



Configuring Avaya Fabric Connect on Avaya Ethernet Routing Switch 4800 Series

Release 5.8
NN47205-507
Issue 02.03
November 2016

Notice

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

Documentation disclaimer

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

Link disclaimer

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

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: <https://support.avaya.com/helpcenter/getGenericDetails?detailId=C20091120112456651010> under the link "Warranty & Product Lifecycle" or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya.

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

Hosted Service

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

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

Licenses

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

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

License types

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

Heritage Nortel Software

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

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the

documentation, Hosted Service, and the product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Virtualization

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

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Avaya’s website at: <https://support.avaya.com/Copyright> or such successor site as designated by Avaya. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

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

Service Provider

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

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

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

Compliance with Laws

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

Preventing Toll Fraud

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

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: <https://support.avaya.com> or such successor site as designated by Avaya.

Security Vulnerabilities

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

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

Downloading Documentation

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

Contact Avaya Support

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

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

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

Contents

Chapter 1: Introduction	7
Purpose.....	7
Chapter 2: New in this release	8
Features.....	8
Other changes.....	8
Chapter 3: SPBM and IS-IS fundamentals	9
MAC-in-MAC encapsulation.....	10
I-SID.....	10
BCBs and BEBs.....	11
Basic SPBM network topology.....	11
IS-IS.....	13
Standard TLVs.....	13
IS-IS hierarchies.....	15
IS-IS PDUs.....	15
IS-IS configuration parameters.....	16
SPBM B-VLAN.....	18
Pre-populated FIB.....	19
RPFC.....	19
SPBM FIB.....	20
Fabric Attach.....	21
Chapter 4: SPBM and IS-IS infrastructure configuration using ACLI	30
Configuring minimum SPBM and IS-IS parameters.....	30
Displaying global SPBM parameters.....	36
Displaying global IS-IS parameters.....	37
Displaying IS-IS areas.....	39
Configuring optional SPBM parameters.....	40
Configuring optional IS-IS global parameters.....	42
Configuring optional IS-IS interface parameters.....	46
Displaying IS-IS interface parameters.....	48
Displaying the multicast FIB, unicast FIB, and unicast tree.....	51
Displaying IS-IS LSDB and adjacencies.....	53
Displaying IS-IS statistics and counters.....	57
Fabric Attach configuration.....	60
Activating FA Proxy mode.....	60
Activating FA Server mode.....	60
Displaying FA-specific settings.....	61
Displaying FA elements.....	61
Displaying FA I-SID/VLAN assignment data.....	61
Displaying FA per-port settings.....	62

Configuring per-port FA settings.....	63
Creating an I-SID/VLAN assignment on an FA Proxy.....	63
Deleting an I-SID/VLAN assignment on an FA Proxy.....	64
Enabling FA Proxy external client proxy support.....	64
Disabling FA Proxy external client proxy support.....	65
Configuring Auto Attach support.....	65
Configuring the FA authentication key.....	68
Enabling FA message authentication support.....	68
Disabling FA message authentication support.....	69
Chapter 5: SPBM and IS-IS infrastructure configuration using EDM.....	70
Configuring required SPBM and IS-IS parameters.....	70
Displaying the SPBM I-SID information.....	74
Displaying Level 1 Area information.....	75
Enabling or disabling SPBM globally.....	76
Configuring SPBM parameters.....	76
Displaying SPBM nicknames.....	77
Configuring interface SPBM parameters.....	78
Configuring SPBM on an interface.....	79
Displaying the unicast FIB.....	79
Displaying the multicast FIB.....	80
Displaying LSP summary information.....	81
Displaying IS-IS adjacencies.....	82
Configuring IS-IS globally.....	83
Configuring system level IS-IS parameters.....	85
Configuring IS-IS interfaces.....	85
Configuring IS-IS interface level parameters.....	87
Configuring an IS-IS Manual Area.....	88
Displaying IS-IS system statistics.....	89
Displaying IS-IS interface counters.....	90
Displaying IS-IS interface control packets.....	90
Fabric Attach configuration.....	91
Configuring Fabric Attach.....	91
I-SID configuration.....	92
Configuring per-port FA settings.....	93
Chapter 6: Layer 2 VSN configuration fundamentals.....	95
SPBM L2 VSN.....	95
SPBM L2 VSN sample operation.....	97
Chapter 7: Layer 2 VSN configuration using ACLI.....	103
Configuring a SPBM Layer 2 VSN C-VLAN.....	103
Configuring a SPBM Layer 2 VSN Switched UNI.....	104
Displaying C-VLAN and Switched UNI I-SID information.....	106
Managing the switch via Layer 2.....	109
Chapter 8: Layer 2 VSN configuration using EDM.....	110

Configuring SPBM Layer 2 VSN C-VLANs.....	110
Displaying the MAC address table for a C-VLAN.....	111
Configuring SPBM switched UNIs.....	112
Managing the switch via Layer 2.....	112
Chapter 9: CFM fundamentals.....	114
MD.....	114
MA.....	115
MEP.....	116
Fault verification.....	117
LBM.....	117
Layer 2 ping.....	117
Fault isolation.....	118
LTM.....	118
Layer 2 traceroute.....	119
Layer 2 tracetree.....	119
MIP.....	119
Nodal MPs.....	120
Configuration considerations.....	120
Chapter 10: CFM configuration using ACLI.....	121
Configuring CFM.....	121
Triggering an LBM Layer 2 ping.....	123
Triggering an LTM Layer 2 traceroute.....	124
Triggering an LTM Layer 2 tracetree.....	126
Chapter 11: CFM configuration using EDM.....	127
Configuring CFM.....	127
Displaying CFM MD.....	128
Displaying CFM MA.....	129
Displaying CFM MEP.....	130
Configuring Layer 2 ping.....	131
Initiating a Layer 2 traceroute.....	133
Viewing Layer 2 traceroute results.....	135
Chapter 12: Resources.....	137
Support.....	137
Searching a documentation collection.....	138
Subscribing to e-notifications.....	139
Glossary.....	142

Chapter 1: Introduction

Purpose

This document provides instructions to configure Avaya VENA Fabric Connect on the Ethernet Routing Switch 4800 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).

Using the document

The document is organized into feature sections:

1. Infrastructure configuration — You must first configure your base SPB and IS-IS architecture described in the infrastructure configuration chapters. The chapter includes initial steps to configure the minimum SPB and IS-IS parameters to enable Fabric Connect on your network, and additional steps to configure optional SPB and IS-IS parameters. For more information, see [Configuring minimum SPBM and IS-IS parameters](#) on page 30
2. Services configuration — After you have completed the infrastructure configuration, you configure the appropriate services for your network to run on top of your base architecture. Services can include: Layer 2 VSNs, IP Shortcut Routing, Layer 3 VSNs, and Inter-VSN Routing.
3. Operations and management — Finally, Ethernet Routing Switch 4800 series provides tools to monitor and troubleshoot your Fabric Connect network.

The document also includes configuration examples at the end of each chapter to show basic configurations to implement Fabric Connect technology.

Chapter 2: New in this release

The following sections detail what is new in *Configuring Avaya Fabric Connect on Avaya Ethernet Routing Switch 4800 Series*, NN47205-507 for Release 5.8.

Features

See the following sections for information about feature changes:

 **Note:**

Release 5.8 features are supported only on ERS 4800 series.

Fabric Attach

With the Fabric Attach feature, you can extend the fabric edge to devices that do not have full Shortest Path Bridging - MAC (SPBM) support. Fabric Attach allows non-SPBM devices to take advantage of full SPBM support if it is available.

For more information about Fabric Attach, see the following:

- [Fabric Attach fundamentals](#) on page 21
- [Configuring Fabric Attach](#) on page 60

Other changes

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

Document title change

Configuring Avaya VENA Fabric Connect on Avaya Ethernet Routing Switch 4000 Series is renamed *Configuring Avaya Fabric Connect on Avaya Ethernet Routing Switch 4800 Series*.

Introduction chapter

Information about Related resources and Support are moved to the last chapter in this document.

Chapter 3: 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).

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.

Avaya ERS 4800 Series 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:

ERS 4800 devices are currently meant to function strictly as BEB devices.

Basic SPBM network 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 L2 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.

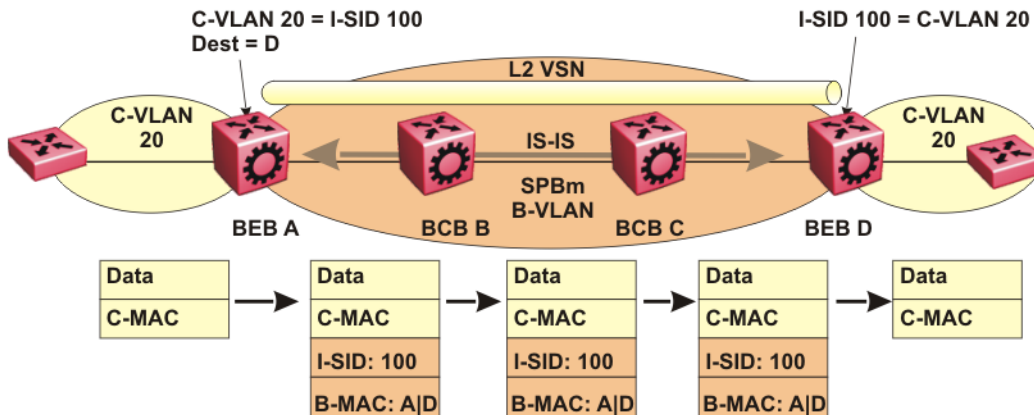


Figure 1: SPBM L2 VSN

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

The BEBs provide the boundary between the SPBM domain and the virtualized services domain. For a Layer 2 VSN service, the BEBs map a C-VLAN to an I-SID based on local service provisioning. Any BEB in the network that has the same I-SID configured can participate in the same Layer 2 VSN. 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. ERS 4800 also supports Sub-TLVs. TLVs exist inside IS-IS packets and Sub-TLVs exist as additional information in TLVs.

Avaya ERS 4800 Series 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. Avaya is in full compliance with the IEEE 802.1aq standard. The following table lists the TLVs that the ERS 4800 Series 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.
143	Multi-topology port aware capability (MT-Port-Capability) TLV This TLV carries the SPB instance ID in a multiple SPB instances environment. This TLV is carried within IS-IS Hello Packets (IIH).	This TLV carries the following SPBM Sub TLVs: <ul style="list-style-type: none"> • MCID Sub TLV: The MCID is a digest of the VLANs and MSTI. Neighboring SPBM nodes must agree on the MCID to form an adjacency. In the current release, the MCID is set to all zeros (0). After the ERS 4800 receives a none-zero MCID Sub TLV, it reflects content back to the neighbor. • SPB B-VID Sub TLV (type 6): The Sub TLV indicates the mapping between a VLAN and its equal cost tree (ECT) algorithm. To form an adjacency, both nodes must have a matching primary (BVLAN, ECT) pair, and secondary (BVLAN, ECT) pair.

Table continues...

TLV	Description	Usage
144	<p>Multi-topology Capability (MT-Capability) TLV.</p> <p>This TLV carries the SPB instance ID in a multiple SPB instance environment. This TLV is carried within LSPs.</p>	<p>This TLV carries the following Sub TLVs:</p> <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.
184	<p>SPBM IP VPN reachability — IS-IS TLV 184 is used to advertise SPBM L3 VSN route information across the SPBM cloud.</p>	<p>IP reachability for Layer 3 VSNS</p>

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 ERS 4800 release only supports one manual area.
- **NSEL** — The last byte (00) is the n-selector. In the ERS 4800 implementation, this part is automatically attached. There is no user input accepted.

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

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

The following are the requirements for system IDs:

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

PSNP interval

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

CSNP periodic and interval rate

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

Parameters for the link state packet (LSP)

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

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

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

Point-to-point mode

All SPBM links are point-to-point links. ERS 4800 does not support broadcast links.

IS-IS interface authentication

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

You can use either one of the following authentication methods:

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

Password considerations

The passwords for all authentications are saved as cleartext in the configuration file on the Avaya Ethernet Routing Switch 4800. The passwords for simple and HMAC-MD5 are displayed in cleartext through ACLI. The HMAC-MD5 packet is encrypted when transmitted over the network.

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

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

Hellos

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

- Interface Hello interval
- Interface Hello multiplier

Interface Hello interval

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

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

Hello multiplier

You can configure the Hello multiplier to specify how many Hellos the Avaya Ethernet Routing Switch 4800 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 Avaya Ethernet Routing Switch 4800 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. ERS 4800 devices are currently meant to function strictly as BEB devices, and will always set the overload bit (this option cannot be modified on this release), and cannot work as transit devices (BCBs).

When an LSP with an overload bit is received from a neighboring transit-capable SPBM device, the Avaya ERS 4800 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 true on ERS4800 devices, and cannot be modified in this release.

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 Multi-Link Trunking on Avaya Ethernet Routing Switch 4800 Series*, NN47205-501.

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.

```
4850GTS-PWR+(config-if)#show isis spbm unicast-fib
=====
                        SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION          BVLAN  SYSID          HOST-NAME      OUTGOING      COST
ADDRESS
-----
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.

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

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

The following text shows an example of the multicast FIB.

```
4850GTS-PWR+(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
-----
```

Fabric Attach

With the Fabric Attach (FA) feature, you can extend the fabric edge to devices that do not have full SPBM support. FA allows non-SPBM devices to take advantage of full SPBM support if it is available.

FA 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. This feature functions in both stand-alone and stacked configurations. For ease of use, existing SPBM commands are leveraged and replicated on non-SPBM devices.

FA Signaling

The FA elements communicate between themselves using FA Signaling. FA Signaling is Avaya's application level protocol that leverages standard network protocols (such as LLDP, 802.1x, RADIUS, and Web Services) 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.
- **FA Proxy**—A non-SPB network device running the FA agent in FA Proxy 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, such as laptops, IP phones, printers or IP cameras, will also be supported in a later release.

Note:

FA Client functionality in Release 5.8 is a technology demonstration feature only.

The ERS 4800 can act as either an FA Proxy or FA Server, with the global SPBM status determining whether the device is a FA Proxy (SPBM disabled) or FA Server (SPBM enabled).

FA Element Discovery

An FA agent which controls FA functionality resides on all FA-capable devices (FA Server, FA Proxy or FA Client). The agent executes as a normal priority task and no agent-specific configuration is necessary.

FA Proxy and FA Server elements control FA through a global FA service 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 agent determines whether FA data (I-SID/VLAN assignments) is exported to an FA Proxy that acts as an external client proxy or an FA Server.

On ERS 4800, the global FA service is always enabled. Per-port settings are, by default, enabled on FA Proxies and disabled on FA Servers.

*** 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 FA Signaling 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

The switch maintains FA service status, FA per-port settings, Auto Attach status, external client proxy status, message authentication requirements and previously configured I-SID/VLAN assignments in the non-volatile storage and restores them during the agent initialization sequence. In a stack environment, FA agent startup and initialization occurs on every unit in the stack using the data restored from non-volatile storage.

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 Server before a system reset.

FA Clients

FA Clients connect to an FA Proxy through standard, non MAC-in-MAC access ports, advertising configured I-SID/VLAN requests that they would like supported by 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 may 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 and tagging of the FA Server uplink port if necessary.

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:**

FA Clients are not available yet at the time of the development of this document. The above description of the FA Clients is made for clarity of the Avaya Fabric Attach architecture.

FA Proxy I-SID/VLAN assignment configuration

I-SID/VLAN binding data is typically configured by an administrator on an FA Proxy. 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.

I-SID/VLAN bindings can also be received by an FA Proxy from FA Clients. If external client proxy support is enabled, FA Client bindings are handled the same way that locally-configured bindings are handled when it comes to standard processing requirements. Processing differs (following FA Server binding acceptance) for client bindings in that the FA Proxy/Client downlink is conditioned in a client-specific way as well.

I-SID/VLAN assignment configuration data is validated before it is accepted. Validation includes checks for:

- VLAN presence, type, and status.
- Duplicate VLAN definitions. Duplicate binding definitions are allowed across multiple FA Clients and the FA Proxy.
- VLAN and I-SID values that exceed acceptable ranges.
- Resource exhaustion, such as LLDP limits reached.

*** Note:**

Data exchanges (I-SID/VLAN assignments) between an FA Proxy and an FA Server or FA Client are supported, as are exchanges between an FA Server and an FA Proxy or FA Client. FA Proxy to FA Proxy and FA Server to FA Server interactions are not supported in Release 5.8.

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. Specifically, an uplink port 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) clear the port VLAN membership when the related I-SID/VLAN assignment is rejected by the FA Server or deleted by the FA Proxy administrator. Once the FA agent establishes the port tagging status, it remains in effect regardless of I-SID/VLAN assignment status.

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.
- Update port VLAN tagging status to ensure egress traffic is tagged.

FA Server I-SID/VLAN assignment addition actions:

- Create SPBM switched UNI VLAN corresponding to I-SID/VLAN assignment VLAN.
- Update downlink port VLAN tagging status to ensure egress traffic is tagged. Tagging status for FA client connections is determined by the client type.
- 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.
- Update downlink port VLAN ID (PVID) for untagged client connections (if a valid default VLAN in the range of 1-4094 was specified by the client and if the VLAN in the binding being activated equals the specified default VLAN).

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.
- Update downlink port VLAN tagging status based on the FA Client type (tagged – ‘tagAll’/ untagged – ‘untagPvidOnly’).
- Update downlink port VLAN ID (PVID) for untagged FA Clients (if a valid default VLAN [1..4094] was specified by the client and if the VLAN in the binding being activated equals the specified default 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 port VLAN membership to exclude I-SID/VLAN assignment VLAN.

*** Note:**

The FA Proxy deletes port-based VLANs created during binding activation if the VLAN is not associated with other ports.

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.
- Default downlink port VLAN ID (PVID) for untagged clients. Downlink port VLAN tagging status remains unchanged.

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.
- Default downlink port VLAN ID (PVID) for untagged clients.

*** Note:**

The FA Proxy deletes port-based VLANs created during binding activation if the VLAN is not associated with other ports.

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.
- 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 pre-determined 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 pre-determined 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.

FA Proxy/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.

Primary server failure is detected using a capabilities advertisement timeout. Once a predefined period of time without FA Server Signaling 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 connections to the same FA server are treated as a single logical connection by the FA Proxy. 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 stand-alone 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) in order 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 - FA Server interface leaves the stack, any I-SID/VLAN assignments accepted by the server are immediately aged-out. I-SID/VLAN assignment data is defaulted (reverts to the '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.

FA Auto Attach

The Auto Attach (AA) feature allows a device to extract management VLAN data from the primary FA Server advertisements and use this data to update the in-use management VLAN and initiate IP address acquisition using DHCP. This information can be cascaded to FA Clients as well

Auto Attach feature control is provided for situations when you prefer or require manual configuration of the settings affected by Auto Attach. You can enable or disable the Auto Attach functionality on an FA Proxy or Server. Auto Attach is enabled by default on FA Proxies and disabled on FA Servers.

Auto Attach must be enabled on the FA Server or FA Proxy before management VLAN information is included in the FA signaling generated by the element. If Auto Attach is disabled, a value indicating that the management VLAN data is not being exported is distributed.

Upon FA Signaling receipt, if AA is enabled and the received management VLAN data is not the same as the currently configured management VLAN, the FA Proxy (and possibly FA Client) reconfigures the management VLAN and potentially initiates IP address acquisition using DHCP. No Auto Attach-specific actions are taken based on received FA Signaling data if Auto Attach is disabled or if the specified management VLAN is not valid (that is, not in the range 1 to 4094).

Auto Attach operation is coupled with FA operation. Even though Auto Attach functionality can be enabled and disabled separately from FA, Auto Attach relies on data that is only available during FA Proxy/Server exchanges and, specifically, once a primary FA Server has been selected. An FA Client can choose to utilize FA-provided management VLAN data once an FA Proxy or FA Server have been discovered as well.

Auto Attach is active on an FA Proxy if Auto Attach is enabled and a primary FA server is selected. On an FA Server, Auto Attach is active if Auto Attach is enabled and FA Proxies or FA Clients have been discovered.

If the management VLAN advertised by the primary FA server differs from the management VLAN currently configured on the FA Proxy, Auto Attach initiates the following actions, if required:

- VLAN creation on an FA Proxy—if the FA server-specified (proposed) management VLAN does not exist on the FA Proxy, Auto Attach creates a port-based VLAN.
- Management VLAN update on an FA Proxy—the designated management VLAN for the device is updated. The proposed management VLAN becomes the designated management VLAN for the device. No operations related to the previous management VLAN (for example port membership updates or VLAN deletion) are performed.
- Port VLAN membership update on an FA Proxy or Server—the uplink port through which the primary FA Server is accessed must be a member of the management VLAN for network accessibility. Auto Attach updates the port VLAN membership accordingly, if an update is required. The downlink port through which the FA Proxy/Client is accessed must also be a member of the management VLAN. Auto Attach updates the port VLAN membership for this interface as well.
- IP address acquisition on an FA Proxy—Auto Attach initiates IP address acquisition through DHCP if an IP address has not already been manually configured. The operation performed is equivalent to the action initiated through the `ip address source dhcp-when-needed` command.

*** Note:**

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

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

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 the I-SID/VLAN assignment data can be used to protect the FA Proxy/Server I-SID/VLAN assignment 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.

On secure (SSH) images, you can enable or disable FA message authentication. On non-secure images, message authentication cannot be enabled.

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 FA Signaling. Upon receipt, the HMAC digest is recomputed and compared against the digest included in FA Signaling. 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.

Chapter 4: SPBM and IS-IS infrastructure configuration using ACLI

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

Configuring minimum SPBM and IS-IS parameters

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

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, Avaya recommends that you change the IS-IS system ID from the default B-MAC value to a recognizable address to easily identify a switch (Log on to IS-IS Router configuration mode and use the `system-id <xxxx.xxxx.xxxx>` command) . This helps to recognize source and destination addresses for troubleshooting purposes.

11. Configure an IS-IS manual area (1-13 bytes in the format <xx.xxxx.xxxx...xxx>. In this release, 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.

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

```
4850GTS-PWR+> enable
4850GTS-PWR+# configure terminal
4850GTS-PWR+(config)# spbm
4850GTS-PWR+(config)# router isis
4850GTS-PWR+(config-isis)# spbm 1
4850GTS-PWR+(config-isis)# exit
4850GTS-PWR+(config)# vlan create 1000 type spbm-bvlan
4850GTS-PWR+(config)# vlan create 2000 type spbm-bvlan
4850GTS-PWR+(config)# router isis
4850GTS-PWR+(config-isis)# spbm 1 b-vid 1000,2000 primary 1000
4850GTS-PWR+(config-isis)# spbm 1 nick-name 1.11.16
4850GTS-PWR+(config-isis)# manual-area c0.2000.0000.0000
```

```

4850GTS-PWR+(config-isis)# exit
4850GTS-PWR+config# interface Ethernet 3
4850GTS-PWR+(config-if)# isis
4850GTS-PWR+(config-if)# isis spbm 1
4850GTS-PWR+(config-if)# isis enable
4850GTS-PWR+(config-if)# exit
4850GTS-PWR+(config)# vlan member remove 1 3
4850GTS-PWR+(config)# router isis enable
4850GTS-PWR+(config)# show isis spbm

```

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

```
4850GTS-PWR+(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

```
4850GTS-PWR+(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. 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. Use the no option to delete the manual area.

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

Variable	Value
<1–100>	Creates the SPBM instance. In this release, only one SPBM instance is supported.
b-vid {<vlan-id [-vlan-id] [...]}]	Sets the ISIS SPBM instance data VLANs. Use the no option to remove the specified B-VLAN from the SPBM instance.
nick-name <x.xx.xx>	Specifies a nickname for the SPBM instance globally. The value is 2.5 bytes in the format <x.xx.xx>. Use the no or default options to delete the configured nickname.
primary <1–4094>	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-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"> port — port-based protocol-decEther2 — Specify a decEther2 protocol-based VLAN. protocol-ipEther2 — Specify an ipEther2 protocol-based VLAN.

Table continues...

Variable	Value
	<ul style="list-style-type: none"> • protocol-ipv6Ether2 — Specify an ipv6Ether2 protocol-based VLAN. • protocol-ipx802.2 — Specify an ipx802.2 protocol-based VLAN. • protocol-ipx802.3 — Specify an ipx802.3 protocol-based VLAN. • protocol-ipxEther2 — Specify an ipxEther2 protocol-based VLAN. • protocol-ipxSnap — Specify an ipxSnap protocol-based VLAN. • protocol-Netbios — Specify a NetBIOS protocol-based VLAN. • protocol-RarpEther2 — Specify a RarpEther2 protocol-based VLAN. • protocol-sna802.2 — Specify a sna802.2 VLAN. • protocol-snaEther2 — Specify an snaEther2 protocol-based VLAN. • protocol-Userdef — Specify a user-defined protocol-based VLAN. Enter optional parameters. <ul style="list-style-type: none"> - all – display all Userdef VLANs - ether – display Ethernet II Userdef VLANs - llc – display LLC Userdef VLANs • protocol-vinesEther2 — Specify a vinesEther2 protocol-based VLAN. • protocol-xnsEther2 — Specify an xnsEther2 protocol-based VLAN. • spbm-bvlan — Specify an SPBM-BVLAN. • spbm-switchedUni — Specify an SPBM-switchedUni • voice-vlan — Specify voice VLAN information

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 ACLI to enter User EXEC mode.
2. At the command prompt, enter the following command to check if SPBM is enabled:

```
show spbm
```

3. At the command prompt, enter the following command:

```
show isis spbm
```

4. 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

```
4850GTS-PWR+(config)#show spbm
SPBM Global: Disabled
SPBM Ethertype: 0x8100
```

```
=====
                        ISIS SPBM Info
=====
```

SPBM INSTANCE	B-VID	PRIMARY VLAN	NICK NAME	LSDB TRAP
1	1000,2000	1000	1.11.16	disable

```
=====
```

```
4850GTS-PWR+# 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.

Displaying global IS-IS parameters

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

Procedure

1. Log on to ACLI 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

```

4850GTS-PWR+#show isis
=====
                ISIS General Info
=====
AdminState : enabled
RouterType : Level 1
System ID  : 0000.0000.0000
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
Num of Interfaces : 2
Num of Area Addresses : 1

4850GTS-PWR+#show isis system-id
=====
                ISIS System-Id
=====
SYSTEM-ID
=====
0014.c7e1.33df

4850GTS-PWR+#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.

Table continues...

Parameter	Description
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.
Num of Interfaces	Indicates the number of interfaces on the router.
Num of Area Addresses	Indicates the number of area addresses on the router.

show isis system-id

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

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

show isis net

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

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

Displaying IS-IS areas

Use the following procedure to display IS-IS areas.

Procedure

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

```
show isis manual-area
```

Example

```
4850GTS-PWR+#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

```
4850GTS-PWR+> enable
4850GTS-PWR+# configure terminal
4850GTS-PWR+(config)# spbm ethertype 0x8100
4850GTS-PWR+(config-isis)# no router isis enable
4850GTS-PWR+(config)# router isis
4850GTS-PWR+(config-isis)# spbm 1 lsdB-trap enable
4850GTS-PWR+(config-isis)# router isis enable
4850GTS-PWR+(config-isis)# exit
4850GTS-PWR+(config)# interface ethernet 3
4850GTS-PWR+(config-if)# no isis enable
4850GTS-PWR+(config-if)# isis spbm 1 interface-type ptpt
4850GTS-PWR+(config-if)# isis spbm 1 l1-metric 500
4850GTS-PWR+(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 x88a8. 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 ptp	<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. Configure the Partial Sequence Number Packet (PSNP) in seconds:

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

- f. Configure the minimum time between retransmission of an LSP:

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

- g. Configure the SPF delay in milliseconds:

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

- h. Configure the name for the system:

```
sys-name WORD <1-255>
```

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

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

 **Note:**

The IS-IS “overload” bit is always set, and no NNI to NNI traffic is forwarded.

Example

```
4850GTS-PWR+> enable
4850GTS-PWR+# configure terminal
4850GTS-PWR+(config)# router isis
4850GTS-PWR+(config-isis)# csnp-interval 10
4850GTS-PWR+(config-isis)# is-type l1
4850GTS-PWR+(config-isis)# max-lsp-gen-interval 800
4850GTS-PWR+(config-isis)# metric wide
4850GTS-PWR+(config-isis)# psnp-interval 10
```



```
4850GTS-PWR+(config-isis)# retransmit-lsp-interval 10
4850GTS-PWR+(config-isis)# default sys-name
4850GTS-PWR+(config-isis)# spf-delay 200
4850GTS-PWR+(config-isis)# default system-id
```

Variable definitions

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

Variable	Value
<1-600>	Specifies the CSNP interval in seconds. This is a system level parameter that applies for level 1 CSNP generation on all interfaces. A longer interval reduces overhead, while a shorter interval speeds up convergence. The default value is 10. Use the no or default options to set this parameter to the default value of 10.

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

Variable	Value
{1}	Sets the router type globally: • l1: Level-1 router type The default value is l1. Use the no or default options to set this parameter to the default value of l1.

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

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

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

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

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

Variable	Value
<1-120>	<p>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.</p> <p>The default value is 2. Use the no or default options to set this parameter to the default value of 2.</p>


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

Variable	Value
<1-300>	<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 Level1 retransmission of LSPs.</p> <p>The default value is 5 seconds. Use the no or default options to set this parameter to the default value of 5.</p>

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

Variable	Value
<0-5000>	<p>Configures the delay, in milliseconds, to pace successive Shortest Path First (SPF) runs. The timer prevents more than two SPF runs from being scheduled back-to-back. The mechanism for pacing SPF allows two back-to-back SPF runs.</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
WORD<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> Note:</p> <p>In this release, no consistency checks appear when you edit sys-name on ERS 4800.</p>

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

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

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


```
4850GTS-PWR+> enable
4850GTS-PWR+# configure terminal
4850GTS-PWR+(config)# interface ethernet 3
4850GTS-PWR+(config-if)# isis
4850GTS-PWR+(config-if)# isis hello-auth type hmac-md5 key test
4850GTS-PWR+(config-if)# isis l1-dr-priority 100
4850GTS-PWR+(config-if)# isis l1-hello-interval 20
4850GTS-PWR+(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 <i>{none simple hmac-md5}</i> [key[key WORD<1-16>] [key-id <1-255>]]	Specifies the authentication type used for IS-IS hello packets on the interface. type can be one of the following: <ul style="list-style-type: none"> • none

Table continues...

Variable	Value
	<ul style="list-style-type: none"> • simple: If selected, you must also specify a key value but the key id is optional. Simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5: If selected, you must also specify a key value but the key-id is optional. MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. There is an optional key ID. <p>The default is none. Use the no or default options to set the hello-auth type to none.</p>
l1-dr-priority <0–127>	<p>Configures the level 1 IS-IS designated router priority to the specified value. The default value is 64.</p> <p>Use the no or default options to set this parameter to the default value of 64.</p> <p> Note:</p> <p>This parameter is not used for SPBM because SPBM only runs on point-to-point interfaces. This parameter is for 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 ACLI to enter User EXEC mode.
2. Display IS-IS interface configuration and status parameters (including adjacencies):


```
show isis interface [l1]
```
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

```
4850GTS-PWR+#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

4850GTS-PWR+#show isis int-auth
=====
                        ISIS Interface Auth
=====
IFIDX      AUTH-TYPE  AUTH-KEYID  AUTH-KEY
-----
Trunk: 3   none      0           0
Port: 21   none      0           0

4850GTS-PWR+#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

4850GTS-PWR+#show isis int-ckt-level
=====
                        ISIS Circuit level parameters
=====
IFIDX      LEVEL      DIS          CKTID
-----
Trunk: 2   Level 1    1
Port: 21   Level 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.

Job aid

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

show isis interface

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

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

show isis int-auth

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

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

show isis int-timers

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

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

show isis int-ckt-level

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

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

Displaying the 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 ACLI 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

```
4850GTS-PWR+#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
=====
```

```
4850GTS-PWR+#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
=====
```

```
4850GTS-PWR+#show isis spbm unicast-tree 1000
Node:0018.b0bb.b3df.00 (4850GTS-PWR+) -> ROOT
Node:0016.ca23.73df.00 (4850GTS-PWR+) -> 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 <1-16777215>	Displays the FIB for the specified I-SID.
nick-name <x.xx.xx>	Displays the FIB for the specified nickname.
summary	Displays a summary of the FIB.

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

Variable	Value
b-mac <0x00:0x00:0x00:0x00:0x00:0x00>	Displays the FIB for the specified 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-INTERFACES	Indicates the multicast destination MAC address of the multicast FIB entry.
ISID	Indicates the I-SID of the multicast FIB entry.
BVLAN	Indicates the B-VLAN of the multicast FIB entry.
SYSID	Indicates the system identifier of the multicast FIB entry.
HOST-NAME	Indicates the host name of the multicast FIB entry.
OUTGOING	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 ACLI 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

```
4850GTS-PWR+#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
```

```
4850GTS-PWR+# 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
-----
```

```
4850GTS-PWR+>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)

```

Variable definitions

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

Variable	Value
level {1 2 12}	Displays the LSDB for the specified level: l1, l2, or l12. * Note: Level 1 is supported in this release.
sysid <xxxx.xxxx.xxxx>	Displays the LSDB for the specified system ID.
lspid <xxxx.xxxx.xxxx.xx-xx>	Displays the LSDB for the specified LSP ID.
tlv <1-184>	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.

Job aid

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

show isis lsdb

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

Parameter	Description
LSP ID	Indicates the LSP ID assigned to external IS-IS routing devices.
LEVEL	Indicates the level of the external router: L1, L2, or L12.
LIFETIME	Indicates the maximum age of the LSP. If the max-lsp-gen-interval is set to 900 (default) then the lifetime value begins to count down from 1200 seconds and updates after 300 seconds if connectivity remains. If the timer counts down to zero, the counter adds on an additional 60 seconds, then the LSP for that router is lost. This happens because of the zero age lifetime, which is detailed in the RFC standards.
SEQNUM	Indicates the LSP sequence number. This number changes each time the LSP is updated.
CHKSUM	Indicates the LSP checksum. This is an error checking mechanism used to verify the validity of the IP packet.
HOST-NAME	Indicates the hostname listed in the LSP. If the host name is not configured, then the system name is displayed.

show isis adjacencies

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

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

Displaying IS-IS statistics and counters

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

Procedure

1. Log on to ACLI 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. The current release does not support level 2 IS-IS. The ACLI 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

```
4850GTS-PWR+# show isis statistics
```

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

```
4850GTS-PWR+#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
```

```
4850GTS-PWR+#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.

Job aid

show isis statistics

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

Parameter	Description
LEVEL	Shows the level of the IS-IS interface.
CORR LSPs	Shows the number of corrupted LSPs detected.
AUTH FAILS	Shows the number of times authentication has failed on the global level.
AREA DROP	Shows the number of manual addresses dropped from the area.
MAX SEQ EXCEEDED	Shows the number of attempts to exceed the maximum sequence number.
SEQ NUM SKIPS	Shows the number of times the sequence number was skipped.
OWN LSP PURGE	Shows how many times the local LSP was purged.
BAD ID LEN	Shows the number of ID field length mismatches.
PART CHANGES	Shows the number of partition link changes.
LSP DB OLOAD	Show the number of times the ERS 4800 was in the overload state.

show isis int-counters

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

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

Table continues...

Parameter	Description
MAX AREA	Shows the maximum area address mismatches.
LAN DIS CHANGES	Shows the number of times the DIS has changed.

show isis int-l1-ctrl-pkts

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

Parameter	Description
IFIDX	Shows the interface index for the Ethernet or MLT interface.
DIRECTION	Shows the packet flow (Transmitted or Received).
HELLO	Shows the number of interface-level Hello packets.
LSP	Shows the number of LSP packets.
CSNP	Shows the number of CSNPs.
PSNP	Shows the number of PSNPs.

Fabric Attach configuration

The following sections describe how to configure Fabric Attach (FA) using ACLI.

Activating FA Proxy mode

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Activate FA Proxy mode:

```
no spbm
OR
default spbm
```

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. Activate FA Server mode and enable the FA service:
`spbm`

Displaying FA-specific settings

Use the following procedure to display FA-specific settings.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display FA-specific settings:
`show fa spbm`

Displaying FA elements

Use the following procedure to display discovered FA elements.

*** Note:**

On an FA Server, element tracking is only supported if Auto Attach is enabled.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display the FA elements:
`show fa elements`

Displaying FA I-SID/VLAN assignment data

Use the following command to display FA-specific I-SID/VLAN assignment data.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display FA specific I-SID/VLAN assignment data:

```
show fa i-sid
```

3. To display FA Proxy I-SID data or FA Server I-SID data, enter the following command:

```
show i-sid
```

*** Note:**

On an FA Proxy, each configured I-SID/VLAN association creates a C-VLAN UNI once the assignment becomes active following acceptance by an FA Server.

On an FA Server, only switched UNIs are created as the result of FA Proxy/FA Server interaction.

Example

ERS 4800 acting as an FA Proxy:

```
4850GTS-PWR+(config)#show i-sid
I-SID      Vid  UNI-type  Ports
-----
500        5    C-VLAN    2/91
600        6    C-VLAN    2/91,3/1
13849     138  C-VLAN    2/91
16000000  1000 C-VLAN    2/91
```

ERS 4800 acting as an FA Server:

```
4850GTS-PWR+(config)#show i-sid
I-SID      Vid  UNI-type  Ports
-----
500        5    switched  1/21
600        6    switched  1/21
13849     138  switched  1/21
16000000  1000 switched  1/21
```

Displaying FA per-port settings

Use the following command to display FA per-port settings.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display FA per-port settings:

```
show fa port-enable
```

Configuring per-port FA settings

Use the following procedure to configure per-port FA operation.

About this task

Use the `fa port-enable` command to determine whether FA data is included in FA Signaling.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure per-port FA operation:

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

Variable Definitions

Use the data in the following table to use the `fa port-enable` command.

Variable	Description
<code>fa port-enable <portlist></code>	Enables FA operation on specified ports.
<code>[no]</code>	Disables per-port FA operation.
<code>default fa port-enable</code>	Restores per-port FA settings to default. The current operational mode determines the default setting: <ul style="list-style-type: none"> • enabled, for FA Proxy mode • disabled, for FA Server mode
<code><portlist></code>	Specifies a port or list of ports for which to enable or disable FA operation.

Creating an I-SID/VLAN assignment on an FA Proxy

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

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create an association between an I-SID and a VLAN:

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

Example

```
4850GTS-PWR+(config)#i-sid 600 vlan 6
4850GTS-PWR+(config)#show fa i-sid
```

I-SID	VLAN	Status
600	6	Pending

Variable Definitions

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

Variable	Description
<code>i-sid <1-16777214> vlan<1-4094></code>	Specifies the VLAN to associate with the I-SID.

Deleting an I-SID/VLAN assignment on an FA Proxy

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

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Remove the I-SID from the specified VLAN:


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

OR

Remove all I-SID/VLAN assignments:

```
default i-sid
```

Variable Definitions

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

Variable	Description
<code>i-sid <1-16777214> vlan<1-4094></code>	Specifies the I-SID to remove from the specified VLAN.

Enabling FA Proxy external client proxy support

Use the following command to enable FA Proxy external client proxy support.

Procedure

1. Enter Global Configuration mode:

- ```
enable
configure terminal
```
2. Enable FA Proxy external client proxy support:

```
fa proxy
```

OR

```
default fa proxy
```

---

## Disabling FA Proxy external client proxy support

Use the following command to disable FA Proxy external client proxy support.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```
2. Disable FA Proxy external client proxy support:

```
no fa proxy
```

---

## Configuring Auto Attach support

Use the following procedure to configure Auto Attach on an FA Proxy or FA Server.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```
2. Configure Auto Attach:

```
[no] [default] fa auto-attach
```

## Variable Definitions

Use the data in the following table to use the `fa auto-attach` command.

| Variable                              | Description                                                                                                              |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <code>fa auto-attach</code>           | Enables Auto Attach support.                                                                                             |
| <code>[no] fa auto-attach</code>      | Disables Auto Attach support.                                                                                            |
| <code>[default] fa auto-attach</code> | Restores the default state for Auto Attach. By default, Auto Attach is enabled on FA Proxies and disabled on FA Servers. |

## Job Aid

The following job aid illustrates Auto Attach operations. FA Proxy/FA Server connectivity is established (FA Proxy/FA Server uplink on unit 2, port 91 and FA Server/FA Proxy downlink on unit 1, port 17). The components start in their FA default state (SPBM enabled on the FA Server and a 4800 unit set as FA Proxy).

### FA Proxy:

```
4850GTS-PWR+(config)#show fa spbm

Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Auto Attach Status: Enabled
Fabric Attach Message Authentication Status: Enabled
Fabric Attach Proxy Status: Enabled
Fabric Attach Primary Server Id: <none>
Fabric Attach Primary Server Descr: <none>
```

```
4850GTS-PWR+(config)#show fa i-sid

I-SID VLAN Source Status


```

```
4850GTS-PWR+(config)#show vlan mgmt
```

```
Management VLAN: 1
```

```
4850GTS-PWR+(config)#show vlan
```

| Id             | Name    | Type                            | Protocol | PID    | Active | IVL/SVL | Mgmt |
|----------------|---------|---------------------------------|----------|--------|--------|---------|------|
| 1              | VLAN #1 | Port                            | None     | 0x0000 | Yes    | IVL     | Yes  |
|                |         | Port Members: 1/ALL,2/ALL,3/ALL |          |        |        |         |      |
| Total VLANs: 1 |         |                                 |          |        |        |         |      |

### FA Server:

```
7024XLS(config)#vlan create 222 type port
7024XLS(config)#vlan mgmt 222
7024XLS(config)#show vlan
```

| Id             | Name      | Type                      | Protocol | PID    | Active | IVL/SVL | Mgmt |
|----------------|-----------|---------------------------|----------|--------|--------|---------|------|
| 1              | VLAN #1   | Port                      | None     | 0x0000 | Yes    | IVL     | No   |
|                |           | Port Members: 1/ALL,2/ALL |          |        |        |         |      |
| 222            | VLAN #222 | Port                      | None     | 0x0000 | Yes    | IVL     | Yes  |
|                |           | Port Members: NONE        |          |        |        |         |      |
| Total VLANs: 2 |           |                           |          |        |        |         |      |

```
7024XLS(config)#show fa spbm
```

```
Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Server
Fabric Attach Auto Attach Status: Disabled
Fabric Attach Message Authentication Status: Enabled
```

```
7024XLS(config)#show fa port 1/17
```

| Unit | Port | IfIndex | Service Advertisement |
|------|------|---------|-----------------------|
| 1    | 17   | 17      | Disabled              |

```
7024XLS(config)#show vlan interface info 1/17
```

| Unit/Port | Filter<br>Untagged<br>Frames | Filter<br>Unregistered<br>Frames | PVID | PRI | Tagging  | Name           |
|-----------|------------------------------|----------------------------------|------|-----|----------|----------------|
| 1/17      | No                           | Yes                              | 1    | 0   | UntagAll | Unit 1,Port 17 |

```
7024XLS(config)#fa auto-attach
7024XLS(config)#fa port 1/17
7024XLS(config)#show fa spbm
```

```
Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Server
Fabric Attach Auto Attach Status: Enabled
Fabric Attach Message Authentication Status: Enabled
```

```
7024XLS(config)#show fa port 1/17
```

| Unit | Port | IfIndex | Service Advertisement |
|------|------|---------|-----------------------|
| 1    | 17   | 17      | Enabled               |

### FA Proxy:

```
4850GTS-PWR+(config)#show fa spbm
```

```
Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Auto Attach Status: Enabled
Fabric Attach Message Authentication Status: Enabled
Fabric Attach Proxy Status: Enabled
Fabric Attach Primary Server Id: fc:a8:41:f3:d0:01
Fabric Attach Primary Server Descr:
Virtual Services Platform 7024XLS HW:RoB FW:10.1.0.6 SW:v10.3.1.031
```

```
4850GTS-PWR+(config)#show vlan mgmt
```

```
Management VLAN: 222
```

```
4850GTS-PWR+(config)#show vlan
```

| Id  | Name                                       | Type | Protocol | PID    | Active | IVL/SVL | Mgmt |
|-----|--------------------------------------------|------|----------|--------|--------|---------|------|
| 1   | VLAN #1<br>Port Members: 1/ALL,2/ALL,3/ALL | Port | None     | 0x0000 | Yes    | IVL     | No   |
| 222 | VLAN #222<br>Port Members: 2/91            | Port | None     | 0x0000 | Yes    | IVL     | Yes  |

Total VLANs: 2

### FA Server:

```
7024XLS(config)#show vlan
```

| Id  | Name                                 | Type | Protocol | PID    | Active | IVL/SVL | Mgmt |
|-----|--------------------------------------|------|----------|--------|--------|---------|------|
| 1   | VLAN #1<br>Port Members: 1/ALL,2/ALL | Port | None     | 0x0000 | Yes    | IVL     | No   |
| 222 | VLAN #222<br>Port Members: 1/17      | Port | None     | 0x0000 | Yes    | IVL     | Yes  |

Total VLANs: 2

```
7024XLS(config)#show vlan interface info 1/17
```

| Filter | Filter |
|--------|--------|
|--------|--------|

| Unit/Port | Untagged Frames | Unregistered Frames | PVID | PRI | Tagging | Name            |
|-----------|-----------------|---------------------|------|-----|---------|-----------------|
| 1/17      | No              | Yes                 | 1    | 0   | TagAll  | Unit 1, Port 17 |

## Configuring the FA authentication key

Use the following command to configure the FA authentication key:

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the FA authentication key:

```
[default] fa authentication-key
```

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

## Variable Definitions

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

| Variable  | Description                                        |
|-----------|----------------------------------------------------|
| [default] | Resets the FA authentication key to default value. |

## Enabling FA message authentication support

Use the following procedure to enable the FA message authentication support on an FA Proxy or FA Server.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable the FA message authentication support:

```
fa message-authentication
```

OR

```
default fa message-authentication
```

---

## Disabling FA message authentication support

Use the following procedure to disable the FA message authentication support on an FA Proxy or FA Server.

### Procedure

1. Enter Global Configuration mode:  
`enable`  
`configure terminal`
2. Disable the FA message authentication support:  
`no fa message-authentication`

# Chapter 5: 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 navigation tree, select **Configuration > VLAN > VLANs**.
2. Click the **Basic** tab.
3. Click **Insert**.
4. In the **Type** field, click **spbm-bvlan**.
5. Click **Insert** to create the primary B-VLAN.
6. Click **Insert**.
7. In the **Type** field, click **spbm-bvlan**.
8. Click **Insert** to create the secondary B-VLAN.
9. In the navigation tree, select **Configuration > IS-IS > SPBM**.
10. From the **Globals** tab, select **enable** to enable SPBM globally, and click **Apply**.
11. Click the **SPBM** tab.
12. Click **Insert** to create an SPBM instance (in this release, only one SPBM instance is supported).
13. In the **Id** field, specify the SPBM instance ID.
14. In the **NodeNickName** field, specify the node nickname (valid value is 2.5 bytes in the format <x.xx.xx>)

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

**\* Note:**

Although it is not strictly required for SPBM operation, Avaya recommends that you change the IS-IS system ID from the default B-MAC value to a recognizable address to easily identify a switch (using the **SystemID** field under the IS-IS Globals tab) . This helps to recognize source and destination addresses for troubleshooting purposes.

25. In the AdminState field, click **on**, and click **Apply**.
26. Under the IS-IS tab, click the **Interfaces** tab.
27. Click **Insert** to create an IS-IS circuit.
28. In the **IfIndex** field, specify the port or MLT on which to create the IS-IS interface.
29. 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 at the end of this procedure.

30. Select the newly created IS-IS circuit entry, and click **SPBM**.
31. In the **Interfaces SPBM** tab, click **Insert**.
32. In the **Spbmid** field, specify a SPBM identifier.
33. In the **State** field, select **enable**.
34. Click **Insert** to enable the SPBM instance on the IS-IS circuit.
35. Under the IS-IS tab, click the **Interfaces** tab.
36. In the **AdminState** field for the IS-IS circuit entry, select **on** to enable the IS-IS circuit.
37. Click **Apply**.

38. From the navigation tree, select **Configuration > VLAN > VLANs**.
39. Click the **Basic** tab.
40. Select the row for VLAN#1, and double-click the **PortMembers** cell.
41. Click the **port number** you specified for the IS-IS interface to remove it from the default VLAN, and click **Ok**.
42. In the toolbar, click **Apply**.

**\* Note:**

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

## SPBM field descriptions

**\* Note:**

The following tables list the minimum required SPBM and IS-IS parameters to allow SPBM to operate on the switch. For more detailed information on all of the parameters see the procedures that follow. For more information on how to configure VLANs, see *Configuring VLANs, Spanning Tree, and Multi-Link Trunking on Avaya Ethernet Routing Switch 4800 Series*, NN47205-501.

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> <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>ID</b>         | Specifies the system ID. Valid value is a 6-byte value in the format <xxxx.xxxx.xxxx> <p> <b>Note:</b><br/>Although it is not strictly required for SPBM operation, Avaya recommends that you change the IS-IS system ID from the default B-MAC value to a recognizable address to easily identify a switch (using the <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.

---

## Displaying the SPBM I-SID information

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

### Procedure

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

2. Click **SPBM**.
3. Click the **I-SID** tab.

---

## I-SID field descriptions

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

| Name                    | Description                                                                                                 |
|-------------------------|-------------------------------------------------------------------------------------------------------------|
| <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**.
3. Click the **L1 Area** tab.

---

## L1 Area field descriptions

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

| Name            | Description                                                                                                                 |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------|
| <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.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > IS-IS**.
2. Click **SPBM**.
3. Click the **Globals** tab.
4. To enable or disable SPBM, click **enable** or **disable** in the **GlobalEnable** field.
5. To configure the global ethertype value, click the desired option in the **GlobalEtherType** field.
6. Click **Apply**.

---

## Globals field descriptions

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

| Name                   | Description                                                                            |
|------------------------|----------------------------------------------------------------------------------------|
| <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. |

---

## Configuring SPBM parameters

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

## Procedure

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

---

## SPBM field descriptions

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

| Name                | Description                                                                                            |
|---------------------|--------------------------------------------------------------------------------------------------------|
| <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.     |

---

## 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 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.                                                   |
| <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 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 l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10.                                                                         |

## Configuring SPBM on an interface

Use the following procedure to configure SPBM on an interface.

### Procedure

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

## Interface SPBM 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 l1-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 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 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 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. |
| <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 field descriptions

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

| Name                 | Description                                                                                                                                                                                                                                       |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CircIndex</b>     |                                                                                                                                                                                                                                                   |
| <b>Index</b>         | A unique value identifying the IS adjacency from all other such adjacencies on this circuit. This value is automatically assigned by the system when the adjacency is created.                                                                    |
| <b>IfIndex</b>       | Specifies the IS-IS interface on which the adjacency is found.                                                                                                                                                                                    |
| <b>Usage</b>         | Specifies how the adjacency is used. On a point-to-point link, this can be level 1 and 2. But on a LAN, the usage is level 1 on the adjacency between peers at level 1, and level 2 for the adjacency between peers at level 2.                   |
| <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.                                                                                                                                    |
| <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.                                                                                              |

*Table continues...*

| Name              | Description                                                                                   |
|-------------------|-----------------------------------------------------------------------------------------------|
| <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 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>  | <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>ID</b>         | <p>Specifies the IS-IS system ID for the switch. Valid value is a 6-byte value in the format &lt;xxxx.xxxx.xxxx&gt;.</p> <p> <b>Important:</b><br/>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</p> |

*Table continues...*

| Name                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     | <p>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> <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 RxmtLspInt.</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>RxmtLspInt</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>                                                                                                                                                                                                                                               |

## Configuring system level IS-IS parameters

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

### Procedure

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

## System Level field descriptions

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

| Name                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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>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 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>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

*Table continues...*

| Name            | Description                                                                                                                                                       |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                 | <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 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> | <p>Specifies the router type globally:</p> <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> <p>The default value is I1.</p>                                              |
| <b>ISPriority</b> | <p>Specifies an integer sub-range for IS-IS priority. Range of 0–127. The default is 0 for SPBM interfaces.</p> <p> <b>Note:</b><br/>ISPriority only applies to broadcast interfaces.</p> |
| <b>HelloTimer</b> | Specifies the level 1 hello interval.                                                                                                                                                                                                                                        |

*Table continues...*

| Name                   | Description                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                        | <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> <p>The default value is 9 seconds.</p> |
| <b>HelloMultiplier</b> | <p>Specifies the level 1 hello multiplier. The default value is 3 seconds.</p>                                                                                                                                                                                                                                                                                                                       |
| <b>DRHelloTimer</b>    | <p>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.</p>                                                                                                                                                                                                                   |

---

## Configuring an IS-IS Manual Area

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

### Procedure

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

---

## Manual Area field descriptions

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

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



## 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 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 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>uitType</b>               |                                                                                                                                                                                 |
| <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>LANDesISChanges</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 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.                                                     |
| <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

Use the procedures in this section to configure Fabric Attach (FA) using Enterprise Device Manager.

---

### Configuring Fabric Attach

Use the following procedure to configure FA settings.

#### Procedure

1. From the navigation tree, select **Edit > Fabric Attach**.
2. Click the **SPBM** tab.
3. To enable or disable Auto Attach support, click **enable** or **disable** in the **AutoAttachService** field.
4. To enable or disable Message authentication support, click **enable** or **disable** in the **MsgAuthStatus** field.
5. Enter the desired password for message authentication in the **MsgAuthKey** field.

6. Confirm the password for message authentication in the **Confirm MsgAuthKey** field.
7. To enable or disable Fabric Attach external client proxy support, click **enable** or **disable** in the **HostProxyStatus** field.
8. Click **Apply**.

## Variable definitions

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

| Variable                  | Value                                                                        |
|---------------------------|------------------------------------------------------------------------------|
| <b>Service</b>            | Displays the service status.                                                 |
| <b>Element Type</b>       | Indicates whether the switch functions as an FA Proxy or FA Server.          |
| <b>PrimaryServerId</b>    | Displays the ID of the primary server.                                       |
| <b>PrimaryServerDescr</b> | Displays the FA Server description.                                          |
| <b>AutoAttachService</b>  | Specifies whether the Auto Attach support is enabled or disabled.            |
| <b>MsgAuthStatus</b>      | Specifies whether the message authentication support is enabled or disabled. |
| <b>MsgAuthKey</b>         | Specifies the password for message authentication.                           |
| <b>Confirm MsgAuthKey</b> | Requires entering again the password for message authentication.             |
| <b>HostProxyStatus</b>    | Specifies whether the external client proxy support is enabled or disabled.  |

---

## I-SID configuration

You can create, delete and view I-SID configuration.

### Displaying FA-specific settings

Use the following procedure to view FA-specific settings:

#### 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.

#### Variable Definitions

| Variable    | Value                                               |
|-------------|-----------------------------------------------------|
| <b>Isid</b> | Indicates the I-SID for this I-SID/VLAN assignment. |

*Table continues...*

| Variable | Value                                              |
|----------|----------------------------------------------------|
| Vlan     | Indicates the VLAN for this I-SID/VLAN assignment. |
| State    | Indicates the assignment state.                    |

## Creating an I-SID/VLAN assignment on an FA Proxy

Use the following procedure to create 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**.

### Variable definitions

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

| Name | Description                                    |
|------|------------------------------------------------|
| Isid | Specifies the I-SID to associate with a VLAN.  |
| Vlan | Specifies the VLAN to associate with an I-SID. |

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

Use the following procedure to delete 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. Select an I-SID/VLAN assignment.
5. Click **Delete**.
6. Click **Yes**.

---

## Configuring per-port FA settings

Use the following procedure to determine whether FA data is included in LLDPDUs.

**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 Fabric Attach information in LLDPDUs, select **enabled** or **disabled** in the **State** field for a specific port or ports.
5. Click **Apply**.

**Variable Definition**

| Variable        | Value                                                                |
|-----------------|----------------------------------------------------------------------|
| <b>IfIndex</b>  | Specifies the interface for which to enable or disable FA operation. |
| <b>State</b>    | Indicates whether FA operation is enabled or disabled.               |
| <b>enabled</b>  | Indicates that FA operation is enabled on corresponding interfaces.  |
| <b>disabled</b> | Indicates that FA operation is disabled on corresponding interfaces. |

# Chapter 6: Layer 2 VSN configuration fundamentals

This section provides fundamentals concepts for Layer 2 Virtual Services Networks (VSN).

---


## SPBM L2 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.



**Figure 3: SPBM L2 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:**

IP forwarding cannot be enabled on an ERS 4800 if SPB is enabled. The only way to manage the switch is via a Layer 2 VLAN.

If the ERS 4800 switch is connected to an SPB network that has IP Shortcuts or L3VSN enabled, you can create a management VLAN on the ERS 4800 with no port members, and assign it to an I-SID for L2 VSN terminated on an ERS 8800, VSP 9000 or VSP 4000 with the same I-SID and IP subnet.


To allow IP connectivity to the ERS 4800, on the ERS 8800, VSP 9000, or VSP 4000 where the L2 VSN is configured, add an IP address to the VLAN that terminates the L2VSN.

---

## SPBM L2 VSN sample operation

The following section shows how a SPBM network is established, in this case, a Layer 2 VSN.

1. *Discover network topology*



**Figure 4: 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 5: 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 6: Shortest path tree for source node A





Figure 7: Shortest path tree for source node B



**Figure 8: 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.

# Chapter 7: Layer 2 VSN configuration using ACLI

This section provides procedures to configure Layer 2 Virtual Services Networks (VSN) using Avaya Command Line Interface (ACLI).

---

## 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

```
4850GTS-PWR+> enable
4850GTS-PWR+# configure terminal
```



```
4850GTS-PWR+(config)# i-sid 200 vlan 200
```

```
4850GTS-PWR+(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 (CVLAN) to associate with the I-SID.</p> <p>Use the <code>no</code> or <code>default</code> options to remove the I-SID from the specified VLAN.</p> <p><b>* Note:</b></p> <p>Ethernet Routing Switch 4800 series reserves I-SID 0x00ffffff. ERS 4800 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> |

## 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

- 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

```
4850GTS-PWR+> enable
```

```
4850GTS-PWR+# configure terminal
```

```
4850GTS-PWR+(config)# vlan create 100 type spbm-switchedUni
```

```
4850GTS-PWR+(config)# i-sid 100 vlan 100 port 1
```

```
4850GTS-PWR+(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.

```
4850GTS-PWR+(config)# i-sid 101 vlan 100 port 2
```

```
4850GTS-PWR+(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                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| i-sid <1-16777215> vlan<2-4094> port <portlist> | <p>Specifies the Switched UNI VLAN to associate with the I-SID. and a port.</p> <p>Use the no or default options to remove the I-SID from the specified VLAN.</p> <p><b>* Note:</b></p> <p>Avaya Ethernet Routing Switch 4800 series reserves I-SID 0x00ffffff. ERS 4800 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

```
4850GTS-PWR+#show vlan i-sid
```

```
=====
 Vlan I-SID
=====
VLAN_ID I-SID

1
2
5 5
10
20
```

```
4850GTS-PWR+#show i-sid
```

```
I-SID Vid UNI-type Ports

```

```
5 5 CVLAN 4
100 100 switched 1
```

```
4850GTS-PWR+#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
200 1.12.45 1000 0016.ca23.73df discover VSP-9001
300 1.12.45 1000 0016.ca23.73df discover VSP-9001
=====
Total number of SPBM ISID entries configed: 6

Total number of SPBM ISID entries discovered: 2

Total number of SPBM ISID entries: 8
=====
```

## 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.                                                                                                                                    |

## Job aid

The following sections describe the fields in the outputs for the C-VLAN I-SID show commands.

### show vlan i-sid

The following table describes the fields in the output for the `show vlan i-sid` command.

| Parameter | Description                                                 |
|-----------|-------------------------------------------------------------|
| VLAN_ID   | Indicates the VLAN IDs.                                     |
| I-SID     | Indicates the I-SIDs associated with the specified C-VLANs. |

### show i-sid

The following table describes the fields in the output for the `show i-sid` command.

| Parameter | Description                                                    |
|-----------|----------------------------------------------------------------|
| I-SID     | Indicates the I-SID IDs.                                       |
| Vid       | Indicates the VLAN IDs.                                        |
| UNI-type  | Indicates the UNI-type as CVLAN or Switched                    |
| Ports     | Indicates ports associated with the specific I-SIDs and VLANs. |

### show isis spbm i-sid

The following describes the fields in the output for the `show isis spbm i-sid` command.

| Parameter                      | Description                                                                                                                                                                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ISID {all   discover   config} | Indicates the IS-IS SPBM I-SID identifier. <ul style="list-style-type: none"> <li>• all: display all SPBM I-SID</li> <li>• discover: display discovered SPBM I-SID</li> <li>• config: display configured SPBM I-SID</li> </ul> |
| SOURCE NAME                    | Indicates the nickname of the node where this I-SID was configured or discovered.<br>* <b>Note:</b><br>SOURCE NAME is equivalent to nickname.                                                                                  |
| VLAN                           | Indicates the B-VLAN where this I-SID was configured or discovered.                                                                                                                                                            |
| SYSID                          | Indicates the system identifier.                                                                                                                                                                                               |
| TYPE                           | Indicates the SPBM I-SID type as either configured or discovered.                                                                                                                                                              |
| HOST_NAME                      | Indicates the host name of the multicast FIB entry.                                                                                                                                                                            |

## 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 L2 VSN terminated on an ERS 8800, VSP 9000 or VSP 4000 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, VSP 9000, or VSP 4000 where the L2 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 ERS8800, VSP 9000, or VSP 4000, assign a VLAN to an I-SID with the same ID as the I-SID with which the management VLAN is associated on the ERS 4800, 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. |

# Chapter 8: 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 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 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 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 L2 VSN terminated on an ERS 8800, VSP 9000, VSP 4000 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, VSP 9000, or VSP 4000 where the L2 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 ERS8800, VSP 9000, or VSP 4000, assign a VLAN to an I-SID with the same ID as the I-SID with which the management VLAN is associated on the ERS 4800, and add an IP address to this VLAN.

---

## Variable definitions

| Variable           | Value                                                             |
|--------------------|-------------------------------------------------------------------|
| I-sid <1-16777214> | Specifies the I-SID with which the management VLAN is associated. |

# Chapter 9: CFM fundamentals

The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of ping and traceroute. To support troubleshooting of the SPBM cloud, Avaya Ethernet Routing switch 4800 Series supports a subset of CFM functionality.

CFM is based on the IEEE 802.1ag standard.

IEEE 802.1ag Connectivity Fault Management (CFM) provides OAM tools for the service layer, which allows you to monitor and troubleshoot an end-to-end Ethernet service instance. CFM is the standard for Layer 2 ping, Layer 2 traceroute, and the end-to-end connectivity check of the Ethernet network.

The 802.1ag feature divides or separates a network into administrative domains called Maintenance Domains (MD). Each MD is further subdivided into logical groupings called Maintenance Associations (MA). A single MD can contain several MAs.

Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Two types of MP exist:

- Maintenance End Point (MEP)
- Maintenance Intermediate Point (MIP)

CFM supports three kinds of standard CFM messages: Continuity Check Message (CCM), Loopback Message (LBM), and Link Trace Message (LTM). Messages are sent between Maintenance Points (MP) in the system.

On Avaya Ethernet Routing Switch 4800 Series, CFM is implemented using the LBM and LTM features only to debug SPBM. CCM messages are not required or supported in the current release.

---

## MD

A Maintenance Domain (MD) is the part of a network that is controlled by a single administrator. For example, a customer can engage the services of a service provider, who, in turn, can engage the services of several operators. In this scenario, there can be one MD associated with the customer, one MD associated with the service provider, and one MD associated with each of the operators.

You assign one of the following eight levels to the MD:

- 0–2 (operator levels)
- 3–4 (provider levels)
- 5–7 (customer levels)

The levels separate MDs from each other and provide different areas of functionality to different devices using the network. An MD is characterized by a level and an MD name (optional).

A single MD can contain several Maintenance Associations (MA).

**\* Note:**

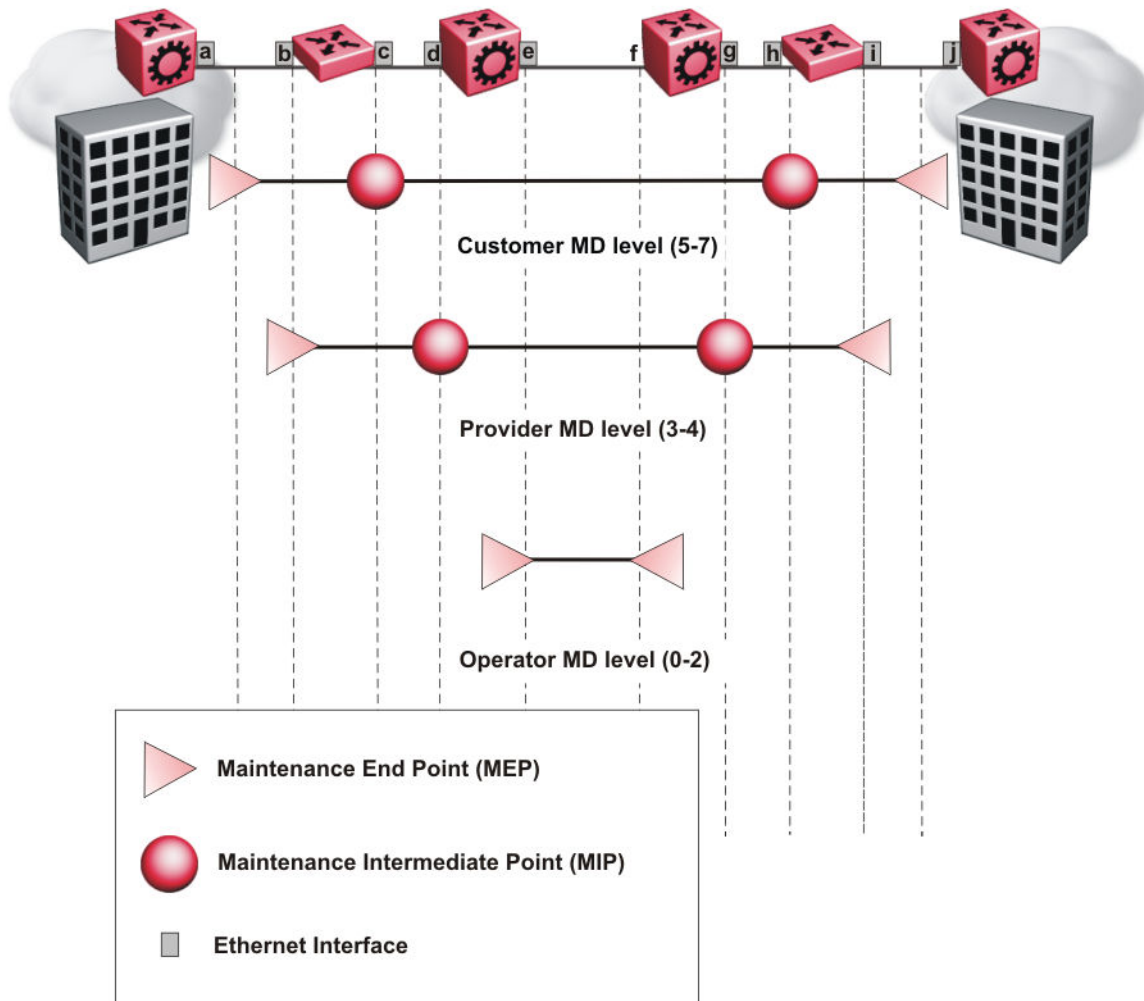
Avaya ERS 4800 Series supports one global MD, named spbm. The spbm MD has a default maintenance level of 4.

---

## MA

A Maintenance Association (MA) represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance association End Points (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

The following figure shows MD level assignment in accordance with the 802.1ag standard. As shown in the figure, MIPs can be associated with MEPs. However, MIPs can also function independently of MEPs.



## MEP

A Maintenance Endpoint (MEP) represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA. MEP functionality can be divided into the following functions:

- Fault Detection
- Fault Verification
- Fault Isolation
- Fault Notification

Fault detection and notification are achieved through the use of Continuity Check Messages (CCM). CCM messages are not supported in the current release.

---

## Fault verification

Fault verification is achieved through the use of Loopback Messages (LBM). An LBM is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a MEP or Maintenance Intermediate Point (MIP) but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR.

---

## LBM

The Loopback Message (LBM) packet is often compared to a ping. A MEP transmits the LBM packet. This packet can be addressed to another MEP or to the MAC address of the MP; in the case of SPBM, this is the SPBM system ID. Only the MP for which the packet is addressed responds with an LBR message. You can trigger an LBM with the `l2ping` command.

- Provides “ICMP ping like” functionality natively at Layer 2.
- DA is the MAC address of the target.
- Includes a transaction identifier that allows the corresponding LBR to be identified when more than one LBM request is waiting for a response.
- Only the target (MIP or MEP) responds.
- Initiator can choose the size and content of the data portion of the LBM frame.
- Can be used to check the ability of the network to forward different sized frames.

---

## Layer 2 ping

The `l2ping` command is a proprietary command that allows a user to trigger an LBM message.

For B-VLANs, specify either the destination MAC address or node name.

The `l2ping` command provides a ping equivalent at Layer 2 for use with nodes on the SPBM B-VLAN in the customer domain.

 **Note:**

Layer 2 ping supports B-VLANs only.

---

## Fault isolation

Fault isolation is achieved through the use of Linktrace Messages (LTM). LTM is intercepted by all the MPs on the way to the destination MP. ERS 4800 supports two types of LTM.

The first type, the unicast LTM, can be addressed to either MEP or MIP MAC address. Each MP on the way decrements the TTL field in the LTM frame, sends Linktrace Reply (LTR), and forwards the original LTM to the destination. The LTM forwards until it reaches the destination or the TTL value is decremented to zero. LTR is a unicast message addressed to the originating MEP.

The second type, the proprietary LTM, is used to map the MAC addresses of the SPBM network; in this case the target MAC is a service instance identifier (I-SID), not an MP.

---

## LTM

CFM offers Linktrace message (LTM) for fault isolation. LTM allow operators, service providers and customers to verify the connectivity that they provide or use and to debug systems.

### Link trace message — unicast

The LTM is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP, which is the SPBM system ID. MPs on the path to the target address respond with an Linktrace reply (LTR). You can trigger an LTM with the `l2traceroute` command.

- LTM trace the path to any given MAC address or System Name.
- DA is unicast
- LTM contains:
  - Time to live (TTL)
  - Transaction Identifier
  - Originator MAC address
  - Target MAC address
- CFM forward the frame like any other data frame.
- MIP or MEP that is not on the path to the target discards the LTM and does not reply.
- MIP that is on the path to the target
  - Forwards the LTM after decrementing the TTL and replacing the SA with its own address.
  - Sends an LTR to the originator.
  - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- If the MIP or MEP is a target
  - Sends an LTR to the originator.
  - Identifies itself in the forwarded LTM and LTR by modifying TLV information.

- A MEP that is not the target but is on the path to the target
  - Generates a reply as described above.
  - It also sets one of the flags fields in the reply to indicate that it is the terminal MEP.

### Link trace message — multicast

The multicast LTM can be used to trace the multicast tree from any node on any I- SID using the nickname MAC address and the I-SID multicast address.

Specifying a multicast target address for an LTM allows for the tracing of the multicast tree corresponding to that destination address (DA). With a multicast target every node that is in the active topology for that multicast address responds with a LTR and also forwards the LTM frame along the multicast path. Missing LTRs from the nodes in the path indicate the point of first failure.

This functionality allows you to better troubleshoot I-SID multicast paths in a SPBM network. You can use the command `l2tracetree` to trace the I-SID tree root.

## Layer 2 traceroute

The `l2traceroute` command is a proprietary command that allows a user to trigger an LTM message.

For B-VLANs, specify either the destination MAC address or node name.

The `l2 traceroute` command provides a trace equivalent at Layer 2 for use with nodes on the SPBM B-VLAN in the customer domain.

### Note:

Layer 2 traceroute supports B-VLANs only.

## Layer 2 tracetree

The `l2tracetree` command is a proprietary command that allows you to trigger a multicast LTM by specifying the B-VLAN and I-SID. Layer 2 tracetree allows you to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

## MIP

Maintenance domain intermediate points (MIPs) do not initialize any CFM messages. MIPs passively receive CFM messages, process the messages received and respond back to the originating MEP. By responding to received CFM messages, MIPs can support discovery of hop-by-

hop path among MEPs, allow connection failures to be isolated to smaller segments of the network to help discover location of faults along the paths. MIP functionality can be summarized as:

- Respond to Loopback (ping) messages at the same level as itself and addressed to it.
- Respond to Linktrace (traceroute) messages.
- Forward Linktrace messages after decrementing the TTL.

---

## Nodal MPs

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. The Nodal MEP provides traceability and troubleshooting at the system level for a given B-VLAN. Each switch has a given MAC address and communicates with other switches. The SPBM instance MAC address is used as the MAC address of the Nodal MP. The Nodal B-VLAN MPs supports eight levels of CFM.

---

## Configuration considerations

When you configure CFM, be aware of the following configuration considerations:

- The Maintenance level for MEPs and MIPs on a given B-VID (in a network) must be configured to the same level for them to respond to a given CFM command.
- CFM is supported only on B-VLANs.



# Chapter 10: CFM configuration using ACLI

This section provides procedures to configure and use Connectivity Fault Management (CFM) using Avaya Command Line Interface (ACLI). The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. To support troubleshooting of the SPBM cloud, Ethernet Routing Switch 4800 Series supports a subset of CFM functionality

**\* Note:**

When you enable CFM in an SBPM network, Avaya recommends that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

---

## Configuring CFM

Use this procedure to configure auto-generated CFM Maintenance End Points (MEPs) and Maintenance Intermediate Point (MIP) level for every SPBM B-VLAN on the ERS 4800. This procedure automatically configures a Maintenance Domain (MD) , Maintenance Associations (MAs), MEP ID, and also associates the MEPs and MIP level to the SPBM VLANs.

### About this task

When you enable CFM, you create a global MD (named spbm) for all the SPBM Nodal MEPs. The spbm MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs created use the MEP ID configured under the global context, which has a default value of 1. You can only modify the global context when CFM is disabled. The Nodal MEPs automatically associate with SPBM VLANs and associate to any SPBM VLAN added later. The MIP level maps to the global level. The MIP level automatically associates with the SPBM VLANs when CFM is enabled, and associate to any SPBM VLAN added later.

### Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the maintenance level for every CFM MEP and MIP level on all SPBM VLANs:

**\* Note:**

You can change the level before or after CFM is enabled. The default level is 4.

```
cfm spbm [level <0-7>]
```

3. Assign a global CFM MEP ID for all CFM SPBM MEPs:

**\* Note:**

You can change the MEP ID only when CFM is disabled.

```
cfm spbm mepid <1-8191>
```

4. Enable the CFM:

```
cfm spbm enable
```

5. Display the global CFM SPBM configuration:

```
show cfm spbm
```

6. If you want to default the CFM MD level, use the following command:

```
default cfm spbm level
```

7. If you want to default the MEP identifier, use the following command:

```
default cfm spbm mepid
```

8. If you want to disable CFM, use one of the following commands:

```
no cfm spbm enable
```

```
default cfm spbm enable
```

### Example

```
4850GTS-PWR+> enable
4850GTS-PWR+# configure terminal
4850GTS-PWR+(config)# cfm spbm level 4
4850GTS-PWR+(config)# cfm spbm mepid 200
4850GTS-PWR+(config)# cfm spbm enable
4850GTS-PWR+(config)# show cfm spbm
```

```
CFM Admin State: Enabled
CFM Spbm Level: 4
CFM Mep Id: 200
```

---

## Variable definitions

Use the data in the following table to use the `cfm spbm` commands.

| Variable                | Value                                                                                                                 |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------|
| cfm spbm level <0-7>    | Specifies the CFM MD level. The default is 4.                                                                         |
| cfm spbm mepid <1-8191> | Specifies the MEP ID. The default is 1.<br><br>* <b>Note:</b><br>You can only modify the MEP ID when CFM is disabled. |
| cfm spbm enable         | Enables CFM globally.                                                                                                 |
| no cfm spbm enable      | Disables CFM globally.                                                                                                |
| default cfm spbm level  | Defaults the CFM MD level.                                                                                            |
| default cfm spbm mepid  | Defaults the CFM MEP ID.                                                                                              |
| default cfm spbm enable | Defaults CFM. Default is globally disabled.                                                                           |
| show cfm spbm           | Displays the current CFM configuration.                                                                               |

## Triggering an LBM Layer 2 ping

Use this procedure to trigger a Layer 2 ping, which acts like native ping. This feature enables CFM to debug Layer 2.

### Before you begin

CFM SPBM must be enabled.

### About this task

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Trigger a Layer 2 ping:

```
l2ping {vlan <1-4094> routernodename WORD<0-255> | vlan <1-4094> mac
<0x00:0x00:0x00:0x00:0x00:0x00>} [burst-count <1-200>] [data-tlv-
size <0-400>] [frame-size <64-1500>] [priority <0-7>] [testfill-
pattern <all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-
random-bit-sequence-crc>] [time-out <1-10>]
```

### Example

```
4850GTS-PWR+# l2ping vlan 500 mac 00.14.0d.bf.a3.df
```

```
Please wait for l2ping to complete or press any key to abort
----00:14:0d:bf:a3:df L2 PING Statistics---- 0(68) bytes of data
1 packets transmitted, 0 packets received, 100.00% packet loss
```

```
4850GTS-PWR+# l2ping vlan 500 routernodename ERS-MONTIO
```

```
Please wait for l2ping to complete or press any key to abort
```

```
----00:14:0d:a2:b3:df L2 PING Statistics---- 0(68) bytes of data
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip (us) min/max/ave/stdv = 26895/26895/26895.00/ 0.00
```

## Variable definitions

Use the data in the following table to configure the `l2ping` parameters.

| Variable                                                                                                   | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vlan <1–4094> routernodename<br>WORD<0–255><br><br>vlan <1–4094> mac<br><0x00:0x00:0x00:0x00:0x00:0x00>    | Specifies the destination for the L2 ping: <ul style="list-style-type: none"> <li>• &lt;1–4094&gt; — Specifies the VLAN ID.</li> <li>• WORD&lt;0–255&gt; — Specifies the Router node name.</li> <li>• &lt;XX:XX:XX:XX:XX:XX&gt; — Specifies the MAC address.</li> </ul>                                                                                                                                                                                                                                                                                  |
| burst-count <1–200>                                                                                        | Specifies the burst count.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| data-tlv-size <0–400>                                                                                      | Specifies the data TLV size. The default is 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| frame-size <64–1500>                                                                                       | Specifies the frame size. The default is 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| testfill-pattern <all-zero all-zero-crc <br>pseudo-random-bit-sequence pseudo-<br>random-bit-sequence-crc> | Specifies the testfill pattern: <ul style="list-style-type: none"> <li>• all-zero — null signal without cyclic redundancy check</li> <li>• all-zero-crc — null signal with cyclic redundancy check with 32-bit polynomial</li> <li>• pseudo-random-bit-sequence — pseudo-random-bit-sequence without cyclic redundancy check</li> <li>• pseudo-random-bit-sequence-crc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial.</li> </ul> A cyclic redundancy check is a code that detects errors.<br>The default is all-zero. |
| priority <0–7>                                                                                             | Specifies the priority. The default is 7.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| time-out <1–10>                                                                                            | Specifies the interval in seconds. The default is 3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Triggering an LTM Layer 2 traceroute

Use this procedure to trigger a Layer 2 traceroute, which acts like native traceroute. This feature enables CFM to debug Layer 2.

**! Important:**

The MAC address must be learned before you can trace a route to a MAC address. For B-VLANs, IS-IS learns the MAC addresses and populates the FDB table.

`linktrace` traces the path up to the closest device to that MAC address that supports CFM.

**Before you begin**

CFM SPBM must be enabled.

**About this task**

The link trace message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID. MPs on the path to the target address respond with an LTR.

**Procedure**

1. Enter Privileged EXEC mode:

```
enable
```

2. Trigger a Layer 2 traceroute:

```
l2traceroute {<vlan <1-4094> routernodename WORD<0-255> | <vlan <1-4094> mac <0x00:0x00:0x00:0x00:0x00:0x00>} [priority <0-7>] [ttl <1-255>]
```

**Example**

```
4850GTS-PWR+# l2traceroute vlan 500 routernodename ERS-MONTIO
```

```
Please wait for l2traceroute to complete or press any key to abort
```

```
l2traceroute to VSP-MONTIO (00:14:0d:a2:b3:df), vlan 500
0 ERS-PETER4 (00:15:9b:11:33:df)
1 ERS-MONTIO (00:14:0d:a2:b3:df)
```

**Variable definitions**

Use the data in the following table to use the `l2traceroute` command.

| Variable                                              | Value                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vlan <1-4094> routernodename<br>WORD<0-255>}          | Specifies the destination for the L2 traceroute: <ul style="list-style-type: none"> <li>• &lt;1-4094&gt; — Specifies the VLAN ID</li> <li>• WORD&lt;0-255&gt; — Specifies the Router Node Name</li> <li>• &lt;XX:XX:XX:XX:XX:XX&gt; — Specifies the MAC address</li> </ul> |
| vlan <1-4094> mac<br><0x00:0x00:0x00:0x00:0x00:0x00>} | • WORD<0-255> — Specifies the Router Node Name<br>• <XX:XX:XX:XX:XX:XX> — Specifies the MAC address                                                                                                                                                                        |
| ttl<1-255>                                            | Specifies the TTL value. The default is 64.                                                                                                                                                                                                                                |
| priority <0-7>                                        | Specifies the priority. The default is 7.                                                                                                                                                                                                                                  |

## Triggering an LTM Layer 2 tracetable

Use this procedure to trigger a Layer 2 tracetable. Layer 2 tracetable allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

### Before you begin

CFM SPBM must be enabled.

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Trigger a Layer 2 tracetable:

```
l2tracetable vlan <1-4094> isid <1-16777215> [routernodename WORD<0-255> | mac <0x00:0x00:0x00:0x00:0x00:0x00>] [priority <0-7>] [ttl <1-255>]
```

### Example

```
4850GTS-PWR+# l2tracetable vlan 2 isid 1 mac 53:55:10:00:00:01
```

```
Please wait for l2tracetable to complete or press any key to abort
```

```
l2tracetable to 53:55:10:00:00:01, vlan 2 i-sid 1 nickname 5.55.10
hops 64
1 ERS-PETER4 00:15:9b:11:33:df -> ERS-MONTIO 00:14:0d:a2:b3:df
2 ERS-MONTIO 00:14:0d:a2:b3:df -> ERS-LEE2 00:15:e8:b8:a3:df
```

## Variable definitions

Use the data in the following table to use the `l2tracetable` command.

| Variable                            | Value                                                                                                                                            |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| vlan<1-4094>isid <1-16777215>       | <ul style="list-style-type: none"> <li>• &lt;1-4094&gt; — Specifies the VLAN ID.</li> <li>• &lt;1-16777215&gt; — Specifies the I-SID.</li> </ul> |
| routernodename WORD<0-255>          | WORD<0-255> — Specifies the Router Node Name.                                                                                                    |
| mac <0x00:0x00:0x00:0x00:0x00:0x00> | <0x00:0x00:0x00:0x00:0x00:0x00> — Specifies the MAC address.                                                                                     |
| ttl <1-255>                         | Specifies the TTL value. The default is 64.                                                                                                      |
| priority <0-7>                      | Specifies the priority value. The default is 7.                                                                                                  |

# Chapter 11: CFM configuration using EDM

This section provides procedures to configure Connectivity Fault management (CFM) using Enterprise Device Manager (EDM).

**\* Note:**

When you enable CFM in an SBPM network, Avaya recommends that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

---

## Configuring CFM

Use this procedure to configure auto-generated CFM Maintenance End Points (MEPs) and Maintenance Intermediate Point (MIP) level for every SPBM B-VLAN on the ERS 4800. This procedure automatically configures a Maintenance Domain (MD), Maintenance Associations (MAs), MEP ID, and also associates the MEPs and MIP level to the SPBM VLANs.

### About this task

When you enable CFM, you create a global MD (named `spbm`) for all the SPBM Nodal MEPs. The `spbm` MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs created use the MEP ID configured under the global context, which has a default value of 1. You can only modify the global context when CFM is disabled. The Nodal MEPs automatically associate with SPBM VLANs and associate to any SPBM VLAN added later. The MIP level maps to the global level. The MIP level automatically associates with the SPBM VLANs when CFM is enabled, and associate to any SPBM VLAN added later.

### Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **Globals** tab.
4. In the **SpbmAdminState** field, click a radio button to enable or disable CFM. specify an index value, name, and level for the MD.
5. In the **SpbmLevel** field, configure the maintenance level for every CFM MEP and MIP level on all the SPBM VLANs.
6. In the **SpbmMepld** field, assign a global CFM MEP ID for all CFM SPBM MEPs.

- On the toolbar, click **Apply**.

---

## Globals field descriptions

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

| Name                  | Description                                                                    |
|-----------------------|--------------------------------------------------------------------------------|
| <b>EtherType</b>      | Read only Ethernet type value. Value of 0x8902                                 |
| <b>SpbmAdminState</b> | Enables or disables the SPBM CFM MD. Click the enable or disable radio button. |
| <b>SpbmLevel</b>      | Specifies the MD level. Default is level 4.                                    |
| <b>SpbmMepId</b>      | Specifies the MEP identifier. Default is 1                                     |

---

## Displaying CFM MD

Use this procedure to display the Connectivity Fault Management (CFM) Maintenance Domain (MD). An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

### Procedure

- From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
- Click **CFM**.
- Click the **MD** tab.
- On the toolbar, click **Refresh** to display the current MD configuration.

---

## MD field descriptions

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

| Name           | Description                                                         |
|----------------|---------------------------------------------------------------------|
| <b>Index</b>   | Specifies a maintenance domain entry index.                         |
| <b>Name</b>    | Specifies the MD name.                                              |
| <b>NumOfMa</b> | Indicates the number of MAs that belong to this maintenance domain. |
| <b>Level</b>   | Specifies the MD maintenance level. The default is 4.               |

*Table continues...*



| Name            | Description                                                         |
|-----------------|---------------------------------------------------------------------|
| <b>NumOfMip</b> | Indicates the number of MIPs that belong to this maintenance domain |
| <b>Type</b>     | Indicates the type of domain.                                       |

---

## Displaying CFM MA

Use this procedure to display a CFM Maintenance Association (MA). An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance Endpoints (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

### Before you begin

You must configure a CFM MD.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **MD** tab.
4. Select an existing MD.
5. On the toolbar, click **MaintenanceAssociation**.

---

## MA field descriptions

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

| Name                    | Description                                                               |
|-------------------------|---------------------------------------------------------------------------|
| <b>DomainIndex</b>      | Specifies the maintenance domain entry index.                             |
| <b>AssociationIndex</b> | Specifies a maintenance association entry index.                          |
| <b>DomainName</b>       | Specifies the MD name.                                                    |
| <b>AssociationName</b>  | Specifies the MA name.                                                    |
| <b>NumOfMep</b>         | Indicates the number of MEPs that belong to this maintenance association. |

## Displaying CFM MEP

Use this procedure to display the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **MD** tab.
4. Select an existing MD, and then click **MaintenanceAssociation**.
5. In the **MA** tab, select an existing MA, and then click **MaintenanceEndpoint**.

## MEP field descriptions

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

| Name                    | Description                                                                                                                                                                                                                                                |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DomainIndex</b>      | Specifies the MD index.                                                                                                                                                                                                                                    |
| <b>AssociationIndex</b> | Specifies the MA index.                                                                                                                                                                                                                                    |
| <b>Id</b>               | Specifies the MEP ID.                                                                                                                                                                                                                                      |
| <b>DomainName</b>       | Specifies the MD name.                                                                                                                                                                                                                                     |
| <b>AssociationName</b>  | Specifies the MA name.                                                                                                                                                                                                                                     |
| <b>AdminState</b>       | Specifies the administrative state of the MEP. The default is disable.                                                                                                                                                                                     |
| <b>MepType</b>          | Specifies the MEP type: <ul style="list-style-type: none"> <li>• trunk</li> <li>• sg</li> <li>• endpt</li> <li>• vlan</li> <li>• port</li> <li>• endptClient</li> <li>• nodal</li> <li>• remotetrunk</li> <li>• remotesg</li> <li>• remoteendpt</li> </ul> |

*Table continues...*

| Name                      | Description                                                                                                       |
|---------------------------|-------------------------------------------------------------------------------------------------------------------|
|                           | <ul style="list-style-type: none"> <li>• remoteVlan</li> <li>• remotePort</li> <li>• remoteEndptClient</li> </ul> |
| <b>ServiceDescription</b> | Specifies the service to which this MEP is assigned.                                                              |

## Configuring Layer 2 ping

Use this procedure to configure a Layer 2 ping. This feature enables CFM to debug Layer 2. It can also help you debug ARP problems by providing the ability to troubleshoot next hop ARP records.

### Before you begin

CFM SPBM must be enabled.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. From the **L2Ping** tab, configure the Layer 2 ping properties.
4. To initiate a Layer 2 ping, highlight an entry and click the **Start** button.
5. To update a Layer 2 ping, click the **Refresh** button.
6. To stop the Layer 2 ping, click the **Stop** button.

## L2Ping field descriptions

Use the data in the following table to use the **L2Ping** tab.

| Name                  | Description                                                                                                                                 |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>VlanId</b>         | Identifies the backbone VLAN.                                                                                                               |
| <b>DestMacAddress</b> | Specifies the target MAC address.                                                                                                           |
| <b>HostName</b>       | Specifies the target host name.                                                                                                             |
| <b>DestIsHostName</b> | Indicates whether the host name is (true) or is not (false) used for L2Ping transmission.                                                   |
| <b>Messages</b>       | Specifies the number of L2Ping messages to be transmitted. The default is 1.                                                                |
| <b>Status</b>         | Specifies the status of the transmit loopback service: <ul style="list-style-type: none"> <li>• ready: the service is available.</li> </ul> |

*Table continues...*

| Name               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <ul style="list-style-type: none"> <li>• transmit: the service is transmitting, or about to transmit, the L2Ping messages.</li> <li>• abort: the service aborted or is about to abort the L2Ping messages.</li> </ul> <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>                                                                                                                                                    |
| <b>ResultOk</b>    | <p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> <li>• true: the L2Ping Messages will be (or have been) sent.</li> <li>• false: the L2Ping Messages will not be sent.</li> </ul> <p>The default is true.</p>                                                                                                                                                                                                                                                                                                                                    |
| <b>Priority</b>    | <p>Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame.</p> <p>The default is 7.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>TimeoutInt</b>  | <p>Specifies the interval to wait for an L2Ping time-out. The default value is 3 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>TestPattern</b> | <p>Specifies the test pattern to use in the L2Ping PDU:</p> <ul style="list-style-type: none"> <li>• allZero: null signal without cyclic redundancy check</li> <li>• allZeroCrc: null signal with cyclic redundancy check with 32-bit polynomial</li> <li>• pseudoRandomBitSequence: pseudo-random-bit-sequence without cyclic redundancy check</li> <li>• pseudoRandomBitSequenceCrc: pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial.</li> </ul> <p>A cyclic redundancy check is a code that detects errors. The default value is allZero.</p> |
| <b>DataSize</b>    | <p>Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The default is 0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>FrameSize</b>   | <p>Specifies the frame size. If the frame size is specified then the data size is internally calculated and the calculated data size is included in the data TLV. The default is 0.</p>                                                                                                                                                                                                                                                                                                                                                                                         |

*Table continues...*

| Name              | Description                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SourceMode</b> | Specifies the source modes of the transmit loopback service: <ul style="list-style-type: none"> <li>• nodal</li> <li>• smltVirtual</li> </ul> The default is nodal. |
| <b>SeqNumber</b>  | The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.                                                            |
| <b>Result</b>     | Displays the Layer 2 Ping result.                                                                                                                                   |

---

## Initiating a Layer 2 traceroute

Use this procedure to trigger a Layer 2 traceroute. This feature enables CFM to debug Layer 2.

If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.

### Important:

The MAC address must be learned before you can trace a route to a MAC address.

For B-VLANs, IS-IS learns the MAC address and populates the FDB table.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

### Before you begin

CFM SPBM must be enabled.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 Traceroute/TraceTree** tab.
4. To configure the traceroute or tracetree, highlight an entry and populate the required column fields.
5. To start the traceroute, click the **Start** button.
6. To update the traceroute, click the **Refresh** button.
7. To stop the traceroute, click the **Stop** button.

## L2Traceroute field descriptions

Use the data in the following table to use the **L2Traceroute** tab.

| Name                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>VlanId</b>         | Specifies a value that uniquely identifies the Backbone VLAN (B-VLAN).                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Priority</b>       | Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>DestMacAddress</b> | Specifies the target MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>HostName</b>       | Specifies the target host name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>DestIsHostName</b> | Specifies whether the host name is (true) or is not (false) used for the L2Trace transmission.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Isid</b>           | Specifies the Service Instance Identifier (I-SID).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>NickName</b>       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>IsTraceTree</b>    | Specifies whether the multicast tree or unicast path is traced. If you configure <b>IsTraceTree</b> to false then EDM performs Traceroute on the unicast path. If you configure <b>IsTraceTree</b> to true then EDM performs TraceTree on the multicast tree.                                                                                                                                                                                                                                                                            |
| <b>Status</b>         | <p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> <li>• ready: the service is available.</li> <li>• transmit: the service is transmitting, or about to transmit, the L2Trace messages.</li> <li>• abort: the service aborted or is about to abort the L2Trace messages.</li> </ul> <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p> |
| <b>ResultOk</b>       | <p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> <li>• true: the L2Trace messages will be (or have been) sent.</li> <li>• false: the L2Trace messages will not be sent.</li> </ul> <p>The default is true.</p>                                                                                                                                                                                                                                                                                           |
| <b>Ttl</b>            | Specifies the number of hops remaining to this L2Trace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

*Table continues...*

| Name              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   | <p>This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The value of the time-to-live (TTL) field in the L2Trace is defined by the originating MEP.</p> <p>The default value is 64.</p>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>SourceMode</b> | Specifies the source mode of the transmit loopback service. The default is nodal.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>SeqNumber</b>  | Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Flag</b>       | <p>L2Trace result flag indicating L2Trace status or error code:</p> <ul style="list-style-type: none"> <li>• none (1): No error</li> <li>• internalError (2): L2Trace internal error</li> <li>• invalidMac (3): Invalid MAC address</li> <li>• mepDisabled (4): MEP must be enabled in order to perform L2Trace</li> <li>• noL2TraceResponse (5): No L2Trace response received</li> <li>• l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent</li> <li>• l2TraceComplete (7): L2Trace completed</li> <li>• l2TraceLookupFailure (8): Lookup failure for L2Trace</li> <li>• l2TraceLeafNode (9): On a leaf node in the I-SID tree</li> <li>• l2TraceNotInTree (10): Not in the I-SID tree</li> </ul> |

---

## Viewing Layer 2 traceroute results

Use this procedure to view Layer 2 traceroute results. This feature enables CFM to debug Layer 2. You can use Layer 2 traceroute to debug ARP problems by troubleshooting next hop ARP records.

### About this task

You can display Layer 2 tracetree results to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

## Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2Traceroute/TraceTree** tab.
4. Click the **Refresh** button to update the results.
5. To view the traceroute results, highlight an entry, and then click **Result**.

---

## L2 Traceroute Result field descriptions

Use the data in the following table to use the **L2 Traceroute Result** tab.

| Name                | Description                                                                                                                                                                                                              |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>VlanId</b>       | A value that uniquely identifies the Backbone VLAN (B-VLAN).                                                                                                                                                             |
| <b>SeqNumber</b>    | The transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which L2Trace's response of the L2Trace is going to be returned. The default is 0.                      |
| <b>Hop</b>          | The number of hops away from L2Trace initiator.                                                                                                                                                                          |
| <b>ReceiveOrder</b> | An index to distinguish among multiple L2Trace responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses. |
| <b>Ttl</b>          | Time-to-Live (TTL) field value for a returned L2Trace response.                                                                                                                                                          |
| <b>SrcMac</b>       | MAC address of the MP that responds to the L2Trace request for this L2TraceReply.                                                                                                                                        |
| <b>HostName</b>     | The host name of the replying node.                                                                                                                                                                                      |
| <b>LastSrcMac</b>   | The MAC address of the node that forwarded the L2Trace to the responding node.                                                                                                                                           |
| <b>LastHostName</b> | The host name of the node that forwarded the L2Trace to the responding node.                                                                                                                                             |



# Chapter 12: Resources

---

## Support

Go to the Avaya Support website at <http://support.avaya.com> for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

---

## Documentation

For a list of the documentation for this product and more information about documents on how to configure other switch features, see *Documentation Reference for Avaya Ethernet Routing Switch 4800 Series*, NN47205–101.

For more information on new features of the switch and important information about the latest release, see *Release Notes for Avaya Ethernet Routing Switch 4800 Series*, NN47205-400.

For more information about how to configure security, see *Configuring Security on Avaya Ethernet Routing Switch 4800 Series*, NN47205-505.

For the current documentation, see the Avaya Support web site: [www.avaya.com/support](http://www.avaya.com/support).

---

## Training

Ongoing product training is available. For more information or to register, see <http://avaya-learning.com/>.

Enter the course code in the **Search** field and click **Go** to search for the course.

| Course code | Course title                                           |
|-------------|--------------------------------------------------------|
| 8D00020E    | Stackable ERS and VSP Products Virtual Campus Offering |

---

## Viewing Avaya Mentor videos

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

### About this task

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

### Procedure

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

 **Note:**

Videos are not available for all products.

---

## Searching a documentation collection

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

### Before you begin

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

### Procedure

1. Extract the document collection zip file into a folder.
2. Navigate to the folder that contains the extracted files and open the file named `<product_name_release>.pdx`.

3. In the Search dialog box, select the option **In the index named <product\_name\_release>.pdx**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
  - Whole Words Only
  - Case-Sensitive
  - Include Bookmarks
  - Include Comments
6. Click **Search**.

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

---

## Subscribing to e-notifications

Subscribe to e-notifications to receive an email notification when documents are added to or changed on the Avaya Support website.

### About this task

You can subscribe to different types of general notifications, for example, Product Correction Notices (PCN), which apply to any product or a specific product. You can also subscribe to specific types of documentation for a specific product, for example, Application & Technical Notes for Virtual Services Platform 7000.

### Procedure

1. In an Internet browser, go to <https://support.avaya.com>.
2. Type your username and password, and then click **Login**.
3. Under **My Information**, select **SSO login Profile**.
4. Click **E-NOTIFICATIONS**.
5. In the GENERAL NOTIFICATIONS area, select the required documentation types, and then click **UPDATE**.

GENERAL NOTIFICATIONS

1/5 Notifications Selected

|                                                 |                                     |
|-------------------------------------------------|-------------------------------------|
| End of Sale and/or Manufacturer Support Notices | <input type="checkbox"/>            |
| Product Correction Notices (PCN)                | <input checked="" type="checkbox"/> |
| Product Support Notices                         | <input type="checkbox"/>            |
| Security Advisories                             | <input type="checkbox"/>            |
| Services Support Notices                        | <input type="checkbox"/>            |

UPDATE >>

6. Click **OK**.
7. In the PRODUCT NOTIFICATIONS area, click **Add More Products**.

PRODUCT NOTIFICATIONS

Show Details

Add More Products

1 Notices

8. Scroll through the list, and then select the product name.
9. Select a release version.
10. Select the check box next to the required documentation types.

| PRODUCTS                                    | My Notifications                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Virtual Services Platform 7000              | <b>VIRTUAL SERVICES PLATFORM 7000</b><br>Select a Release Version<br>All and Future<br>Administration and System Programming <input type="checkbox"/><br>Application Developer Information <input type="checkbox"/><br>Application Notes <input type="checkbox"/><br>Application and Technical Notes <input checked="" type="checkbox"/><br>Declarations of Conformity <input type="checkbox"/><br>Documentation Library <input checked="" type="checkbox"/><br><b>SUBMIT &gt;&gt;</b> |
| Virtualization Provisioning Service         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Visual Messenger™ for OCTEL® 250/350        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Visual Vectors                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Visualization Performance and Fault Manager |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Voice Portal                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Voice over IP Monitoring                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| W310 Wireless LAN Gateway                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| WLAN 2200 Series                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| WLAN Handset 2200 Series                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

11. Click **Submit**.

# Glossary

|                                             |                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ACLI</b>                                 | Avaya Command Line Interface (ACLI) is a text-based, common command line interface used for device configuration and management across Avaya products.                                                                                                                                                                                                                               |
| <b>ACLI 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>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>Auto Attach (AA)</b>                     | A Fabric Attach feature that allows a device to extract management VLAN data from the primary FA server advertisements and use this data to update the in-use management VLAN and initiate IP address acquisition using DHCP.                                                                                                                                                        |
| <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>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.

**designated router (DR)**

A single router elected as the designated router for the network. In a broadcast or nonbroadcast multiple access (NBMA) network running the Open Shortest Path First (OSPF) protocol, a DR ensures all network routers synchronize with each other and advertises the network to the rest of the Autonomous System (AS). In a multicast network running Protocol Independent Multicast (PIM), the DR acts as a representative router for directly connected hosts. The DR sends control messages to the rendezvous point (RP) router, sends register messages to the RP on behalf of directly connected sources, and maintains RP router status information for the group.

**Enterprise Device Manager (EDM)**

A web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.

**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.

|                                                              |                                                                                                                                                                                                                                                                            |
|--------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Local Area Network (LAN)</b>                              | 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).                                                                           |
| <b>media</b>                                                 | 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.

**Spanning Tree Protocol (STP)**

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.

**Split MultiLink Trunking (SMLT)**

An extension to IEEE 802.1AX (link aggregation), provides nodal and link failure protection and flexible bandwidth scaling to improve on the level of Layer 2 resiliency.

**stack**

Stackable Avaya Ethernet Routing Switches 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.

**time-to-live (TTL)**

The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.

**trunk**

A logical group of ports that behaves like a single large port.

**Virtual Local Area Network (VLAN)**

A Virtual Local Area Network is a group of hosts that communicate as if they are attached to the same broadcast domain regardless of their physical location. VLANs are layer 2 constructs.

**Virtual Private Network (VPN)**

A Virtual Private Network (VPN) requires remote users to be authenticated and ensures private information is not accessible to unauthorized parties. A VPN can allow users to access network resources or to share data.