strict. |

## Configuring FA VLANs

Use this procedure to create or delete FA VLANs on an FA Proxy or FA Standalone Proxy.

### Before you begin

Disable SPBM globally on switch.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. To create FA VLANs, enter the following command:

```
fa vlan <LINE>
```

3. To delete FA VLANs, enter the following command:

```
no fa vlan <LINE>
```

4. To delete all configured FA VLANs, enter the following command:

```
default fa vlan
```

### Example

The following is an example of creating an FA VLAN and verifying the configuration.

```
Switch(config)#fa vlan 6
Switch(config)#show fa vlan
```

| VLAN | Source | Status  |
|------|--------|---------|
| 6    | Proxy  | Pending |

```
Binding Count: 1
```

## Variable Definitions

The following table describes the parameters for the `fa vlan` command.

| Variable | Value                                                                                                 |
|----------|-------------------------------------------------------------------------------------------------------|
| [<LINE>] | Specifies an individual VLAN ID or a range of VLAN IDs to create. A VLAN ID can range from 1 to 4094. |

## Displaying Fabric Attach VLAN information

Use this procedure to display Fabric Attach-specific VLAN information.

### Procedure

1. Enter Privileged EXEC mode:
 

```
enable
```
2. To display Fabric Attach VLAN information, enter the following command:
 

```
show fa vlan [<1-4094>]
```

### Example

The following example displays sample output for the `show fa vlan` command.

```
Switch(config)#show fa vlan
```

```

VLAN Source Status

1007 Proxy Pending
1008 Proxy Pending

```

## Disabling Management VLAN Distribution

Use this procedure to exclude management VLAN data in the FA Element TLV. When this option is not specified, management VLAN data in the FA Element TLV is included, by default.

### Procedure

1. Enter Global Configuration mode:
 

```
enable
```

```
configure terminal
```
2. Enter the following command to disable management VLAN distribution:
 

```
fa zero-touch disable-mgmt-vlan-distribution
```
3. Enter the following command to disable zero-touch operations, including management VLAN distribution:
 

```
no fa zero-touch
```
4. Enter the following command to enable management VLAN distribution:



```
default fa zero-touch
```

### Example

```
Switch enable
Switch config term
Switch (config)# fa zero-touch disable-mgmt-vlan-distribution
```

```
Switch (config)# show fa agent

Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Zero Touch Status: Enabled
Fabric Attach Mgmt VLAN Distribution: Disabled
Fabric Attach Auto Provision Setting: Proxy
Fabric Attach Provision Mode: Disabled
Fabric Attach Client Proxy Status: Enabled
Fabric Attach Standalone Proxy Status: Disabled
Fabric Attach Agent Timeout: 240 seconds
Fabric Attach Extended Logging Status: Disabled
Fabric Attach Primary Server Id: <none>
Fabric Attach Primary Server Descr: <none>
```

## Enabling or disabling FA Zero Touch support

Use this procedure to enable or disable the global FA Zero Touch support on an FA Proxy, FA Server, or FA Standalone Proxy. By default, FA Zero Touch support is enabled.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. To enable FA Zero Touch support on an FA Proxy, FA Server, or FA Standalone Proxy, enter the following command:

```
fa zero-touch
```

3. To disable FA Zero Touch support on an FA Proxy, FA Server, or FA Standalone Proxy, enter the following command:

```
no fa zero-touch
```

4. To reset the FA Zero Touch support state to default, enter the following command:

```
default fa zero-touch
```

## Configuring FA Zero Touch Client

Use the following procedure to manipulate Fabric Attach Zero Touch Client (ZTC) specifications on a FA Proxy or FA Server.

### \* Note:

The `auto-client-attach` option must be enabled before Zero Touch Client specifications can be applied (either during discovery or retroactively).

## Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. To enable an FA Zero Touch Client, enter the following command:

```
fa zero-touch-client standard {camera | ona-sdn | ona-spb-over-ip |
phone | router | security-dev | srvr-endpt | switch | video |
virtual-switch | wap-type1 | wap-type2} vlan <VLAN> [i-sid <ISID>]
[priority <Priority>][keep-static]
```

3. To delete a specific FA Zero Touch Client, enter the following command:

```
no fa zero-touch-client standard <ClientName>
```

4. To clear all FA Zero Touch Client settings, enter the following command:

```
default fa zero-touch-client
```

## Variable definitions

The following table describes the parameters for the `fa-zero-touch-client` command.

| Variable          | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| standard          | Specifies the Standard (pre-defined) client type. The following client types are available: <ul style="list-style-type: none"> <li>• 6 - Wireless AP (Type 1)</li> <li>• 7 - Wireless AP (Type 2)</li> <li>• 8 - Switch</li> <li>• 9 - Router</li> <li>• 10 - IP Phone</li> <li>• 11 - IP Camera</li> <li>• 12 - IP Video</li> <li>• 13 - Security Device</li> <li>• 14 - Virtual Switch</li> <li>• 15 - Sever Endpoint</li> <li>• 16 - ONA (SDN)</li> <li>• 17 - ONA (SpbOlp)</li> </ul> |
| vlan ID <1–4094>  | Specifies the VLAN ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| ISID <0–16777214> | Specifies the Client I-SID for I-SID/VLAN binding generation.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| priority <0–7>    | Specifies the Client port priority.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| keep-static       | Specifies whether static VLANs should be kept or removed on the client port for the duration of the client connection.                                                                                                                                                                                                                                                                                                                                                                    |

## Displaying FA Zero Touch Client

Use the following procedure to display Fabric Attach Zero Touch Client (ZTC) specifications on a FA Proxy or FA Server.

### Procedure

1. Enter Privileged EXEC mode:  
enable
2. Enter the following command:  
show fa zero-touch-client

### Example

The following example displays sample output for the show fa zero-touch-client.

```
Switch(config)#show fa zero-touch-client

Fabric Attach Zero Touch Client Auto-Attach Specifications

Type Client Name VLAN I-SID Priority Static

6 wap-type1 123 11111 NA remove
11 camera 200 2000 5 remove
17 ona-spb-over-ip 4001 40001 7 keep

Zero Touch Client Auto-Attach Specification count: 3
Switch(config)#
```

## Configuring FA Zero Touch options

Use this procedure to configure FA Zero Touch option settings.

### Procedure

1. Enter Global Configuration mode:  
enable  
configure terminal
2. To enable an FA Zero Touch option, enter the following command:  
fa zero-touch-options {{auto-client-attach | auto-mgmt-vlan-fa-client | auto-port-mode-fa-client | auto-pvid-mode-fa-client | auto-trusted-mode-fa-client |}} [client-type {hint | <6-17>}] | ip-addr-dhcp}

#### \* Note:

The auto-client-attach option must be enabled before Zero Touch Client specifications can be applied (either during discovery or retroactively).

**\* Note:**

The `auto-port-mode-fa-client` option is incompatible with both the `auto-pvid-mode-fa-client` and `auto-client-attach` options.

**\* Note:**

The `auto-mgmt-vlan-fa-client` option is incompatible with the `auto-pvid-mode-fa-client` and the `auto-port-mode-fa-client` options, as well as with the Zero Touch Client (ZTC) auto-attach support.

- To disable a specific FA Zero Touch option, enter the following command:

```
no fa zero-touch-options {{auto-port-mode-fa-client | auto-mgmt-vlan-fa-client| auto-pvid-mode-fa-client | auto-trusted-mode-fa-client | auto-client-attach} | ip-addr-dhcp}
```

- To clear all FA Zero Touch option settings, enter the following command:

```
default fa zero-touch-options
```

**Example**

```
Switch(config)#fa zero-touch-options auto-mgmt-vlan-fa-client client-type 8,9
Switch(config)#show fa zero-touch-options
```

```
Fabric Attach Zero Touch Options:
auto-mgmt-vlan-fa-client
auto-port-mode-fa-client
auto-pvid-mode-fa-client
auto-trusted-mode
```

```
4850GTS-PWR+(config)#show fa zero-touch-options client-data
```

Zero Touch Client Data

| Type | Client Name     | Applicable Zero Touch Options    |
|------|-----------------|----------------------------------|
| 6    | wap-type1       | auto-port-mode                   |
| 7    | wap-type2       |                                  |
| 8    | switch          | auto-mgmt-vlan                   |
| 9    | router          | auto-mgmt-vlan auto-trusted-mode |
| 10   | phone           |                                  |
| 11   | camera          | auto-trusted-mode                |
| 12   | video           |                                  |
| 13   | security-dev    |                                  |
| 14   | virtual-switch  | auto-port-mode                   |
| 15   | svr-endpt       | auto-pvid-mode                   |
| 16   | ona-sdn         | auto-port-mode                   |
| 17   | ona-spb-over-ip |                                  |

**Variable Definitions**

The following table describes the parameters for the `fa zero-touch-options` command.

| Variable                              | Value                                          |
|---------------------------------------|------------------------------------------------|
| <code>auto-port-mode-fa-client</code> | Automates the configuration of EAP port modes. |
| <code>auto-pvid-mode-fa-client</code> | Automates client PVID/Mgmt VLAN updates.       |

*Table continues...*

| Variable                    | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| auto-trusted-mode-fa-client | Automates the FA Client connection default QoS treatment.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| auto-mgmt-vlan-fa-client    | Automates management VLAN updates.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| ip-addr-dhcp                | Automates DHCP IP address acquisition.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| auto-client-attach          | Automates client attach configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| client-type <6–17>          | <p>Specifies an FA client type or a list of FA client types. Following are the available client types:</p> <ul style="list-style-type: none"> <li>• 6—Wireless AP (Type 1)</li> <li>• 7—Wireless AP (Type 2)</li> <li>• 8—Switch</li> <li>• 9—Router</li> <li>• 10—IP Phone</li> <li>• 11—IP Camera</li> <li>• 12—IP Video</li> <li>• 13—Security Device</li> <li>• 14—Virtual Switch</li> <li>• 15—Server Endpoint</li> <li>• 16—ONA (SDN)</li> <li>• 17—ONA (SpbOlp)</li> </ul> |

**\* Note:**

Default FA client types WAP Type 1 (6) and Switch (8) are associated with the client type-specific Zero Touch options if no client-type data is provided with the CLI commands.

## Displaying FA Zero Touch option settings

Use this procedure to verify the FA Zero Touch option settings.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. To display the FA Zero Touch option settings, enter the following command:

```
show fa zero-touch-options [client-data]
```

### Example

The following is an example of configuring and displaying FA Zero Touch options.

```
Switch(config)#fa zero-touch-options auto-port-mode-fa-client client-type 6,14-16
Switch(config)#show fa zero-touch-options
```

```
Fabric Attach Zero Touch Options:
```

```
ip-addr-dhcp
auto-port-mode-fa-client
```

The following is an example of displaying client data.

```
Switch(config)#show fa zero-touch-options client-data

Zero Touch Client Data

Type Client Name Applicable Zero Touch Options

6 wap-type1 auto-port-mode
7 wap-type2
8 switch
9 router
10 phone
11 camera
12 video
13 security-dev
14 virtual-switch auto-port-mode
15 srvr-endpt auto-port-mode
16 ona-sdn auto-port-mode
17 ona-spb-over-ip

Type Client Description Origin

6 Wireless AP (Type 1) Standard
7 Wireless AP (Type 2) Standard
8 Switch Standard
9 Router Standard
10 IP Phone Standard
11 IP Camera Standard
12 IP Video Standard
13 Security Device Standard
14 Virtual Switch Standard
15 Server Endpoint Standard
16 ONA (SDN) Standard
17 ONA (SpbOIp) Standard
Zero Touch Client Data

Switch(config)#
```

### Variable Definitions

The following table describes the parameters for the `show fa zero-touch-options` command.

| Variable    | Value                 |
|-------------|-----------------------|
| client-data | Displays client data. |

### Displaying Fabric Attach elements

Use this procedure to display discovered Fabric Attach elements.

#### Procedure

1. Enter Privileged EXEC mode:  
enable
2. To display the discovered FA elements, enter the following:

```
show fa elements [<portlist> | trunk <trunknumber> | element-type
{server | proxy | client} | auth-status {auth-pass | auth-fail |
not-auth} | client-type <6-17>]
```

**Example**

The following example displays sample output for the `show fa elements` command.

```
=====
Fabric Attach Discovered Elements
=====
UNIT/ MGMT ELEM ASGN
PORT TYPE VLAN STATE SYSTEM ID AUTH AUTH

1/5 Server 1234 T / S 6c:fa:58:dc:fc:00:00:00:01:05 AP AP
1/36 Client 1 U / D fc:a8:41:fa:f8:00:00:00:00:24 AP N
=====
Fabric Attach Authentication Detail
=====
UNIT/ ELEM OPER ASGN OPER
PORT EXPANDED TYPE AUTH STATUS AUTH STATUS

1/5 Server (Auth) successAuth successAuth
1/36 Switch successAuth none
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 number of Fabric Attach discovered elements displayed

```

| Field                 | Definition                                                     |
|-----------------------|----------------------------------------------------------------|
| State                 | FA Element TLV state field data                                |
| Elem Auth             | FA Element TLV authentication status                           |
| Asgn Auth             | FA I-SID/VLAN Assignment TLV authentication status             |
| Elem Oper Auth Status | FA Element TLV authentication status detail data               |
| Asgn Oper Auth Status | FA I-SID/VLAN Assignment TLV authentication status detail data |

**Configuring FA Standalone Proxy mode**

Use this procedure to enable or disable the FA Standalone Proxy mode on the switch.

**Before you begin**

Disable SPBM globally on switch.

**Procedure**

1. Enter Global Configuration mode:
 

```
enable
configure terminal
```
2. To enable FA Standalone Proxy mode, enter the following command:

```
fa standalone-proxy
```

3. To disable FA Standalone Proxy mode, enter the following command:

```
no fa standalone-proxy
```

4. To restore the FA Standalone Proxy mode to default, enter the following command:

```
default fa standalone-proxy
```

## Displaying FA uplink values

Use this procedure to display FA static uplink values used in FA Standalone Proxy mode.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. To display FA static uplink values, enter the following command:

```
show fa uplink
```

### Example

The following example displays sample output for the `show fa uplink` command.

```
Switch(config)#show fa uplink
Fabric Attach Static Uplinks:
 port - 0
 trunk - 8 (dynamic MLT [LAG admin key 300] - active)
```

## Configuring the static uplink for FA Standalone Proxy mode

Use this procedure to specify a port or trunk to use as a static uplink associated with FA Standalone Proxy operation.

### Before you begin

Disable SPBM globally on switch.

### Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. To specify a port uplink or a trunk uplink to use in FA Standalone Proxy mode, enter the following command:

```
fa uplink {port <port> | trunk <trunkId>}
```

3. To clear static uplink data, enter the following command:

```
no fa uplink
```



## Variable Definitions

The following table describes the parameters for the `fa uplink` command.

| Variable  | Value                                             |
|-----------|---------------------------------------------------|
| <port>    | Specifies the port to use as a static uplink.     |
| <trunkId> | Specifies the trunk ID to use as a static uplink. |

## Configuring Fabric Attach extended-logging

Use the following procedure to configure Fabric Attach extended-logging.

### Procedure

1. Enter Global Configuration mode:
 

```
enable
configure terminal
```
2. Enable Fabric Attach extended-logging:
 

```
fa extended-logging
```
3. Disable Fabric Attach extended-logging:
 

```
no fa extended-logging
```

## Configuring the FA timeout

### Procedure

1. Enter Global Configuration mode:
 

```
enable
configure terminal
```
2. To configure the FA timeout, enter the following command:
 

```
fa timeout <45-480>
```
3. To reset the timeout to its default value, enter the following command:
 

```
default fa timeout
```

## Clearing FA statistics

Use the following procedure to clear FA summary and per-port statistics counters. You can clear global counters, counters for an individual port or range of ports, or all ports.

### Procedure

1. Enter Global Configuration mode:
 

```
enable
configure terminal
```

2. Clear FA statistics:

```
clear fa statistics [summary | <PortList>]
```

### Variable Definitions

The following table describes the parameters for the `clear fa statistics` command.

| Variable   | Value                                                            |
|------------|------------------------------------------------------------------|
| <portlist> | Specifies a port or a list of ports for which to clear counters. |

### Displaying FA statistics

Use the following procedure to display the FA summary and per-port statistics counters. You can display global counters, counters for an individual port or range of ports, or all ports. When no port data is specified, only data for ports that are FA-enabled is displayed.

#### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Display FA statistics:

```
show fa statistics [summary | <portlist>]
```

### Variable Definitions

The following table describes the parameters for the `show fa statistics` command.

| Variable   | Value                                                                         |
|------------|-------------------------------------------------------------------------------|
| <portlist> | Specifies a port or a list of ports for which to display statistics counters. |

---

## SPBM and IS-IS infrastructure configuration using EDM

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

---

### Configuring required SPBM and IS-IS parameters

Use the following procedure to configure the minimum required SPBM and IS-IS parameters to allow SPBM to operate on the switch. SPBM uses the Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol to provide a loop free Ethernet topology that creates a shortest path topology from every node to every other node in the network based on node MAC addresses.

## Procedure

1. From the **Globals** tab, select **enable** to enable SPBM globally, and click **Apply**.
2. Click the **SPBM** tab.
3. Click **Insert** to create an SPBM instance. In this release, only one SPBM instance is supported.
4. In the **Id** field, specify the SPBM instance ID.
5. In the **NodeNickName** field, specify the node nickname (valid value is 2.5 bytes in the format <x.xx.xx>)
6. Click **Insert**.
7. In the **Vlans** field, specify the IDs of the SPBM B-VLANs to add to the SPBM instance.
8. In the **PrimaryVlan** field, specify which of the SPBM B-VLANs specified in the previous step is the primary B-VLAN.
9. Click **Apply**.
10. In the navigation tree, select **Configuration > IS-IS > IS-IS**.
11. Click the **Manual Area** tab.
12. In the Manual Area tab, click **Insert** to add a manual area. In this release, only one manual area is supported.
13. Specify the Manual Area Address (valid value is 1–13 bytes in the format <xx.xxxx.xxxx...xxxx>).
14. Click **Insert**.
15. 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.

16. In the AdminState field, click **on**, and click **Apply**.
17. Under the IS-IS tab, click the **Interfaces** tab.
18. Click **Insert** to create an IS-IS circuit.
19. In the **IfIndex** field, specify the port or MLT on which to create the IS-IS interface.
20. Click **Insert**.

 **Note:**

By default, all ports are enabled in VLAN 1. You can remove the port for the IS-IS interface from VLAN 1 and disable Spanning Tree participation at end of this procedure.

21. Select the newly created IS-IS circuit entry, and click **SPBM**.

22. In the **Interfaces SPBM** tab, click **Insert**.
23. In the **SpbmId** field, specify a SPBM identifier.
24. In the **State** field, select **enable**.
25. Click **Insert** to enable the SPBM instance on the IS-IS circuit.
26. Under the IS-IS tab, click the **Interfaces** tab.
27. In the **AdminState** field for the IS-IS circuit entry, select **on** to enable the IS-IS circuit.
28. Click **Apply**.
29. From the navigation tree, select **Configuration > VLAN > VLANs**.
30. Click the **Basic** tab.
31. Select the row for VLAN#1, and double-click the **PortMembers** cell.
32. Click the **port number** you specified for the IS-IS interface to remove it from the default VLAN, and click **Ok**.
33. In the toolbar, click **Apply**.

 **Note:**

Ensure you remove the port specified for the IS-IS interface from all non-SPBM VLANs.

34. From the navigation tree, select **Configuration > VLAN > VLANs**.
35. Click the **Basic** tab.
36. Click **Insert**.
37. In the **Type** field, click **spbm-bvlan**.
38. Click **Insert** to create the primary B-VLAN.
39. Click **Insert**.
40. In the **Type** field, click **spbm-bvlan**.
41. Click **Insert** to create the secondary B-VLAN.
42. In the navigation tree, select **Configuration > IS-IS > SPBM**.

## SPBM Globals Tab 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 MultiLink Trunking on Ethernet Routing Switch 4900 and 5900 Series](#).

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

| Name                   | Description                                                                            |
|------------------------|----------------------------------------------------------------------------------------|
| <b>GlobalEnable</b>    | Enables or disables SPBM globally.                                                     |
| <b>GlobalEtherType</b> | 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                                                                                            |
|---------------------|--------------------------------------------------------------------------------------------------------|
| <b>Id</b>           | Specifies the SPBM instance ID. In this release, only one SPBM instance is supported.                  |
| <b>NodeNickName</b> | Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>. |
| <b>PrimaryVlan</b>  | Specifies the primary SPBM B-VLANs to add to the SPBM instance.                                        |
| <b>Vlans</b>        | Specifies the SPBM B-VLANs to add to the SPBM instance.                                                |
| <b>LsdbTrap</b>     | Enables or disables LSDB trap for the SPBM instance.                                                   |

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

| Name        | Description                                                                                                                                                      |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b> | Specifies the type of VLAN: <ul style="list-style-type: none"> <li>• byPort</li> <li>• byProtocolId</li> <li>• spbm-bvlan</li> <li>• spbm-switchedUni</li> </ul> |



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

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

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

| Name              | Description                                                                                                     |
|-------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>AdminState</b> | Specifies the global status of IS-IS on the switch: on or off. The default is off.                              |
| <b>LevelType</b>  | Sets the router type globally: <ul style="list-style-type: none"> <li>• level1 — Level-1 router type</li> </ul> |

*Table continues...*

| Name      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <ul style="list-style-type: none"> <li>level2 — Level-2 router type</li> <li>Level1and2 — Level-1 and Level-2 router type</li> </ul> <p> <b>Note:</b><br/>level2 and level1and2 are not supported in this release.</p>                                                                                                                                                                                                                                         |
| <b>ID</b> | <p>Specifies the system ID. Valid value is a 6-byte value in the format &lt;xxxx.xxxx.xxxx&gt;</p> <p> <b>Note:</b><br/>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 <b>ID</b> 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                                                                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IfIndex</b>    | 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. |
| <b>AdminState</b> | Specifies the administrative state of the circuit: on or off. The default is off.                                                                                                                                         |

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

| Name         | Description                                                  |
|--------------|--------------------------------------------------------------|
| <b>State</b> | Specifies whether the SPBM interface is enabled or disabled. |

## Job aid

### Important:

After you configure the SPBM nickname and enable 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 might not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps:

1. Disable IS-IS.
2. Change the system ID.
3. Change the nickname to a temporary one.

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

---

## Configuring IP Multicast over 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.

### 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).
- You must add IST to the C-VLAN for an SMLT topology.

### Procedure

1. Determine if any I-SIDs are within the default range reserved for multicast. From the navigation tree, expand the following folders: **Configuration > IS-IS > SPBM**.
2. Click the **I-SID** tab to determine if the I-SIDs are within the default range reserved for multicast.
3. From the navigation tree, expand the following folders: **Configuration > IS-IS > SPBM**.
4. Click the **SPBM** tab.
5. If you want to enable multicast on an SPBM instance that already exists, in the **Multicast** 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 **Multicast** box, select **enable** to enable IP Multicast over Fabric Connect globally.
8. Click **Insert**.
9. Click **Apply**.

## SPBM Tab Field Descriptions

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

| Name                        | Description                                                                                  |
|-----------------------------|----------------------------------------------------------------------------------------------|
| <b>Id</b>                   | Specifies the SPBM instance ID. In this release, only one SPBM instance is supported.        |
| <b>NodeNickName</b>         | Specifies a nickname for the SPBM instance globally.                                         |
| <b>PrimaryVlan</b>          | Specifies the primary SPBM B-VLAN to add to the SPBM instance.                               |
| <b>Vlans</b>                | Specifies the SPBM B-VLANs to add to the SPBM instance.                                      |
| <b>LsdbTrap</b>             | Specifies if the LSDB update trap is enabled on this SPBM instance. The default is disabled. |
| <b>IpShortcut</b>           | Specifies if SPBM IP Shortcuts is enabled. The default is disabled.                          |
| <b>Multicast</b>            | Specifies if IP multicast over SPBM is enabled. The default is disabled.                     |
| <b>McastFwdCacheTimeout</b> | 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 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).
- You must add IST to the C-VLAN for an SMLT topology.

### Procedure

1. From the navigation tree, expand the following folders: **Configuration > IS-IS > SPBM**.
2. Click the **SPBM** tab.
3. Select **enable** or **disable** in the **Multicast** column in the table.
4. Select **enable** or **disable** in the **LsdbTrap** column in the table.
5. Click **Apply**.



## Displaying the SPBM I-SID information

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

### Procedure

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

## I-SID Tab Field Descriptions

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

| Name                    | Description                                                                                                 |
|-------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>SysId</b>            | Indicates the system identifier.                                                                            |
| <b>Vlan</b>             | Indicates the B-VLAN where this I-SID was configured or discovered.                                         |
| <b>McastDestMacAddr</b> | Indicates the multicast destination MAC address based on the NickName and I-SID to build the Multicast-FIB. |
| <b>Isid</b>             | Indicates the IS-IS SPBM I-SID identifier.                                                                  |
| <b>NickName</b>         | Indicates the nickname of the node where this I-SID was configured or discovered.                           |
| <b>HostName</b>         | Indicates the host name listed in the LSP, or the system name if the host name is not configured.           |
| <b>Type</b>             | Indicates the SPBM I-SID type; either configured or discovered.                                             |

## Displaying Level 1 Area information

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

### Important:

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

### Procedure

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

- Click the **L1 Area** tab.

## L1 Area Tab Field Descriptions

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

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

## Enabling or disabling SPBM globally

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.

### Before you begin

Configure the loopback port.

### Procedure

- In the navigation tree, expand the following folders: **Configuration > IS-IS**.
- Click **SPBM**.
- Click the **Globals** tab.
- To enable or disable SPBM, select **enable** or **disable** option from the **GlobalEnable**.
- To configure the global ethertype value, select the desired option from the **GlobalEtherType**.
- To configure the next loopback status on a port, select the desired option from the **SpbmLoopbackPortNextState**.

**\* Note:**

The SpbmLoopbackPortNextState option is available only when SPBM is enabled.

- Click **Apply**.

## Globals Tab Field Descriptions

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

| Name                   | Description                                                                            |
|------------------------|----------------------------------------------------------------------------------------|
| <b>GlobalEnable</b>    | Enables or disables SPBM globally. The default is disabled.                            |
| <b>GlobalEtherType</b> | Specifies the global ethertype value as 0x8100 or 0x88a8. The default value is 0x8100. |

*Table continues...*

| Name                             | Description                                                                                                                                                                                                                                                                |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                  | Specifies the current loopback port setting (one of the stack ports or the last front panel port).                                                                                                                                                                         |
| <b>SpbmLoopbackPortNextState</b> | <p>Specifies the loopback port setting after the next reboot (stack ports or the last two front panel ports).</p> <p><b>* Note:</b><br/>This option is not available in ERS 4900 Series.</p> <p><b>* Note:</b><br/>This option is available only when SPBM is enabled.</p> |

## Configuring SPBM parameters

Use the following procedure to configure SPBM global parameters. SPBM uses the Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol to provide a loop free Ethernet topology that creates a shortest path topology from every node to every other node in the network based on node MAC addresses.

### Procedure

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

## SPBM Tab Field Descriptions

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

| Name                | Description                                                                                            |
|---------------------|--------------------------------------------------------------------------------------------------------|
| <b>Id</b>           | Specifies the SPBM instance ID. In this release, only one SPBM instance is supported.                  |
| <b>NodeNickName</b> | Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>. |
| <b>PrimaryVlan</b>  | Specifies the primary SPBM B-VLANs to add to the SPBM instance.                                        |
| <b>Vlans</b>        | Specifies the SPBM B-VLANs to add to the SPBM instance.                                                |
| <b>LsdbTrap</b>     | Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable.     |

*Table continues...*

| Name                        | Description                                                       |
|-----------------------------|-------------------------------------------------------------------|
| <b>IpShortcut</b>           | Indicates whether IP Shortcut is activated for the SPBM instance. |
| <b>Multicast</b>            | Indicates whether multicast is activated for the SPBM instance.   |
| <b>McastFwdCacheTimeout</b> | Specifies the forwarding cache timeout for the SPBM instance.     |

---

## Displaying SPBM nicknames

Use the following procedure to display SPBM nicknames.

### Procedure

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

## Nickname Tab Field Descriptions

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

| Name                  | Description                                                                                          |
|-----------------------|------------------------------------------------------------------------------------------------------|
| <b>Level</b>          | Indicates the level at which this LSP appears. The only possible value in the current release is L1. |
| <b>ID</b>             | Indicates the 8 byte LSP ID, consisting of the SystemID, Circuit ID, and Fragment Number.            |
| <b>LifetimeRemain</b> | Indicates the remaining lifetime in seconds for the LSP.                                             |
| <b>NickName</b>       | Indicates the nickname for the SPBM node.                                                            |
| <b>HostName</b>       | Indicates the hostname listed in the LSP, or the system name if the host name is not configured.     |

---

## Configuring interface SPBM parameters

Use the following procedure to configure SPBM interface parameters.

### Procedure

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

## Interface SPBM Tab Field Descriptions

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

| Name                | Description                                                                                                                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>        | Specifies an Index value for the SPBM interface.                                                                                                                                                     |
| <b>SpbmId</b>       | Specifies an ID value for the SPBM interface.                                                                                                                                                        |
| <b>State</b>        | Specifies whether the SPBM interface is enabled or disabled.                                                                                                                                         |
| <b>Type</b>         | Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT: ptpt or bcast. In this release, only the point-to-point (ptpt) interface type is supported. |
| <b>WideL1Metric</b> | Configures the SPBM instance I1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10.                                                                         |

## Configuring SPBM on an interface

Use the following procedure to configure SPBM on an interface.

### Procedure

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

## Interfaces SPBM Tab Field Descriptions

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

| Name                | Description                                                                                                                                                                                 |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>        | Specifies an Index value for the SPBM interface.                                                                                                                                            |
| <b>Id</b>           | Specifies the SPBM instance ID.                                                                                                                                                             |
| <b>State</b>        | Specifies whether the SPBM interface is enabled or disabled. The default is disabled.                                                                                                       |
| <b>Type</b>         | Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT. In this release, only the pt-pt interface type is supported. The default is pt-pt. |
| <b>WideL1Metric</b> | Configures the SPBM instance I1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10.                                                                |

## 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. 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. The 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. A unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

### Procedure

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

## Unicast FIB Tab Field Descriptions

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

| Name                      | Description                                                                 |
|---------------------------|-----------------------------------------------------------------------------|
| <b>SysId</b>              | Specifies the system ID of the node where the unicast FIB entry originated. |
| <b>Vlan</b>               | Specifies the VLAN of the unicast FIB entry.                                |
| <b>DestinationMacAddr</b> | Specifies the destination MAC Address of the unicast FIB entry.             |
| <b>OutgoingPort</b>       | Specifies the outgoing port of the unicast FIB entry.                       |
| <b>HostName</b>           | Specifies the host name of the node where unicast FIB entry originated.     |
| <b>Cost</b>               | 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. 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. The B-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. A unicast path now exists from every node to

every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

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

### Procedure

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

## Multicast FIB Tab Field Descriptions

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

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

## Displaying LSP summary information

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

### Procedure

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

## LSP Summary Tab Field Descriptions

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

| Name         | Description                                                                               |
|--------------|-------------------------------------------------------------------------------------------|
| <b>Level</b> | Specifies the level at which this LSP appears.                                            |
| <b>ID</b>    | Specifies the 8 byte LSP ID, consisting of the SystemID, Circuit ID, and Fragment Number. |

*Table continues...*

| Name                  | Description                                                                    |
|-----------------------|--------------------------------------------------------------------------------|
| <b>Seq</b>            | Specifies the sequence number for this LSP.                                    |
| <b>Checksum</b>       | Specifies the 16 bit Fletcher Checksum for this LSP.                           |
| <b>LifetimeRemain</b> | The remaining lifetime in seconds for this LSP.                                |
| <b>HostName</b>       | The hostname listed in LSP, or the system name if host name is not configured. |

## Displaying IS-IS adjacencies

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

### Procedure

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

## Adjacency Tab Field Descriptions

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

| Name                 | Description                                                                                                                                                                                                                                       |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>     | Specifies the IS-IS interface on which the adjacency is found.                                                                                                                                                                                    |
| <b>Level</b>         | Indicates the level of the IS-IS interface (Level 1 [default] or Level 2).                                                                                                                                                                        |
| <b>State</b>         | Specifies the state of the adjacency: <ul style="list-style-type: none"> <li>• down</li> <li>• initializing</li> <li>• up</li> <li>• failed</li> </ul>                                                                                            |
| <b>LastUpTime</b>    | Indicates when the adjacency most recently entered the state <b>up</b> , 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 <b>up</b> . |
| <b>NeighPriority</b> | Specifies the priority of the neighboring Intermediate System for becoming the Designated Intermediate System.                                                                                                                                    |

*Table continues...*



| Name              | Description                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HoldTimer</b>  | 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. |
| <b>NeighSysID</b> | Specifies the system ID of the neighboring Intermediate System.                                                                                      |
| <b>HostName</b>   | Specifies the host name listed in the LSP, or the system name if host name is not configured.                                                        |

## Configuring IS-IS globally

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

### Procedure

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

## Globals Tab Field Descriptions

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

| Name              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AdminState</b> | Specifies the global status of IS-IS on the switch: on or off. The default is off.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>LevelType</b>  | Sets the router type globally: <ul style="list-style-type: none"> <li>• level1 — Level-1 router type</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>ID</b>         | Specifies the IS-IS system ID for the switch. Valid value is a 6-byte value in the format <xxx.xxxx.xxxx>. <p><b>!</b> <b>Important:</b></p> <p>After you configure the SPBM nickname and enable 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 might not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps:</p> <ol style="list-style-type: none"> <li>1. Disable IS-IS.</li> <li>2. Change the system ID.</li> <li>3. Change the nickname to a temporary one.</li> </ol> |

*Table continues...*

| Name                       | Description                                                                                                                                                                                                                                    |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                            | <ol style="list-style-type: none"> <li>4. Enable IS-IS.</li> <li>5. Disable IS-IS.</li> <li>6. Change the nickname to the original nickname.</li> <li>7. Enable IS-IS.</li> </ol>                                                              |
| <b>MaxLSPGenInt</b>        | <p>Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate system. The value must be greater than any value configured for RxmtLsplnt.</p> <p>The default value is 900 seconds.</p>                            |
| <b>Csnplnt</b>             | <p>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> <p>The default value is 10.</p>                                       |
| <b>RxmtLsplnt</b>          | <p>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> <p>The default value is 5 seconds.</p> |
| <b>PSNPInterval</b>        | <p>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> <p>The default value is 2.</p>                                         |
| <b>SpfDelay</b>            | <p>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> <p>The default value is 100 milliseconds.</p>                               |
| <b>HostName</b>            | <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>              |
| <b>IpSourceAddressType</b> | <p>Specifies the IP source address type used in IP Shortcut. The only option in this release is ipv4.</p>                                                                                                                                      |
| <b>IpSourceAddress</b>     | <p>Specifies the CLIP interface as the source address for SPBM IP shortcuts.</p>                                                                                                                                                               |

## Configuring system level IS-IS parameters

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

## Procedure

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

## System Level Tab Field Descriptions

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

| Name                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>            | Specifies the level: I1 or I2.<br><br>In this release, only I1 is supported.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>State</b>            | 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 <b>SetOverload</b> . If the state is 'waiting' or 'overloaded', you originate LSPs with the Overload bit set. |
| <b>MinLSPGenInt</b>     | Specifies the minimum time between successive generation of LSPs with the same LSPID. This a system level parameter that applies to both L1 and L2 LSP generation.<br><br>The default value is 30 seconds.                                                                                                                                                                                                                                                                                      |
| <b>SetOverload</b>      | Indicates whether there is an overload condition.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>SetOverloadUntil</b> | Indicates the overload-on-startup value, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>MetricStyle</b>      | Specifies the IS-IS metric type. Available values are narrow, wide or both. Only wide is supported in this release.                                                                                                                                                                                                                                                                                                                                                                             |

## Configuring IS-IS interfaces

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


### Procedure

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

5. Click **Apply**.

## Interfaces Tab Field Descriptions

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

| Name              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>      | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>IfIndex</b>    | Specifies the interface on which the circuit is configured (port or MLT).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Type</b>       | Specifies the IS-IS circuit type. In this release, only the point-to-point (PtToPt) interface type is supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>AdminState</b> | Specifies the administrative state of the circuit: on or off.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>OperState</b>  | Specifies the operational state of the circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>AuthType</b>   | <p>Specifies the authentication type:</p> <ul style="list-style-type: none"> <li>• none</li> <li>• 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.</li> <li>• 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.</li> </ul> <p>The default is none.</p> |
| <b>AuthKey</b>    | Specifies the authentication key.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>KeyId</b>      | Specifies the authentication key ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>LevelType</b>  | <p>Sets the router type globally:</p> <ul style="list-style-type: none"> <li>• level1 — Level-1 router type</li> <li>• level2 — Level-2 router type</li> <li>• Level1and2 — Level-1 and Level-2 router type</li> </ul> <p> <b>Note:</b><br/>level2 and level1and2 is not supported in this release.</p>                                                                                                                                                                                                                                                                                                                         |
| <b>NumAdj</b>     | Specifies the number of adjacencies on this circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>NumUpAdj</b>   | Specifies the number of adjacencies that are up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Configuring IS-IS interface level parameters


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

### Procedure

1. From the navigation tree, choose **Configuration > IS-IS**.
2. Click **IS-IS**.
3. Click the **Interfaces Level** tab.
4. Configure the IS-IS interface level parameters.
5. Click **Apply**.

### Interfaces Level Tab Field Descriptions

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

| Name                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>           | Indicates the identifier of the circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and does not have any relation to any protocol value.                                                                                                                                                                                                                                                 |
| <b>LevelIndex</b>      | Specifies the router type globally: <ul style="list-style-type: none"> <li>• I1: Level1 router type</li> <li>• I12: Level1/Level2 router type. Not supported in this release.</li> </ul> The default value is I1.                                                                                                                                                                                                                   |
| <b>ISPriority</b>      | Specifies an integer sub-range for IS-IS priority. Range of 0–127. The default is 0 for SPBM interfaces. <p> <b>Note:</b><br/>ISPriority only applies to broadcast interfaces.</p>                                                                                                                                                               |
| <b>HelloTimer</b>      | Specifies the level 1 hello interval. <p>Specifies the maximum period, in seconds, 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.</p> The default value is 9 seconds. |
| <b>HelloMultiplier</b> | Specifies the level 1 hello multiplier. The default value is 3 seconds.                                                                                                                                                                                                                                                                                                                                                             |

*Table continues...*

| Name                | Description                                                                                                                                                                 |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DRHelloTimer</b> | Specifies the period, in seconds, between Hello PDUs on multiaccess networks when this Intermediate System is the Designated Intermediate System. The default is 3 seconds. |

---

## Configuring an IS-IS Manual Area

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

### Procedure

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

## Manual Area Tab Field Descriptions

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

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

---

## 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 tree, choose **Configuration > IS-IS**.
2. Click **Stats**.
3. Click the **System Stats** tab.

## System Stats Tab Field Descriptions

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

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

---

## Displaying IS-IS interface counters

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

### Procedure

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

## Interface Counters Tab Field Descriptions

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

| Name                         | Description                                                                                                                                                                     |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>                 | Shows a unique value identifying the IS-IS interface.                                                                                                                           |
| <b>CircuitType</b>           | Indicates the level type for the IS-IS interface. Only Level1 is supported in the current release.                                                                              |
| <b>AdjChanges</b>            | Shows the number of times an adjacency state change has occurred on this circuit.                                                                                               |
| <b>InitFails</b>             | 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. |
| <b>RejAdjs</b>               | Shows the number of times an adjacency has been rejected on this circuit.                                                                                                       |
| <b>IDFieldLenMismatches</b>  | Shows the number of times an IS-IS control PDU with an ID field length different to that for this system has been received.                                                     |
| <b>MaxAreaAddrMismatches</b> | 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.                                               |
| <b>AuthFails</b>             | Shows the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.                                                         |
| <b>LANDesiChanges</b>        | 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 tree, expand the following folders: **Configuration > IS-IS**.
2. Click **Stats**.
3. Click the **Interface Control Packets** tab.

## Interface Control Packets Tab Field Descriptions

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

| Name             | Description                                                                                        |
|------------------|----------------------------------------------------------------------------------------------------|
| <b>Index</b>     | Shows a unique value identifying the Intermediate-System-to-Intermediate-System (IS-IS) interface. |
| <b>Level</b>     | Indicates the level at which this LSP appears.                                                     |
| <b>Direction</b> | Indicates whether the switch is sending or receiving the PDUs.                                     |

*Table continues...*



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

## Fabric Attach configuration using Enterprise Device Manager

Use the procedures in this section to configure Fabric Attach (FA) using Enterprise Device Manager.

### Configuring Fabric Attach

Use this procedure to configure Fabric Attach.

#### Procedure

1. From the navigation tree, select **Edit > Fabric Attach**.
2. Click the **Agent** tab.
3. To set the Auto Provision mode to FA Proxy, click **proxy** in the **AutoProvision** field.
4. To enable or disable FA Standalone Proxy mode, click **enable** or **disable** in the **StandaloneProxy** field.
5. To enable or disable external client proxy support, click **enable** or **disable** in the **ClientProxy** field.
6. Specify the port to use as a static uplink associated with FA Standalone Proxy operation in the **UplinkPort** field.
7. Specify the trunk to use as a static uplink associated with FA Standalone Proxy operation in the **UplinkTrunk** field.
8. Specify the agent timeout in the **Timeout** field.
9. To enable or disable extended logging, click **enable** or **disable** in the **ExtendedLogging** field.
10. Click **Apply**.

#### Agent Tab Field Descriptions

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

| Name                   | Description                                                                                  |
|------------------------|----------------------------------------------------------------------------------------------|
| <b>Service</b>         | Displays the service status.                                                                 |
| <b>ElementType</b>     | Displays the element type.                                                                   |
| <b>ProvisionMode</b>   | Displays the provision mode status                                                           |
| <b>AutoProvision</b>   | Displays the Auto Provision mode.                                                            |
| <b>StandaloneProxy</b> | Specifies whether FA Standalone Proxy mode is enabled or disabled. The default is disabled.  |
| <b>ClientProxy</b>     | Specifies whether external client proxy is enabled or disabled. The default is enabled.      |
| <b>UplinkPort</b>      | Specifies the port to use as a static uplink associated with FA Standalone Proxy operation.  |
| <b>UplinkTrunk</b>     | Specifies the trunk to use as a static uplink associated with FA Standalone Proxy operation. |
| <b>Timeout</b>         | Specifies the agent timeout in seconds. The default value is 240 seconds.                    |
| <b>ExtendedLogging</b> | Specifies whether extended logging is enabled or disabled. The default is disabled.          |

## Configuring an I-SID/VLAN assignment

Use the following procedure to configure an I-SID/VLAN assignment on an FA Proxy.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Fabric Attach**.
3. In the work area, click the **I-SID** tab.
4. Click **Insert**.
5. Specify an I-SID in the **Isid** field.
6. Specify a VLAN in the **Vlan** field.
7. Click **Insert**.

## Configuring per-port FA settings

Use the following procedure to enable or disable FA Signaling or to configure FA message authentication.

### Procedure

1. From the navigation tree, select **Edit**.
2. In the Edit tree, double-click **Fabric Attach**.
3. On the work area, click the **Ports** tab.
4. To enable or disable the transmission of FA information in FA Signaling, select **enabled** or **disabled** in the **State** field for a specific port or ports.

5. To enable or disable message authentication, select **enabled** or **disabled** in the **MsgAuthStatus** field for a specific port or ports.
6. To configure the authentication key, enter an alphanumeric string of up to 32 characters in the **MsgAuthKey** field for a specific port or ports.
7. To configure the authentication key usage, select **strict** or **standard** in the **MsgAuthKeymode** field for a specific port or ports.
8. Click **Apply**.

## Field Descriptions

The following table describes the fields associated with FA Signaling or configuration of FA message authentication.

| Name                  | Description                                                                                                                                                                                                                                   |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IfIndex</b>        | Indicates the interface for which to configure FA operation and message authentication.                                                                                                                                                       |
| <b>State</b>          | Enables or disables FA operation on the interface.                                                                                                                                                                                            |
| <b>MsgAuthKey</b>     | Configures the authentication key for the specified interface.                                                                                                                                                                                |
| <b>MsgAuthStatus</b>  | Enables or disables FA message authentication on the interface.                                                                                                                                                                               |
| <b>MsgAuthKeymode</b> | Specifies the Authentication key usage setting — the user-defined authentication key (strict) or both the user-defined and default authentication keys (standard) are used for FA TLV data authentication.<br><br>Default key-mode is strict. |

## Displaying Fabric Attach elements

Use the following procedure to view discovered FA elements.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration>Edit**.
2. Click **Fabric Attach**.
3. In the work area, click the **Elements** tab.

## Elements Tab Field Descriptions

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

| Name           | Description                                                          |
|----------------|----------------------------------------------------------------------|
| <b>IfIndex</b> | Indicates the interface through which the FA element was discovered. |
| <b>Type</b>    | Indicates the FA element type.                                       |
| <b>Vlan</b>    | Indicates the management VLAN advertised by the FA element.          |

*Table continues...*

| Name                       | Description                                                                                                                                     |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Id</b>                  | Indicates the FA Element System ID, which is the unique system identifier used for connection management and limited device state distribution. |
| <b>State</b>               | Indicates the state flag data associated with the discovered FA element.                                                                        |
| <b>Auth</b>                | Indicates the authentication status for the discovered element.                                                                                 |
| <b>OperAuthStatus</b>      | Displays FA Element TLV authentication status detail data.                                                                                      |
| <b>AsgnsAuth</b>           | Indicates FA I-SID/VLAN Assignment TLV authentication status.                                                                                   |
| <b>AsgnsOperAuthStatus</b> | Displays FA I-SID/VLAN Assignment TLV authentication status detail data.                                                                        |

## Automating configurations for FA Clients

Use the following procedure to automate configurations for specific types of FA Clients.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Fabric Attach**.
3. In the work area, click the **Zero Touch** tab.
4. To enable or disable Zero Touch support, click enable or disable in the **ZeroTouchService** field.
5. To enable or disable Fabric Attach Mgmt VLAN Distribution, click enable or disable in the **ZeroTouchMgmtVlanDist** field.
6. To enable Zero Touch options, select the appropriate check-box in the **OptionFlags** field.
7. Specify the FA Client type ID for the selected OptionFlag:
  - Specify the FA Client type ID in the **autoPortModeFaClient** field to automate the configuration of EAP port modes.
  - Specify the FA Client type ID in the **autoTrustedModeFaClient** field to automate the FA Client connection default QoS treatment.
  - Specify the FA Client type ID in the **autoPvidModeFaClient** field to automate client PVID/ Mgmt VLAN updates.
  - Specify the FA Client type ID in the in the **autoClientAttach** field to automate the FA Client Attach field.
  - Specify the FA Client type ID in the in the **autoMgmtVlanFaClient** field to automate the FA Client auto mgmt Vlan.
8. Click **Apply**.

## Field Descriptions

The following table describes the fields associated with automate configurations for specific types of FA Clients.

Use the data in the following table to use the **Zero Touch** tab.

| Name                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ZeroTouchService</b> | Specifies whether Zero Touch support is enabled or disabled. The default is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>OptionFlags</b>      | Indicates the configured FA option flags for Zero Touch. <ul style="list-style-type: none"> <li>ipAddrDhcp — automates DHCP IP address acquisition.</li> <li>autoPortModeFaClient: Automates the configuration of EAP port modes.</li> <li>autoTrustedModeFaClient: Automates the FA Client connection default QoS treatment.</li> <li>autoPvidModeFaClient: Automates client PVID/Mgmt VLAN updates.</li> <li>autoClientAttach: Automates Zero Touch Client Attach configuration.</li> <li>autoMgmtVlanFaClient: Automates the FA Client auto mgmt Vlan.</li> </ul> |
| <b>Type</b>             | Indicates the configured FA Client type ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Descr</b>            | Indicates the configured FA Client type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Configuring Zero Touch Client Auto Attach

### Procedure


1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Fabric Attach**.
3. In the work area, click the **Zero Touch Client Auto Attach** tab.
4. Click **Insert**.
5. Select the Zero Touch Client Auto Attach type from the **Type** list and click **Ok**.
6. Specify a VLAN in the **Vlan** field.
7. Specify I-SID in the **Isid** field.
8. Specify a priority in the **PortPriority** field.
9. Select keepStaticVlan to keep the static VLANs or select removeStaticVlans to remove the static VLANs from the **ExcludeStatic** options.

Depending on the selection, the static VLANs are kept or removed on the client port for the duration of the client connection.

10. Click **Insert**.

## Zero Touch Client Auto Attach Tab Field Descriptions

Use the data in the following table to use the **Zero Touch Client Auto Attach** tab.

| Name                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ClientName</b>    | <p>Specifies an FA client type or a list of FA client types. Following are the available client types:</p> <ul style="list-style-type: none"> <li>• 6—Wireless AP (Type 1)</li> <li>• 7—Wireless AP (Type 2)</li> <li>• 8—Switch</li> <li>• 9—Router</li> <li>• 10—IP Phone</li> <li>• 11—IP Camera</li> <li>• 12—IP Video</li> <li>• 13—Security Device</li> <li>• 14—Virtual Switch</li> <li>• 15—Server Endpoint</li> <li>• 16—ONA (SDN)</li> <li>• 17—ONA (SpbOIp)</li> </ul>                                                                                        |
| <b>Type</b>          | <p>Specifies the Zero Touch Client Auto Attach type. If this type matches the FA interface type, Zero Touch Client Auto Attach specifications are applied to the port.</p> <p>Type 0 applies the specifications to non-EAP enabled, non-FA Client ports.</p> <p>Type 1 applies the specifications to non-EAP enabled, FA Client (any) ports.</p> <p> <b>Note:</b></p> <p>Zero Touch Client Auto Attach specification processing terminates if no applicable interfaces are found.</p> |
| <b>Vlan</b>          | Specifies the VLAN ID. The value range is from 1 to 4094.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Isid</b>          | Specifies the Isid value. The value range is from 0 to 16777214.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>PortPriority</b>  | Specifies 802.1p user priority. The value range is from 1 to 7.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>ExcludeStatic</b> | <p>Specifies whether static VLANs should be kept or removed on the client port for the duration of the client connection.</p> <p>Default is RemoveStaticVlans.</p>                                                                                                                                                                                                                                                                                                                                                                                                       |

## Displaying Fabric Attach statistics

Use any one of the following procedures to view the Fabric Attach statistics:

- [Displaying Fabric Attach statistics for ports](#) on page 135
- [Displaying Fabric Attach in a graph](#) on page 136
- [Displaying Fabric Attach statistics for chassis](#) on page 138

## Displaying Fabric Attach statistics for ports

### About this task

Use the following procedure to view FA statistics based on port index.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**
2. Click **Fabric Attach**.
3. In the work area, click the **Port Stats** tab.
4. Select a port row.
5. **(Optional)** Click **Graph** to view the statistics.
6. **(Optional)** Click **Clear Counters** to clear the counters and start over at zero.

### Port Stats Tab Field Descriptions

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

| Name             | Description                                                                                                                                                                                                                                                                             |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PortIndex        | Indicates the port for which FA statistics are displayed.                                                                                                                                                                                                                               |
| DiscElemReceived | Indicates the number of FA Element TLVs received on the identified port.                                                                                                                                                                                                                |
| DiscElemExpired  | Indicates the number of discovered FA elements from received FA Element TLVs that have expired on the identified port. This counter is not incremented when elements are deleted for reasons other than expiration.                                                                     |
| DiscElemDeleted  | Indicates the number of discovered FA elements from received FA Element TLVs that have been deleted on the identified port. This counter is only incremented when elements are deleted for reasons other than expiration.                                                               |
| DiscAuthFailed   | Indicates the number of received FA Element TLVs for which authentication was attempted and failed on the identified port.                                                                                                                                                              |
| AsgnReceived     | Indicates the number of I-SID/VLAN bindings received in FA I-SID/VLAN Assignment TLVs on the identified port.                                                                                                                                                                           |
| AsgnAccepted     | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are accepted (activated) on the identified port. This counter is incremented when the binding transitions from a non-accepted state such as 'pending' or 'rejected' to the accepted state. |

*Table continues...*

| Name           | Description                                                                                                                                                                                                                                                                 |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AsgnRejected   | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are rejected on the identified port. This counter is incremented when the binding transitions from a non-rejected state such as 'pending' or 'accepted' to the rejected state. |
| AsgnExpired    | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have expired on the identified port. This counter is not incremented when bindings are deleted for reasons other than expiration.                                              |
| AsgnDeleted    | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have been deleted on the identified port. This counter is only incremented when bindings are deleted for reasons other than expiration.                                        |
| AsgnAuthFailed | Indicates the number of received FA I-SID/VLAN Assignment TLVs for which authentication was attempted and failed on the identified port.                                                                                                                                    |

## Displaying Fabric Attach statistics in a graph

### About this task

Use the following procedure to view FA statistics.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**
2. Click **Fabric Attach**.
3. In the work area, click the **Port Stats** tab.
4. Select a port row.
5. Click **Graph**.

The FA Stats tab appears.

6. From the work area, select **Poll Interval**.

The table data refreshes automatically based on the value selected in the Poll Interval field.

7. **(Optional)** Click **Clear Counters** to clear the counters and start over at zero.

### FA Stats Tab Field Descriptions

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



| Name             | Description                                                                                                                                                                                                                                                                             |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DiscElemReceived | Indicates the number of FA Element TLVs received on the identified port.                                                                                                                                                                                                                |
| DiscElemExpired  | Indicates the number of discovered FA elements from received FA Element TLVs that have expired on the identified port. This counter is not incremented when elements are deleted for reasons other than expiration.                                                                     |
| DiscElemDeleted  | Indicates the number of discovered FA elements from received FA Element TLVs that have been deleted on the identified port. This counter is only incremented when elements are deleted for reasons other than expiration.                                                               |
| DiscAuthFailed   | Indicates the number of received FA Element TLVs for which authentication was attempted and failed on the identified port.                                                                                                                                                              |
| AsgnReceived     | Indicates the number of I-SID/VLAN bindings received in FA I-SID/VLAN Assignment TLVs on the identified port.                                                                                                                                                                           |
| AsgnAccepted     | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are accepted (activated) on the identified port. This counter is incremented when the binding transitions from a non-accepted state such as 'pending' or 'rejected' to the accepted state. |
| AsgnRejected     | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are rejected on the identified port. This counter is incremented when the binding transitions from a non-rejected state such as 'pending' or 'accepted' to the rejected state.             |
| AsgnExpired      | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have expired on the identified port. This counter is not incremented when bindings are deleted for reasons other than expiration.                                                          |
| AsgnDeleted      | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have been deleted on the identified port. This counter is only incremented when bindings are deleted for reasons other than expiration.                                                    |
| AsgnAuthFailed   | Indicates the number of received FA I-SID/VLAN Assignment TLVs for which authentication was attempted and failed on the identified port.                                                                                                                                                |

## Displaying Fabric Attach statistics for chassis

### About this task

Use the following procedure to view FA statistics.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Graph**
2. Click **Chassis**.
3. In the work area, click the **Fabric Attach** tab.
4. Select a port row.
5. The table data refreshes automatically based on the value selected in the **Poll Interval** field.
6. **(Optional)** Click **Clear Counters** to clear the counters and start over at zero.

### FA Stats Tab Field Descriptions

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

| Name             | Description                                                                                                                                                                                                                                                                             |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DiscElemReceived | Indicates the number of FA Element TLVs received on the identified port.                                                                                                                                                                                                                |
| DiscElemExpired  | Indicates the number of discovered FA elements from received FA Element TLVs that have expired on the identified port. This counter is not incremented when elements are deleted for reasons other than expiration.                                                                     |
| DiscElemDeleted  | Indicates the number of discovered FA elements from received FA Element TLVs that have been deleted on the identified port. This counter is only incremented when elements are deleted for reasons other than expiration.                                                               |
| DiscAuthFailed   | Indicates the number of received FA Element TLVs for which authentication was attempted and failed on the identified port.                                                                                                                                                              |
| AsgnReceived     | Indicates the number of I-SID/VLAN bindings received in FA I-SID/VLAN Assignment TLVs on the identified port.                                                                                                                                                                           |
| AsgnAccepted     | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are accepted (activated) on the identified port. This counter is incremented when the binding transitions from a non-accepted state such as 'pending' or 'rejected' to the accepted state. |
| AsgnRejected     | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are                                                                                                                                                                                        |

*Table continues...*

| Name           | Description                                                                                                                                                                                                                          |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                | rejected on the identified port. This counter is incremented when the binding transitions from a non-rejected state such as 'pending' or 'accepted' to the rejected state.                                                           |
| AsgnExpired    | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have expired on the identified port. This counter is not incremented when bindings are deleted for reasons other than expiration.       |
| AsgnDeleted    | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have been deleted on the identified port. This counter is only incremented when bindings are deleted for reasons other than expiration. |
| AsgnAuthFailed | Indicates the number of received FA I-SID/VLAN Assignment TLVs for which authentication was attempted and failed on the identified port.                                                                                             |

## Displaying Fabric Attach statistics summary

### About this task

Use the following procedure to view FA statistics summary.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Fabric Attach**.
3. In the work area, click the **Stats Summary** tab.
4. Select the **ClearGlobalErrorCounters(Summary)** check-box to clear the global error counters.

### Stats Summary Tab Field Descriptions

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

| Name                              | Description                                                                                                                                                                                                         |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ClearGlobalErrorCounters(Summary) | Clears the global error counters.                                                                                                                                                                                   |
| DiscElemReceived                  | Indicates the number of FA Element TLVs received on the identified port.                                                                                                                                            |
| DiscElemExpired                   | Indicates the number of discovered FA elements from received FA Element TLVs that have expired on the identified port. This counter is not incremented when elements are deleted for reasons other than expiration. |

*Table continues...*

| Name            | Description                                                                                                                                                                                                                                                                             |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DiscElemDeleted | Indicates the number of discovered FA elements from received FA Element TLVs that have been deleted on the identified port. This counter is only incremented when elements are deleted for reasons other than expiration.                                                               |
| DiscAuthFailed  | Indicates the number of received FA Element TLVs for which authentication was attempted and failed on the identified port.                                                                                                                                                              |
| AsgnReceived    | Indicates the number of I-SID/VLAN bindings received in FA I-SID/VLAN Assignment TLVs on the identified port.                                                                                                                                                                           |
| AsgnAccepted    | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are accepted (activated) on the identified port. This counter is incremented when the binding transitions from a non-accepted state such as 'pending' or 'rejected' to the accepted state. |
| AsgnRejected    | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that are rejected on the identified port. This counter is incremented when the binding transitions from a non-rejected state such as 'pending' or 'accepted' to the rejected state.             |
| AsgnExpired     | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have expired on the identified port. This counter is not incremented when bindings are deleted for reasons other than expiration.                                                          |
| AsgnDeleted     | Indicates the number of I-SID/VLAN bindings from received FA I-SID/VLAN Assignment TLVs that have been deleted on the identified port. This counter is only incremented when bindings are deleted for reasons other than expiration.                                                    |
| AsgnAuthFailed  | Indicates the number of received FA I-SID/VLAN Assignment TLVs for which authentication was attempted and failed on the identified port.                                                                                                                                                |

# Chapter 4: Layer 2 VSN Configuration

This chapter provides conceptual and procedural information related to the configuration and management of Layer 2 Virtual Services Networks (VSN).

---

## Layer 2 VSN configuration fundamentals

This section provides fundamentals concepts for Layer 2 Virtual Services Networks (VSN).

---

### SPBM Layer 2 VSN

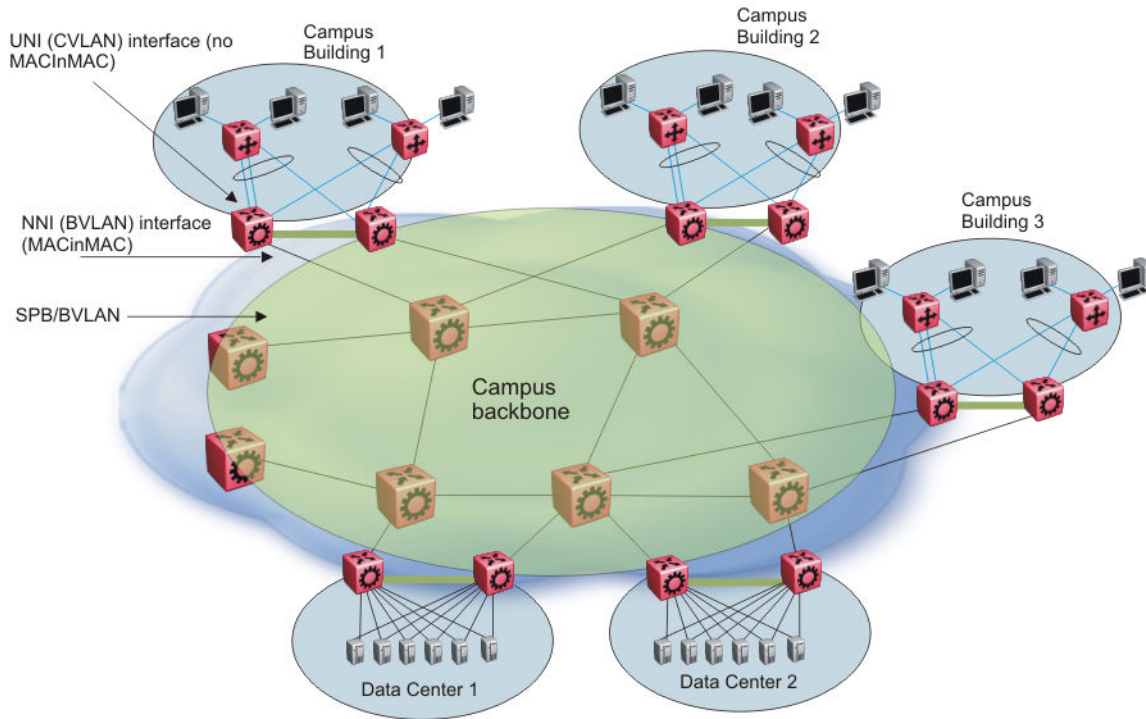
Shortest Path Bridging MAC (SPBM) supports Layer 2 VSN functionality where customer VLANs (C-VLANs) and Switched UNIs are bridged over the SPBM core infrastructure.

At the Backbone Edge Bridges (BEBs), customer VLANs (C-VLAN) and Switched UNIs 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 or I-SID to Switched UNI 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.

## Layer 2 VSN Configuration



**Figure 5: SPBM Layer 2 VSN in a campus**

One of the key advantages of the SPBM Layer 2 VSN is that you can achieve network virtualization provisioning 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 you add connectivity services to the SPBM network. For example, when you create new virtual server instances that require their own VLAN instances, you can provision at the network edge only and do not need configure 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.

### C-VLAN UNI

C-VLAN UNIs are created by the association of VLANs to I-SIDs. A VLAN with an I-SID configured becomes a C-VLAN. All ingress traffic of the VLAN from any member ports belong to the configured I-SID. C-MAC learning occurs inside the I-SID, on both UNI and NNI side (C-MAC + I-SID pointing to UNI port from the UNI side traffic, or C-MAC + I-SID pointing to a remote SPBM node - where the source C-MAC is connected).

Broadcast, unknown multicast and unknown unicast traffic in the I-SID is replicated to all local I-SID endpoints, including all C-VLAN member ports along with switched UNIs, and to all remote endpoints carried by the I-SID's multicast group. For UNI originated broadcast traffic, the originating endpoint is excluded from flooding, and the ingress port for broadcast traffic coming in on an NNI is excluded from flooding.

## Switched UNI

Switched UNI allows 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.

**\* Note:**

If the switch is connected to an SPB network that has IP Shortcuts or L3VSN enabled, you can create a management VLAN on the switch with no port members, and assign it to an I-SID for Layer 2 VSN terminated on an ERS 8800 with the same I-SID and IP subnet.


To allow IP connectivity to the switch, on the ERS 8800 where the Layer 2 VSN is configured, add an IP address to the VLAN that terminates the L2VSN.

---

## SPBM Layer 2 VSN sample operation

The following section shows how a SPBM network is established, in this case, a Layer 2 VSN. This release supports only Layer 2 VSN.

1. *Discover network topology*



**Figure 6: SPBM topology discover**

IS-IS runs on all nodes of the SPBM domain. IS-IS is the basis of SPBM, the IS-IS adjacency must be formed first. After the neighboring nodes see hellos from each other, the nodes 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 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 7: SPBM BMAC and I-SID population**

BMAC and I-SID information floods throughout the network to announce new I-SID memberships. In this case, VLAN 20 is mapped to I-SID 100.

#### \* Note:

I-SIDs are only used for virtual services (Layer 2 VSNs and Layer 3 VSNs). If IP Shortcuts only is enabled on the BEBs, I-SIDs are never exchanged in the network as IP Shortcuts allow for IP networks to be transported across IS-IS.

Each node populates its FDB with the BMAC information derived from the IS-IS shortest path tree calculations. 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, three shortest path trees are calculated (not counting the Equal Cost Trees (ECTs)).




Figure 8: Shortest path tree for source node A





Figure 9: Shortest path tree for source node B



**Figure 10: 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, therefore 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 the traffic is sent as a multicast destination frame, where the multicast MAC is created from the Nick-name of the source bridge and the I-SID. 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 that 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.


## SPBM IP shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

You can use SPBM IP Shortcuts as an alternative to SPBM Layer 2 VSN for forwarding traffic in an SPBM setup. Both features can function together in the same setup.

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 11: 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.

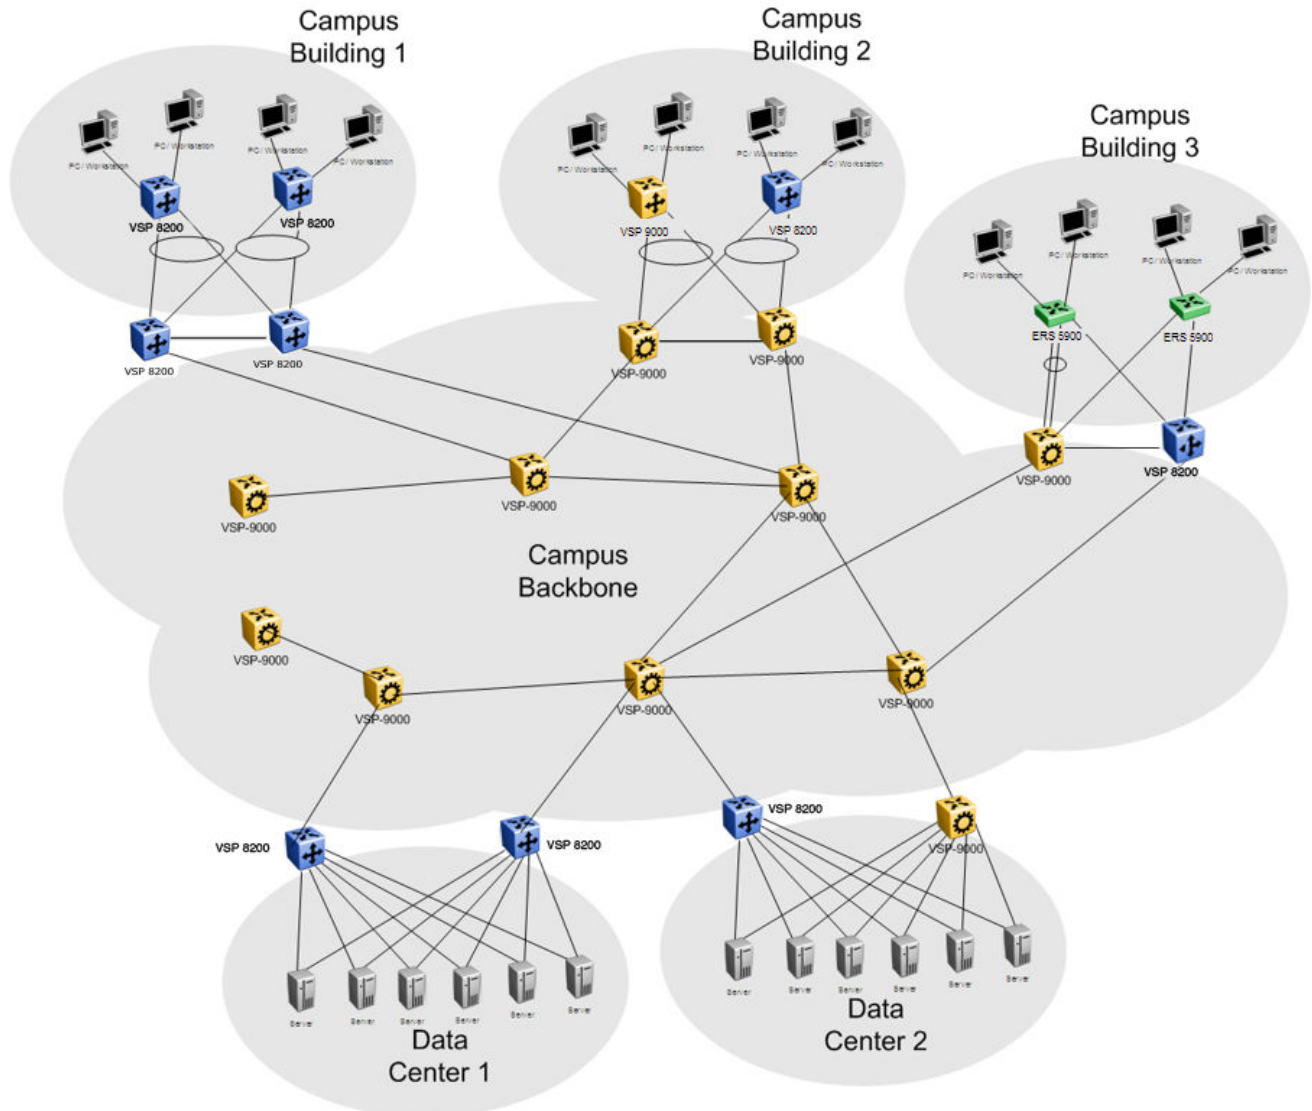
In contrast to IP routing or Multiprotocol Label Switching (MPLS), SPBM IP shortcuts provide a simpler method of forwarding IP packets in an Ethernet network using the preestablished Ethernet FIBs on the BEBs. SPBM allows a network to make the best use of routing and forwarding techniques, where only 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 12: 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 and static routes into IS-IS.

### ! Important:

By mapping the management VLAN to an I-SID (VSN) you can manage the switch both from the UNI side and from the NNI side. You cannot manage the switch from C-VLANs (VSNs) other than the management VLAN (VSN).

The switch supports routing for non-VSN VLANs over the SPBM Fabric, using IP shortcuts. However, the switch does not support the following:

- Inter-VSN unicast routing
- Unicast routing between non-SPB VLANs and C-VLANs (L2 VSN)

### \* Note:

In this release, redistribute routing protocols are not supported. Only direct and static redistribution is supported.

### \* Note:

You must obtain an appropriate license in order to enable IP Shortcuts.

### \* Note:

IP Shortcut unicast traffic is transmitted as unencapsulated IP packet whereas L2VSN and IPSC multicast are Mac-in-Mac encapsulated. CoS/DSCP marking for IPSC unicast packets adheres to the default untrusted QoS policy for IP packets in non-SPB environment.

## Limitations

The following limitations apply to the SPBM IP shortcuts feature:

- The switch supports 256 Layer 3 VLANs. With routing enabled, you can configure up to 256 VLANs on the UNI side of the switch.
- IP Shortcut routes count against the limit of supported IPv4 routes (up to 4096 routes on ERS 5900 or up to 2048 routes on ERS 4900).
- The NNI side of the switch supports up to 1000 Layer 2 VSNs and up to 1000 nodes.
- IP routing cannot be enabled before SPBM and IPSC are enabled.

## 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.

You must enable IP routing.

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 for EDM and CLI.

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.

## Limitations

The following limitations apply to the *IP Multicast over Fabric Connect within the GRT* feature:

- Only IPv4 multicast traffic is supported
- Only the existing IGMP functionality is supported
- The SPB BEB devices interact with the receivers and senders connected either directly, or through a Layer 2 aggregation or Layer 2 switch
- There is no interaction with PIM/Multicast routers on the access networks
- You must reserve a loopback port, using one of the stack ports in standalone mode or the last front panel port in stack mode
- The maximum number of allowed streams is 1024.
- The *Unknown Multicast no Flood* feature is mutual exclusive with *SPBM Multicast*
- The switch does not support SPBM Multicast on Switched UNI VLANs
- Multicast is not supported for multiple C-VLANs added to the same ISID. Only one C-VLAN with snooping enabled can be added to an I-SID
- You cannot configure VLAN ID 4060, as it is used as an internal VLAN for supporting multicast over SPB for GRT scope. VLAN 4060 is also reserved for non-SPB environment.

## Scaling limits

The switch supports up to 256 VLANs with multicast enabled. 1024 local or remote streams are supported on a BEB.

## ECMP support for IP Shortcuts

The Equal Cost Multipath (ECMP) feature supports IP Shortcuts.

With ECMP, the switch can determine up to four equal-cost paths 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 traffic. Equal Cost Multipath is formed using routes from the same protocol.

### Note:

IP routing cannot be enabled before SPBM and IPSC are enabled.

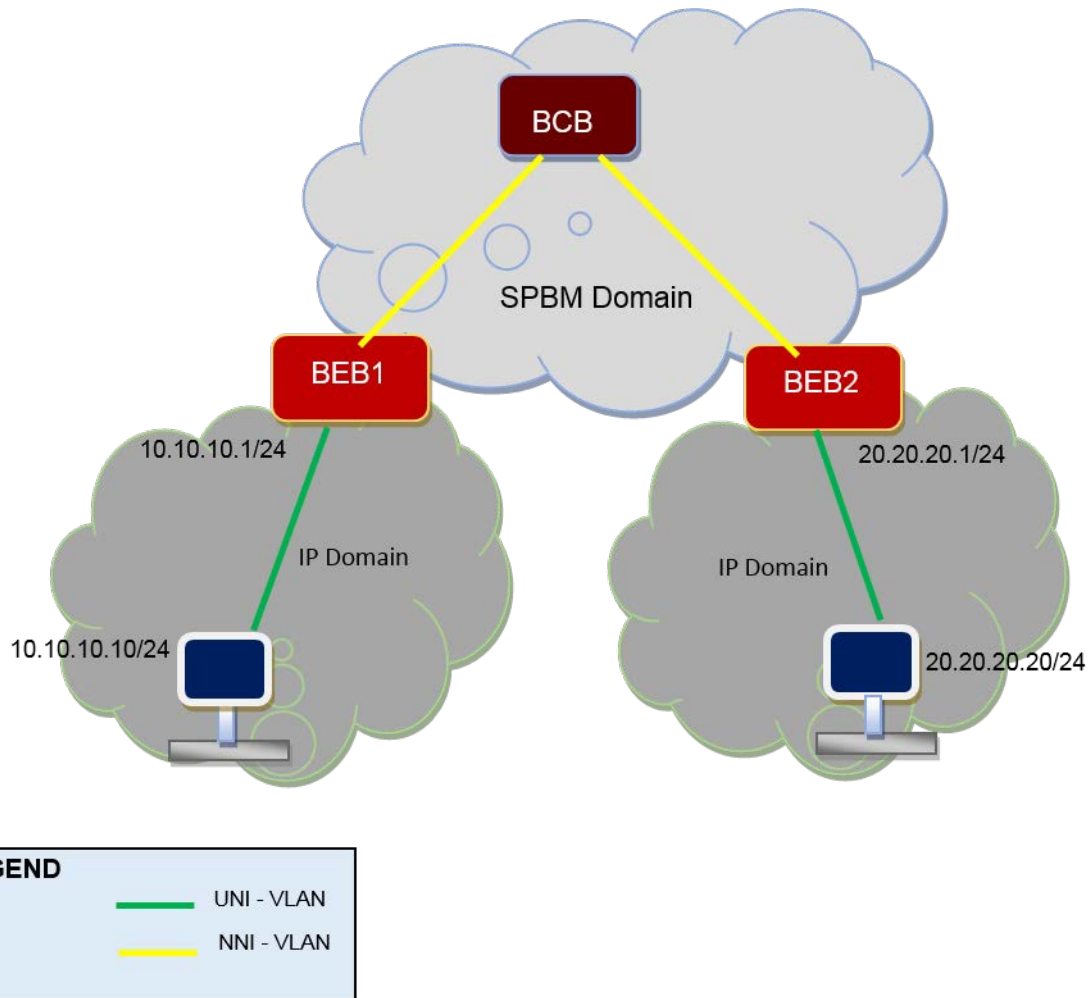
## Topology

This section explains how SPBM IP shortcut connects two IP domains across the SPBM domain.

- [Topology without IP Shortcuts](#) on page 154
- [SPBM topology for protocol route redistribution](#) on page 155
- [SPBM ECMP topology](#) on page 156

Topology without IP Shortcuts

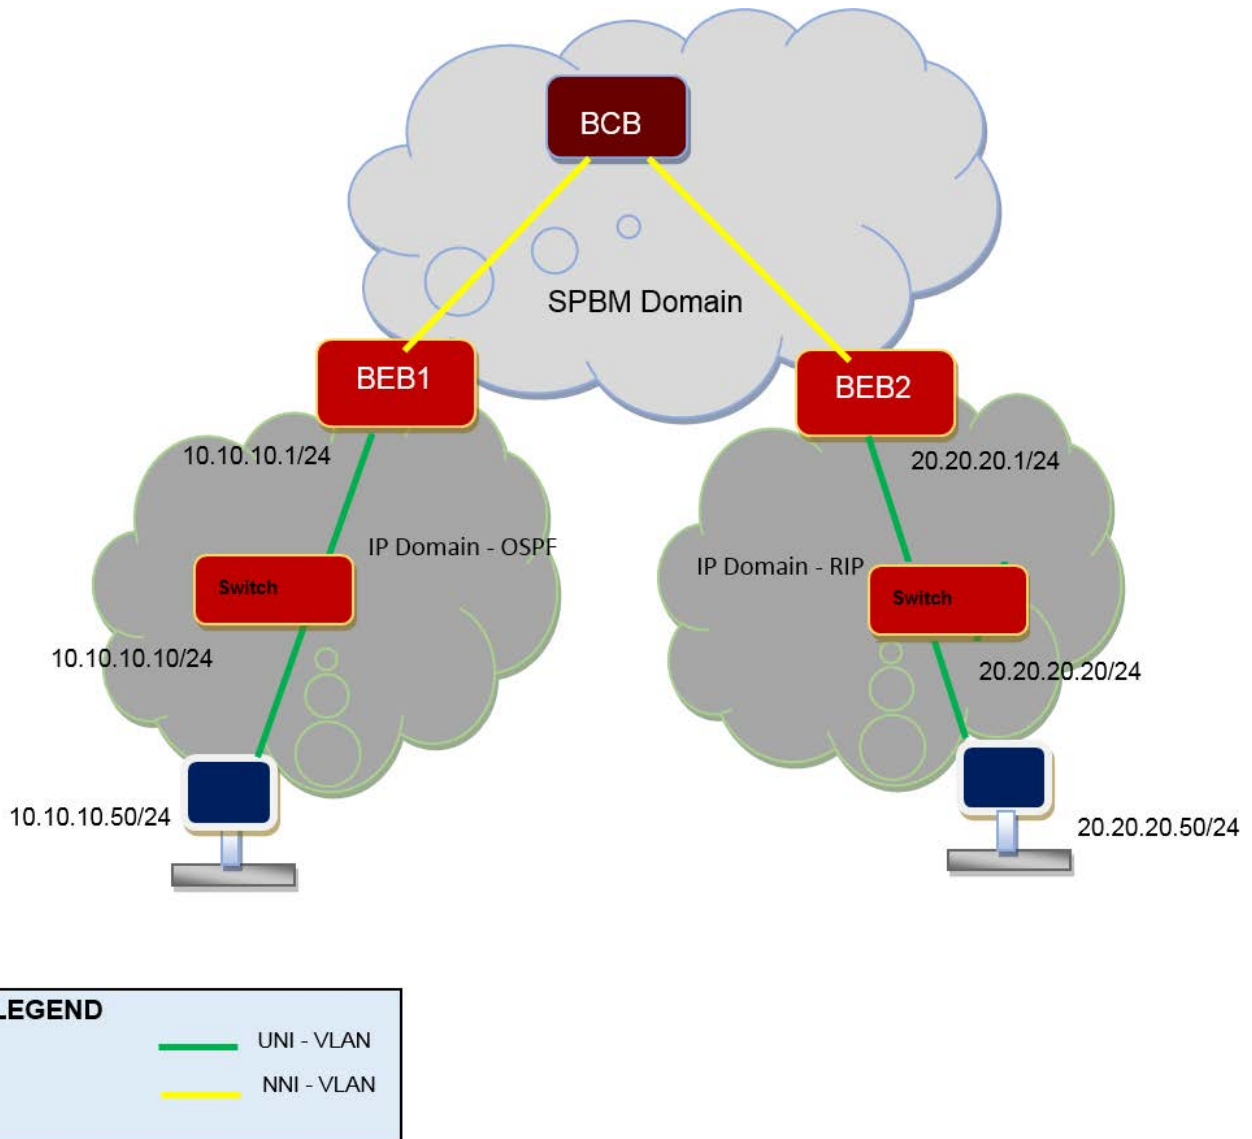
In the following topology, without SPBM IP shortcut the two IP domains cannot communicate with each other.



**Figure 13: Communication between domains without IP Shortcuts**

SPBM topology for protocol route redistribution

In the following topology diagram, the SPB cloud connects the BEB1, BCB1 and BEB2 by running ISIS between the core (BCB1) and the edges (BEB1 and BEB2). This provides Layer 2 connectivity from BEB1 to BEB2.




**Figure 14: SPBM topology for protocol route redistribution**

For the IP Layer connection between two IP domains, the IP routes from both the IP domains must be exchanged. After the BEBs know the IP and Layer 2 reachability, BEBs use the source and destination MACs learnt through Layer 2 ISIS learning (BEB1\_MAC, BEB2\_MAC) and forward the IP data packet to BCB. BCB switches the data packet based on the destination BEB MAC (regular Layer 2 switching).

**SPBM ECMP topology**

In the following diagram, to support ECMP, BEB configures primary and secondary BVLAN and connects to two BCBs. BEB1 Route Table Manager (RTM) has two route entries to 20.20.20.20/24 with different next-hop. One is BCB1 and other is BCB2.



SPBM assigns one BVLAN for each BCB. There are always two paths to reach remote BEB, one on each BVLAN. The ISIS route learned from remote BEB results in two ECMP paths. There are two route entries for each remote ISIS route in the RTM for each BEB.

---

## Layer 2 VSN configuration using CLI

This section provides procedures to configure Layer 2 Virtual Services Networks (VSN) using Command Line Interface (CLI).

---

### Configuring a SPBM Layer 2 VSN C-VLAN

Shortest Path Bridging MAC (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):

```
i-sid <1-16777214> vlan <1-4094>
```

3. Display C-VLAN information:

```
show i-sid <1-16777214>
```

**Example**

```
Switch> enable
Switch# configure terminal
Switch(config)# i-sid 200 vlan 200
Switch(config)# show i-sid 200
```

| I-SID | Vid | UNI-type | Ports |
|-------|-----|----------|-------|
| 200   | 200 | C-VLAN   | 7     |

**Variable definitions**

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

| Variable                                                  | Value                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>i-sid &lt;1-16777214&gt; vlan &lt;1-4094&gt;</code> | <p>Specifies the customer VLAN (C-VLAN) 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><b>* Note:</b></p> <p>The switch reserves I-SID 0x00ffffff and uses this I-SID to advertise the virtual B-MAC in a SMLT dual-homing environment. The platform clears the receive and transmit</p> |

| Variable | Value                                                                               |
|----------|-------------------------------------------------------------------------------------|
|          | bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service. |

## 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).
- 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.
- You must enable IGMP Snooping on the C-VLANs where IGMP receivers are present or multicast traffic is present.

#### \* Note:

To enable IGMP Snooping on a VLAN interface, enable SPBM multicast.

### About this task

When IGMP snooping is enabled on C-VLAN, 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 multicast over SPBM.

SPBM supports enabling IGMP Snooping on a C-VLAN, but it does not support enabling Protocol Independent Multicast (PIM) on a C-VLAN. If you enable IGMP snooping on a C-VLAN, then its operating mode is Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network.

In this release, the switch only supports IPv4 multicast traffic.

### Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4094>
```

2. Enable IGMP snooping:

```
ip igmp snooping
```

3. **(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 (but it is highly recommended). 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.

### 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 snooping
Switch:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
```

## 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. To enter User EXEC mode, log on to the switch.
2. 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 spb-mcast-summary
```

```
=====
 SPB Multicast - Summary
=====
SCOPE SOURCE GROUP DATA LSP HOST
I-SID ADDRESS ADDRESS I-SID BVID FRAG NAME

80 192.0.2.1 203.0.113.3 16300014 1000 0x0 MERS2-8606
80 192.0.2.1 203.0.113.4 16300015 1000 0x0 MERS2-8606
80 192.0.2.3 203.0.113.3 16300001 1001 0x0 MERS4-8606
80 192.0.2.3 203.0.113.4 16300002 1001 0x0 MERS4-8606
200 192.0.2.4 203.0.113.2 16000001 1000 0x1 4826GTS
80 192.0.2.5 203.0.113.2 16000003 1000 0x1 4826GTS

6 out of 6 Total Num of Entries
```



## Variable definitions

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

| Variable                      | Value                                                                        |
|-------------------------------|------------------------------------------------------------------------------|
| host-name <i>WORD</i> <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.    |

## 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
```

### \* Note:

You must configure basic SPBM and IS-IS infrastructure. For more information, see [Configuring minimum SPBM and IS-IS parameters](#) on page 56.

## 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 [vlan <1-4084>]
```
3. Display information about the IGMP cache:
 

```
show ip igmp cache
```
4. Display information about the IGMP group:

```
show ip igmp group [count][group {A.B.C.D}][member-subnet
{A.B.C.D/<0-32>}]
```

5. Display information about IGMP snoop information:

```
show ip igmp snoop
```

**Example**

The following example displays the output for the **show ip igmp interface** command.

```
Switch>enable
Switch#show ip igmp interface vlan 1
 Query Oper
VLAN Intvl Vers Vers Querier Query Wrong LastMbr Send

1 125 2 2 0.0.0.0 100 0 0 2 10 No
```

The following example displays the output for the **show ip igmp cache** command.

```
Switch#show ip igmp cache
Group Address Vlan ID Last Reporter Expiration Vl Host Timer Type

239.255.255.250 1 172.16.120.253 160 0 Dynamic
```

The following example displays the output for the **show ip igmp group** command.

```
Switch#show ip igmp group
Group Address VLAN Member Address Expiration Type In Port

239.255.255.250 1 172.16.120.253 188 Dynamic 1
```

**Variable definitions**

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

| Variable      | Value               |
|---------------|---------------------|
| vlan <1–4084> | Specifies the VLAN. |

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. |

**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. To enter User EXEC mode, log on to the switch.
2. Display IS-IS Link State Database information by Type-Length-Value (TLV):

```
show isis lsdb [tlv <1-144>]
```

3. Display IS-IS Link State Database information by Link State Protocol ID:

```
show isis lsdb [lsp-id <xxxx.xxxx.xxxx.xx-xx>] [detail] [tlv <1-144>]
```

## Example

```
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)

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
```

## Layer 2 VSN Configuration

```
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 in the current release. |
| lspid <xxxx.xxxx.xxxx.xx-xx> | Specifies information about the IS-IS Link State database by LSP ID.                                                                                                                                               |
| sysid <xxxx.xxxx.xxxx>       | Specifies information about the IS-IS Link State database by System ID.                                                                                                                                            |
| tlv <1-186>                  | Specifies information about the IS-IS Link State database by TLV.                                                                                                                                                  |

## Configuring the IP Multicast over Fabric Connect forward cache timeout value

Use this procedure to configure the timeout value. 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.
- 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 <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
no spbm <1-100> multicast fwd-cache-timeout
```

**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.                                                                                 |
| <10-86400> | Specifies the IP Multicast over Fabric Connect forward-cache timeout in seconds. The default is 210 seconds. |

**Configuring the loopback port**

A loopback port must be configured to enable IP Multicast over Fabric Connect.

**About this task**

Use this procedure to configure one port in loopback (one of the stack ports or the last front panel port) and to enable IP Multicast over Fabric Connect if it is not previously enabled.

**\* Note:**

After you set the stacking port in loopback, the unit is not able to be stacked.

After you configure the front panel port in loopback, the last port on each unit in stack is occupied. This port cannot be used for user traffic and no longer appears in the available interface list.

**\* Note:**

The ERS 5900 unit or stack does not reset to partial default when the loopback port is configured on reserved-port front-panel [(51,52), where last port and last port-1 are used]. Configuration settings are maintained after the unit or stack restarts.

**Procedure**

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the loopback port.

On units without an internal loopback port, enter the following command:

```
spbm reserved-port {front-panel | stack}
```

OR

On units with an internal loopback port, enter the following command:

```
spbm reserved-port {internal | stack}
```

**\* Note:**

The device must reset in order for the configuration change to become effective.

3. (Optional) Disable the stack port or front panel port:

```
no spbm reserved-port
default spbm reserved-port
```

### Next steps

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

## Variable definitions

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

| Variable           | Value                              |
|--------------------|------------------------------------|
| <i>front-panel</i> | Set the loopback front panel port. |
| <i>stack</i>       | Set the loopback stack port.       |

## Viewing the loopback port settings

### About this task

Use this procedure to display the current and the next reserved loopback port settings.

### Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the current and the next reserved loopback port settings:

```
show spbm reserved-port
```

### Example

Display the current and the next reserved loopback port settings:

```
show spbm reserved-port
SPBM Current Reserved Port: front-panel(51,52)
SPBM Next Boot Reserved Port: front-panel(51,52)
#
```

## Configuring a SPBM Layer 2 VSN Switched UNI

Shortest Path Bridging MAC (SPBM) supports Layer 2 Virtual Service Network (VSN) functionality where Switched UNIs are bridged over the SPBM core infrastructure.

At the BEBs, Switched UNIs 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-Switched UNI VLAN provisioning.

### Before you begin

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

### About this task

To configure a Switched UNI, you must create a Switched UNI VLAN, and map an I-SID to the Switched UNI VLAN and a port.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a Switched UNI VLAN:

```
vlan create <2-4094> type spbm-switchedUni
```

3. Map a Switched UNI VLAN to a Service Instance Identifier (I-SID):

```
i-sid <1-16777214> vlan <2-4094> port <portlist>
```

**\* Note:**

You can run this command again to map a Switched UNI VLAN to multiple I-SIDs.

4. Display the Switched UNI information:

```
show i-sid <1-16777214>
```

**\* Note:**

You can verify the Switched UNI VLAN using `show i-sid` only. The `show vlan i-sid` command does not display Switched UNI details.

### Example

```
Switch> enable
Switch# configure terminal
Switch(config)# vlan create 100 type spbm-switchedUni
```



```
Switch(config)# i-sid 100 vlan 100 port 1
```

```
Switch(config)# show i-sid 100
```

| I-SID | Vid | UNI-type | Ports |
|-------|-----|----------|-------|
| 100   | 100 | switched | 1     |

You can map a Switched VLAN UNI to multiple I-SIDs.

```
Switch(config)# i-sid 101 vlan 100 port 2
```

```
Switch(config)# show i-sid
```

| I-SID | Vid | UNI-type | Ports |
|-------|-----|----------|-------|
| 100   | 100 | switched | 1     |
| 101   | 100 | switched | 2     |

## Variable definitions

Use the data in the following table to use the `i-sid vlan` command to configure a Switched UNI.

| Variable                                                                        | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>i-sid &lt;1-16777215&gt; vlan &lt;2-4094&gt; port &lt;portlist&gt;</code> | <p>Specifies the Switched UNI VLAN to associate with the I-SID. and a port.</p> <p>Use the <code>no</code> or <code>default</code> options to remove the I-SID from the specified VLAN.</p> <p><b>* Note:</b></p> <p>The switch reserves I-SID 0x00ffffff and 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 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>]
```

**Example**

The following example displays the output for the **show vlan i-sid** command.

```
Switch#show vlan i-sid
=====
Vlan I-SID
=====
VLAN_ID I-SID

1
2
5 5
10
20
```

The following example displays the output for the **show i-sid** command.

```
Switch>show i-sid
I-SID Vid UNI-type Ports

5 5 CVLAN 4
100 100 switched 1
```

The following example displays the output for the **show isis spbm i-sid** command.

```
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 ERS-4000
300 1.11.16 1000 0014.c7e1.33df config ERS-4000
400 1.11.16 1000 0014.c7e1.33df config ERS-4000
200 1.11.16 2000 0014.c7e1.33df config ERS-4000
300 1.11.16 2000 0014.c7e1.33df config ERS-4000
400 1.11.16 2000 0014.c7e1.33df config ERS-4000

Total number of SPBM ISID entries configed: 6

Total number of SPBM ISID entries: 6

```

**Variable definitions**

Use the data in the following table to use the **show vlan i-sid** commands.

| Variable              | Value                                                                                          |
|-----------------------|------------------------------------------------------------------------------------------------|
| <1-4094> <1-16777215> | Displays I-SID information for the specified C-VLAN. You can specify the VLAN ID and I-SID 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"> <li>all: displays all I-SID entries</li> <li>config: displays configured I-SID entries</li> <li>discover: displays discovered I-SID entries</li> </ul> |
| vlan <1-4094>                    | 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.                                                                                                                                    |

---

## Managing the switch via Layer 2

Use this procedure to manage the switch via Layer 2.

### About this task

To manage the switch via Layer 2, create a management VLAN on the switch with no port members, and assign it to an I-SID for Layer 2 VSN terminated on an ERS 8800 with the same I-SID and IP subnet.

To allow IP connectivity to the switch, add an IP address to the VLAN that terminates the L2VSN on the ERS 8800 where the Layer 2 VSN is configured.

### Procedure

1. Enter Global Configuration Mode:

```
enable
configure terminal
```

2. To create a management VLAN, enter the following commands at the command prompt:

```
vlan create <vlan_ID> type port
vlan mgmt <vlan_ID>
```

3. To assign the management VLAN to an I-SID, enter the following command at the command prompt:

```
i-sid <1-16777214> vlan <vlan_ID>
```

**Next steps**

On the ERS 8800 assign a VLAN to an I-SID with the same ID as the I-SID with which the management VLAN is associated on the switch, and add an IP address to this VLAN.

**Variable definitions**

| Variable           | Value                                                             |
|--------------------|-------------------------------------------------------------------|
| vlan_ID            | Specifies the management VLAN ID. Range is <2-4094>.              |
| i-sid <1-16777214> | Specifies the I-SID with which the management VLAN is associated. |

**SPBM IP Shortcuts configuration using CLI**

This section provides procedures to configure SPBM IP Shortcuts using 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.

**! Important:**

To enable IP Shortcuts you must obtain an appropriate license.

**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 <1-16>
```
2. Configure a CLIP interface to use as the source address for SPBM IP shortcuts:
 

```
ip address <A.B.C.D> <A.B.C.D>
```
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 {direct | static}
```
9. Enable routes to be announced into the SPBM network
 

```
redistribute {direct | static} enable
```
10. If you want to delete the configuration, use the no option:
 

```
no redistribute {direct | static}
no redistribute {direct | static} enable
```
11. Exit to Global Configuration mode:
 

```
exit
```
12. Apply the configured redistribution:
 

```
isis apply redistribute {direct | static}
```

**Example**

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface loopback 1
```

## Layer 2 VSN Configuration

```
Switch:1(config-if)# ip address 10.0.0.2 255.0.0.0
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 B-VID PRIMARY NICK LSDB IP MULTICAST
INSTANCE VLAN VLAN NAME TRAP
=====
4 400-401 400 1.10.01 disable enable enable
=====
```

```
Switch:1(config-isis)# redistribute static
Switch:1(config-isis)# redistribute static enable
Switch:1(config-isis)# exit
Switch:1(config)# isis apply redistribute static
Switch:1(config)# show isis redistribute
```

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

```
=====
 ISIS Redistribute List - GlobalRouter
=====
SOURCE MET MTYPE SUBNET ENABLE LEVEL

Direct 1 Type 1 Allow True
Static 1 Type 1 Allow True
=====
```

### Variable definitions

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

| Variable            | Value                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------|
| <A.B.C.D> <A.B.C.D> | Specifies an IP address and subnet mask. 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.<br>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                                                                                                                                                    |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| {direct   static} | Specifies the protocol.                                                                                                                                  |
| enable            | Enables the redistribution of the specified protocol into the SPBM network.<br>The default is disabled. Use the no option to disable the redistribution. |

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

| Variable          | Value                   |
|-------------------|-------------------------|
| {direct   static} | Specifies the protocol. |

## Configuring ECMP paths

Use the following procedure to set the maximum number of ECMP paths.

### Before you begin

Enable SPBM and IP Shortcut.

### Procedure

1. Enter Global Configuration mode:
 

```
enable
configure terminal
```
2. Enter the following command:
 

```
[no] [default] isis maximum-path <1-4>
```
3. Enter the following command to view the configuration:
 

```
show ecmp
```

### Example

The following is a sample output of the **show ecmp** command.

```
Switch# show ecmp
Protocol MAX-PATH

static: 1
rip: 1
ospf: 1
isis: 1
```

## Variable definitions

Use the data in the following table to use the `isis maximum-path` command.

| Variable | Value                                              |
|----------|----------------------------------------------------|
| <1–4>    | Specifies the ECMP path value. Default value is 1. |

## 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.
- You must globally enable IP routing.
- You must configure IP addresses on the VLAN interfaces.
- An appropriate license is required in order to enable IP Multicast within GRT on VLAN interface.
- You can enable IP Multicast within GRT on C-VLANs (VLANs with I-SIDs configured).

### Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- 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 VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4094>
```

2. Create an IP interface on the VLAN:

```
ip address <A.B.C.D> <A.B.C.D>
```

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 for EDM and CLI.

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<br>INTVL | STATUS | VERS.<br>VERS | OPER<br>VERS | QUERIER | QUERY<br>MAXRSPT | WRONG<br>QUERY | JOINS | ROBUST | LASTMEM<br>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               |            |

## 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. To enter User EXEC mode, log on to the switch.
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 the IP Multicast over Fabric Connect multicast group and source address information:

```
show isis spbm ip-multicast-route [group {A.B.C.D}] [source {A.B.C.D}] [source-beb WORD<0-255>]
```

5. Display summary information for each S, G, V tuple with the corresponding scope, data I-SID, and the host name of the source:

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

### 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.<br>Specifies the source BEB name. |
| vlan                                                      | Displays IP Multicast over Fabric Connect route information by VLAN.                                                                                                                            |
| 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 <i>WORD</i> <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"> <li>• routed—For GRT and Layer 3 VSN.</li> <li>• snoop—For Layer 2 VSN.</li> </ul>                                                                                                                                                                                                                    |
| 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 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. |

*Table continues...*

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                    |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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.                                                                                                                                                                                                                                                                                          |
| 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 vlan <1-4094>
```

Ensure that the output displays `routed-spb` under `MODE`.

3. Display information about the IGMP cache:

```
show ip igmp cache
```

4. Display information about the IGMP group:

```
show ip igmp group {count | group ABCD | member-subnet A.B.C.D/
<0-32>}
```

5. Display information about the IGMP sender:

```
show ip igmp sender {count | group ABCD | member-subnet A.B.C.D/
<0-32>}
```

## Variable definitions

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

| Variable      | Value               |
|---------------|---------------------|
| vlan <0-4094> | Specifies the VLAN. |

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. |

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. |

## Job aid

The following table describes the fields for the `show ip igmp interface vlan` command.

| Parameter   | Description                                                                          |
|-------------|--------------------------------------------------------------------------------------|
| VLAN        | Indicates the VLAN where IGMP is configured.                                         |
| Query Intvl | Indicates the frequency at which IGMP host query packets transmit on this interface. |

*Table continues...*

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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.                                                                                                                                                                                                                                 |
| LastMbr 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. |
| Send Query    | Indicates whether send query is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                          |
| 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 cache` command.

| Parameter     | Description                                                                                                                   |
|---------------|-------------------------------------------------------------------------------------------------------------------------------|
| Group Address | 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                                                                                                                                                                                                                     |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VLAN ID       | Indicates the logical interface (VLAN) which received group reports from various sources.                                                                                                                                       |
| Last Reporter | 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.                                                                                                                                                   |
| V1 Host Timer | 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.                                                                                                                                                      |

The following table describes the fields for the **show ip igmp group** command.

| Parameter      | Description                                                                                                                   |
|----------------|-------------------------------------------------------------------------------------------------------------------------------|
| Group Address  | Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address. |
| VLAN           | Indicates the logical interface (VLAN) which received group reports from various sources.                                     |
| Member Address | 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.                                                    |
| In Port        | Indicates the physical interface or a logical interface (VLAN) which received group reports from various sources.             |

The following table describes the fields for the **show ip igmp sender** command.

| Parameter      | Description                           |
|----------------|---------------------------------------|
| Group Address  | Indicates the IP multicast address.   |
| Interface      | Indicates the interface index number. |
| Member Address | Indicates the IP address of the host. |
| Port           | Indicates the IGMP sender ports.      |

---

## 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 C-VLANs

After you configure 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 tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. To map a C-VLAN to a Service instance identifier (I-SID), in the **I-sid** column, 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.

---

### Displaying the MAC address table for a C-VLAN

Use the following procedure to view the MAC Address table for a C-VLAN.



## Procedure

1. In the navigation tree, expand the following folders: **Configuration > IS-IS**.
2. Click **SPBM**.
3. In the work area, click the **MAC Addresses** tab.

## MAC Addresses Tab Field Descriptions

Use the data in the following table to use the **MAC Addresses** tab.

| Name             | Description                                                                                                                                                                                                                                                                    |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Isid</b>      | Indicates the I-SID for this MAC address.                                                                                                                                                                                                                                      |
| <b>Addr</b>      | Indicates the customer MAC address for which the bridge has forwarding and/or filtering information                                                                                                                                                                            |
| <b>CPort</b>     | Either displays the value 0, or indicates the port on which a frame came from.                                                                                                                                                                                                 |
| <b>CVlanId</b>   | Indicates the VLAN ID for this MAC address.                                                                                                                                                                                                                                    |
| <b>BDestAddr</b> | Indicates the provider MAC address for which the bridge has forwarding and/or filtering information.                                                                                                                                                                           |
| <b>Type</b>      | Indicates the MAC address learned type as local (C-VLAN or Switched UNI) or remote (B-VLAN). <ul style="list-style-type: none"> <li>• Type remote shows a BDestAddr associated, but no CVlanID.</li> <li>• Type local shows a CVlanID associated, but no BDestAddr.</li> </ul> |
| <b>Status</b>    | Indicates the status of this entry: <ul style="list-style-type: none"> <li>• other</li> <li>• invalid</li> <li>• learned</li> <li>• self</li> <li>• mgmt</li> </ul>                                                                                                            |

## 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 loopback port.


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

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
  2. Click **VLANs**.
  3. Click the **Basic** tab.
  4. Select a VLAN.
  5. Click **IP**.
  6. Select the **SnoopEnable** check box.
  7. **(Optional)** Select the **ProxySnoopEnable** check box.
  8. 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.
  9. **(Optional)** If you want to enable snoop querier, select **SnoopQuerierEnable**.
  10. **(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.
11. Click **Apply**.

## Displaying IS-IS redistribution

Use this procedure to view IS-IS redistribution. The routes are not redistributed into IS-IS automatically. To advertise the routes, you must explicitly redistribute direct or static protocols into IS-IS.

### Procedure

1. From the navigation tree, double-click **IS-IS**.
2. Click **IS-IS**.
3. Click the **Redistribute** tab.

## IS-IS Redistribute Tab Field Descriptions

Use the data in the following table to configure the **IS-IS Redistribute** tab.

| Name               | Description                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RouteSource</b> | Specifies the source protocol for the route redistribution entry.                                                                                                                               |
| <b>Enable</b>      | Enables or disables a redistribution entry. The default is disable.                                                                                                                             |
| <b>Metric</b>      | Specifies the metric for the redistributed route. The value can be a range between 0 to 65535. The default value is 0. You should use a value that is consistent with the destination protocol. |
| <b>MetricType</b>  | Specifies the metric type.                                                                                                                                                                      |
| <b>Subnets</b>     | Indicates whether the subnets are advertised individually or aggregated to their classful subnet.                                                                                               |

## Configuring SPBM switched UNIs

Use the following procedure to configure SPBM switched UNIs by mapping I-SIDs, VLANs, and ports.

### About this task

The VLAN must be type spbm-switchedUni. The port does not need to be a member of the VLAN, it is automatically added to the associated VLAN when you create the Switched UNI.

### Procedure

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

## Switched UNIs Tab Field Descriptions

Use the data in the following table to use the **Switched UNIs** tab.

| Name        | Description                              |
|-------------|------------------------------------------|
| <b>Isid</b> | Specifies the I-SID of the switched UNI. |
| <b>Port</b> | Specifies the port of the switched UNI.  |
| <b>Vlan</b> | Specifies the VLAN of the switched UNI.  |

---

## Managing the switch via Layer 2

Use this procedure to manage the switch via Layer 2.

### Before you begin

Create a management VLAN on the switch.

### About this task

To manage the switch via Layer 2, create a management VLAN on the switch with no port members, and assign it to an I-SID for Layer 2 VSN terminated on an ERS 8800 or ERS 4800 with the same I-SID and IP subnet.

To allow IP connectivity to the switch, add an IP address to the VLAN that terminates the L2VSN on the ERS 8800 where the Layer 2 VSN is configured.

### Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. Click the **Basic** tab.
4. To map the management VLAN to an I-SID, specify the I-SID to associate with the management VLAN in the **I-sid** column.
5. Click **Apply**.

### Next steps

On the ERS 8800 assign a VLAN to an I-SID with the same ID as the I-SID with which the management VLAN is associated on the switch, and add an IP address to this VLAN.

## Field Descriptions

| Name               | Description                                                       |
|--------------------|-------------------------------------------------------------------|
| I-sid <1-16777214> | Specifies the I-SID with which the management VLAN is associated. |

## SPBM IP Shortcuts configuration using EDM

This section provides procedures to configure SPBM IP Shortcuts using Enterprise Device Manager (EDM).

### Configuring IS-IS redistribution parameters using EDM


Use the following procedure to configure IS-IS redistribution parameters.

#### Procedure

1. In the navigation tree, expand the following folder: **Configuration>IS-IS** .
2. Click **IS-IS**.
3. In the work area, click the **Redistribute** tab.
4. On the toolbar, click **Insert**.
5. Choose the route source protocol in the **RouteSource** field.
6. Enable or disable the redistribution entry in the **Enable** field.
7. Click **Insert**.

#### Redistribute Tab Field Descriptions

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

| Name               | Description                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RouteSource</b> | Specifies the route source protocol for redistribution. <ul style="list-style-type: none"> <li>• direct</li> <li>• static</li> </ul>                                                                                                                                                            |
| <b>Enable</b>      | Enable or disable the redistribution entry. <ul style="list-style-type: none"> <li>• enable</li> <li>• disable</li> </ul>                                                                                                                                                                       |
| <b>Metric</b>      | Specifies the metric for the redistributed route. The range is from 0 to 65535 and the default value is 0. <p> <b>Note:</b><br/>You should use a value that is consistent with the destination protocol.</p> |
| <b>MetricType</b>  | Specifies the metric type.                                                                                                                                                                                                                                                                      |
| <b>Subnets</b>     | Indicates whether the subnets are advertised individually or aggregated to their classful subnet.                                                                                                                                                                                               |

## 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.

### \* 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 ports.
- You must enable IP Multicast over Fabric Connect globally.
- If there is no IP interface on the VLAN you must 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. From the navigation tree, double-click **Configuration**.
2. In the Configuration tree, click **VLAN**.
3. Click **VLANs**.
4. Choose a VLAN, and then click the **SpbMcast** 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 for EDM and CLI.

5. To enable or disable IP multicast over SPB within the GRT on the VLAN, click **Enable** or **Disable**.
6. Click **Apply**.

## Configuring ECMP paths using EDM

### Procedure

1. From the navigation tree, double-click **Configuration**.
2. In the Configuration tree, click **IP**.
3. Click **IP**.
4. In the double-click the cell in the **MaxPth** column.
5. Click **Apply**.

### ECMP Tab Field Descriptions

Use the data in the following table to describe the fields for **ECMP** tab.

| Name                    | Description                                                                                                                              |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Routing Protocol</b> | Specifies the switch protocol. <ul style="list-style-type: none"> <li>• static</li> <li>• rip</li> <li>• ospf</li> <li>• isis</li> </ul> |
| <b>MaxPath</b>          | Specifies the maximum number of ECMP paths. Default is 1.                                                                                |

# Glossary

|                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Address Resolution Protocol (ARP)</b>    | Maps an IP address to a physical machine address, for example, maps an IP address to an Ethernet media access control (MAC) address.                                                                                                                                                                                                                                                                                                                                   |
| <b>Autonomous System (AS)</b>               | 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.                                                                                                                                                                                                                                                          |
| <b>Bridge Protocol Data Unit (BPDU)</b>     | A data frame used to exchange information among the bridges in local or wide area networks for network topology maintenance.                                                                                                                                                                                                                                                                                                                                           |
| <b>Bridging</b>                             | A forwarding process, used on Local Area Networks (LAN) and confined to network bridges, that works on Layer 2 and depends on the Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP). Bridging is also known as MAC forwarding.                                                                                                                                                                                                                       |
| <b>CLI</b>                                  | Command Line Interface (CLI) is a text-based, common command line interface used for device configuration and management across Extreme Networks products.                                                                                                                                                                                                                                                                                                             |
| <b>CLI modes</b>                            | Differing command modes are available within the text-based interface, dependant on the level of user permissions determined by logon password. Each successive mode level provides access to more complex command sets, from the most restrictive—show level only, to the highest configuration levels for routing parameters, interface configuration, and security.                                                                                                 |
| <b>cyclic redundancy check (CRC)</b>        | Ensures frame integrity is maintained during transmission. The CRC performs a computation on frame contents before transmission and on the receiving device. The system discards frames that do not pass the CRC.                                                                                                                                                                                                                                                      |
| <b>Designated Intermediate System (DIS)</b> | 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. |
| <b>designated router (DR)</b>               | 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.

**Fabric Attach (FA)**

A feature used to extend the fabric edge to devices that do not have full SPBM support. Fabric Attach also decreases the configuration requirements on the SPBM devices by off-loading some configuration to the attached non-SPBM devices and by automating certain configuration steps that occur most often.

**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.

**Internet Control Message Protocol (ICMP)**

A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.

**Internet Protocol version 4 (IPv4)**

The protocol used to format packets for the Internet and many enterprise networks. IPv4 provides packet routing and reassembly.

**Internet Protocol version 6 (IPv6)**

An improved version of the IP protocol, IPv6 improves the IPv4 limitations of security and user address numbers.

**Layer 2**

Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.

**Layer 3**

Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).

**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).

**media**

A substance that transmits data between ports; usually fiber optic cables or category 5 unshielded twisted pair (UTP) copper wires.

|                                                              |                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Media Access Control (MAC)</b>                            | Arbitrates access to and from a shared medium.                                                                                                                                                                                                                                                                                                           |
| <b>Message Digest 5 (MD5)</b>                                | A one-way hash function that creates a message digest for digital signatures.                                                                                                                                                                                                                                                                            |
| <b>MultiLink Trunking (MLT)</b>                              | 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.                                                                                                |
| <b>multiple spanning tree instance (MSTI)</b>                | One of a number of spanning trees calculated by the Multiple Spanning Tree Protocol (MSTP) within an MST region, to provide a simple and fully connected active topology for frames that belong to a VLAN mapped to the MSTI.                                                                                                                            |
| <b>Open Shortest Path First (OSPF)</b>                       | A link-state routing protocol used as an Interior Gateway Protocol (IGP).                                                                                                                                                                                                                                                                                |
| <b>operation, administration, and maintenance (OA&amp;M)</b> | All the tasks necessary for providing, maintaining, or modifying switching system services.                                                                                                                                                                                                                                                              |
| <b>port</b>                                                  | A physical interface that transmits and receives data.                                                                                                                                                                                                                                                                                                   |
| <b>Protocol Data Units (PDUs)</b>                            | 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.                                                                                                                                                                         |
| <b>request for comments (RFC)</b>                            | A document series published by the Internet Engineering Task Force (IETF) that describe Internet standards.                                                                                                                                                                                                                                              |
| <b>routing switch</b>                                        | Virtualizes the physical router interfaces to switches. A virtual router port, or interface, acts as a router port to consolidate switching and routing functions in the broadcast domain, or between broadcast domains, and enable IP routing for higher traffic volumes.                                                                               |
| <b>shortest path first (SPF)</b>                             | 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.                                                                                                                                                                    |
| <b>spanning tree</b>                                         | A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function. |

|                                          |                                                                                                                                                                                                                          |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Spanning Tree Protocol (STP)</b>      | MAC bridges use the STP to exchange information across Local Area Networks to compute the active topology of a bridged Local Area Network in accordance with the Spanning Tree Protocol algorithm.                       |
| <b>Split MultiLink Trunking (SMLT)</b>   | 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.                                                |
| <b>stack</b>                             | Stackable Extreme Networks Ethernet Routing Switch can be connected in a stack configuration of two or more units, up to eight units maximum. A switch stack operates and is managed as a single virtual switch.         |
| <b>time-to-live (TTL)</b>                | 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.                                                      |
| <b>trunk</b>                             | A logical group of ports that behaves like a single large port.                                                                                                                                                          |
| <b>Virtual Local Area Network (VLAN)</b> | 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.                              |
| <b>Virtual Private Network (VPN)</b>     | 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. |