Configuration — Shortest Path Bridging MAC (SPBM)
Avaya Ethernet Routing Switch 8800/8600

7.1
NN46205-525, 01.03
February 2011

Notice

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

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya generally makes available to users of its 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 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 Web sites 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 its Hardware and Software (“Product(s)”). 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 Web site: http://support.avaya.com. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/ ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. 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 AUTHORIZED RESELLER; 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”).

Copyright

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

Third-party components

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements (“Third Party Components”), which may contain terms that expand or limit rights to use certain portions of the Product (“Third Party Terms”). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: http://support.avaya.com/Copyright.

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 Web site: http://support.avaya.com. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation 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, and “Linux” is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support Web site: http://support.avaya.com.

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your Product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: http://support.avaya.com.

Contents

Chapter 1: New in this release

Avaya Ethernet Routing Switch 8800/8600 Configuration – Shortest Path Bridging (NN46205-510) is a new document for Release 7.1 of the Ethernet Routing Switch 8800/8600.

Chapter 2: SPBM fundamentals

Release 7.1 of the Ethernet Routing Switch 8800/8600 supports the IEEE 802.1aq standard of Shortest Path Bridging MAC (SPBM). SPBM makes network virtualization much easier to deploy within the Enterprise environment, reducing the complexity of the network while at the same time providing greater scalability.

Note:
The Ethernet Routing Switch 8800/8600 implementation of SPBM is a pre-standard implementation because the IEEE 802.1aq standard is not yet ratified.

SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet based link state protocol which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.

SPBM layers the Ethernet network into customer/edge and provider/backbone domains with complete isolation among their MAC addresses. This base technology provides the foundation for the integration in a single control plane of all the functions that MPLS requires multiple layers and protocols to support. At the same time SPBM provides a clear separation of infrastructure and service layer. Once you create an SPBM infrastructure, you can add additional services (such as VLAN extensions or VRF extensions) by end-point-point provisioning only. SPBM introduces a service identifier called I-SID.

SPBM supports the following implementation options:

- L2 Virtual Services Network (VSN): bridges customer VLANs (C-VLANs) over the SPBM core infrastructure
- L3 Virtual Services Network (VSN): provides IP connectivity over SPBM for VRFs.
- Native IP Shortcuts: forwards standard IP packets over IS-IS in the SPBM core
- IP VPN Lite over SPBM: switches IP-in-IP packets in the SPBM core
- Inter-ISID routing: provides routing between Layer 2 VLANs with different I-SIDs

The following sections provide more details about the fundamentals of SPBM as well as the various SPBM implementation options.

MAC-in-MAC encapsulation

To forward customer traffic across the service provider backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC

encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses.

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

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

I-SID

The SPBM B-MAC header includes a service instance identifier (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. These I-SIDs are used to virtualize VLANs (L2 Virtual Services Network [VSN]) or VRFs (L3 Virtual Services Network [VSN]) across the MAC-in-MAC backbone.

With L2 VSN, the I-SID is associated with a customer VLAN, which is then virtualized across the backbone. With L3 VSN, the I-SID is associated with a customer VRF, which is also virtualized across the backbone.

Note:

I-SID configuration is required only for virtual services such as L2 VSN and L3 VSN. With IP Native shortcuts, no I-SID is required as forwarding is done using the (Global Routing Table) GRT.

BCBs and BEBs

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

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

Basic SBPM network topology

The following figure shows a basic SPBM network topology, specifically an L2 VSN. Switches A and D are the Backbone Edge Bridges (BEB) that provide the boundary between the customer VLANs (C-VLAN) and the Backbone. Switches B and C are the Backbone Core Bridges (BCB) that form the core of the SPBM network.

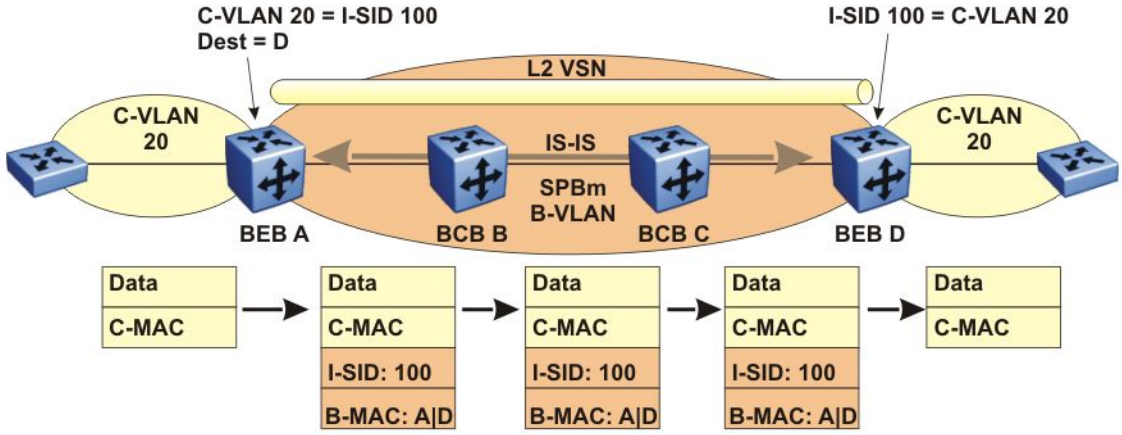

Figure 1: SPBM L2 VSN

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

The BEBs provide the boundary between the SPBM domain and the virtualized services domain. For an L2 VSN service, the BEBs map a C-VLAN into 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 L2 VSN.

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

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

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

IS-IS

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

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

Unlike in an IP OSPF environment, the SPBM use of IS-IS does not require transport of any IP addressing for topology calculations. In the SPBM environment for L2 VSNs, IS-IS carries only pure Layer 2 information with no requirement for an underlying IP control plane or forwarding path. IS-IS runs directly over Layer 2.

Note:
(SPBM carries L3 information for L3 VSNs.)

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

- Discovers the network topology
- Builds shortest path trees between the network nodes:
  - Used for forwarding unicast traffic
  - Used for determining the forwarding table for multicast traffic
- Communicates network information in the control plane:

  Instance Service Identifier (I-SID) information

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. For example, 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 PDUs (LSP). When a new service instance is provisioned on a node, its membership is flooded throughout the topology using an IS-IS advertisement.

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

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

- Flooding is disabled
- Broadcasting is disabled
- Source address learning is disabled
- Unknow mac discard is disabled

Essentially the VLAN becomes a header indicating the SPBM network to use.

Modification of the VLAN behavior is necessary to ensure proper control over the SPBM traffic.

Pre-populated FIB

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

SPBM does not allow any broadcast flooding of traffic on the B-VLAN in order to prevent looping accomplished through flooding packets with unknown destinations (although multicast traffic is supported). As such, MAC addresses must be distributed within SPBM. This is accomplished by carrying the necessary B-MAC addresses inside the IS-IS link state database. To that end, SPBM supports an IS-IS TLV that advertises the I-SID and B-MAC information across the network. This functionality enables the powerful end-point-provisioning of SPBM.

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

RPFC

A loop prevention mechanism is required at Layer 2 to stop wayward traffic from crippling the network. Reverse Path Forwarding Check (RPFC) is the chosen method of squelching 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 squelched imposing a guaranteed loop control. If there's no VLAN FDB entry to the source MAC address with the outgoing port as the ingress port, the frame will be dropped.

SPBM sample operation—L2 VSN

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

1. *Discover network topology*

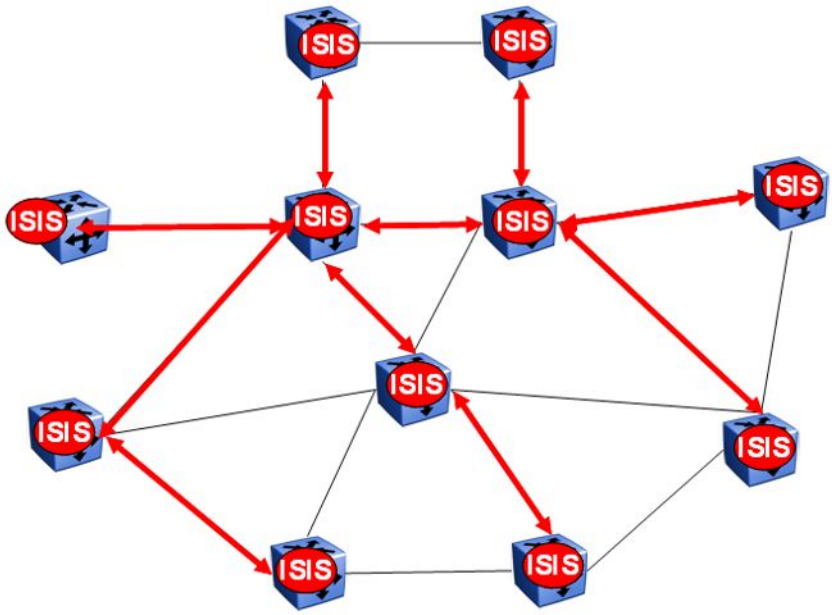

Figure 2: SPBM topology discover

IS-IS runs on all nodes of the SPBM domain. Since IS-IS is the basis of SPBM, the IS-IS adjacency must be formed first. Once the neighboring nodes see hellos from

each other they will 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 will send Link State Protocol Data Units (LSPDUs) which contain connectivity information for the SPBM node. These nodes will also send copies of all other LSPDUs they have in their databases. This establishes a network of connectivity providing the necessary information for each node to find the best and proper path to all destinations in the network.

Each node has a System ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network.

2. *Each IS-IS node automatically builds trees from itself to all other nodes*

   When the network topology is discovered and stored in the IS-IS link state database (LSDB), each node calculates shortest path trees for each source node. A unicast path now exists from every node to every other node

   With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes. Multicast FIB is not produced until L2 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.

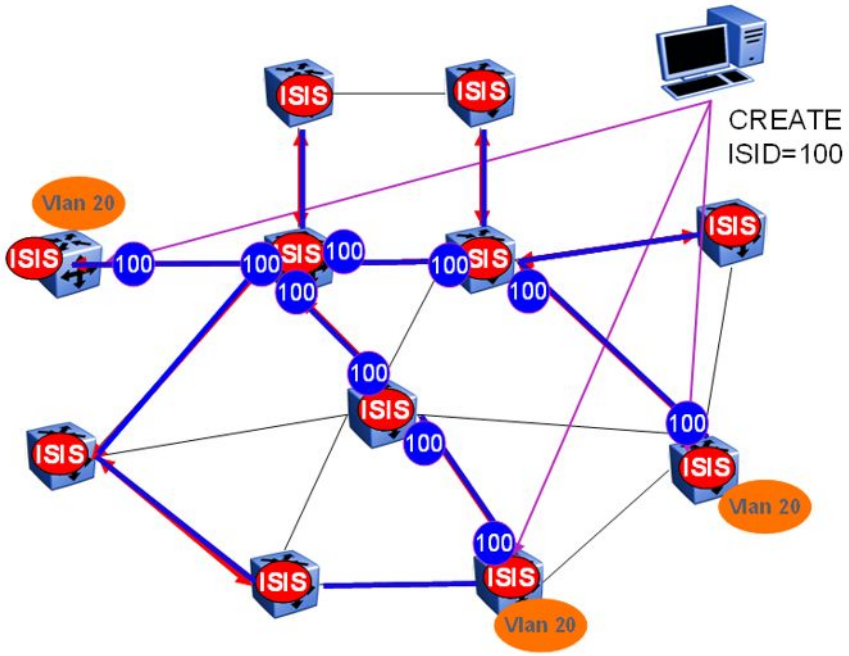

Figure 3: SPBM BMAC and I-SID population

BMAC and I-SID information is flooded throughout the network to announce new I-SID memberships. In this case, VLAN 20 is mapped to I-SID 100.

Note:

I-SIDs are only used for virtual services (L2 VSNs and L3 VSNs). If Native IP Shortcuts only is enabled on the BEBs, I-SIDs are never exchanged in the network as IP Shortcuts allows for the GRT IP networks to be transported across IS-IS.

Each node populates its FDB with the BMAC information derived from the IS-IS shortest path tree calculations. Thus there is no traditional flooding and learning mechanism in place for the B-VLAN, but FDBs are programmed by the IS-IS protocol.

4. *When a node receives notice of a new service AND is on the shortest path, it updates the FDB*

   In this scenario, where there are three source nodes having a membership on I-SID 100, there are three shortest path trees calculated (not counting the Equal Cost Trees (ECTs)). The following diagrams depict the traffic flow for this formed ELAN.

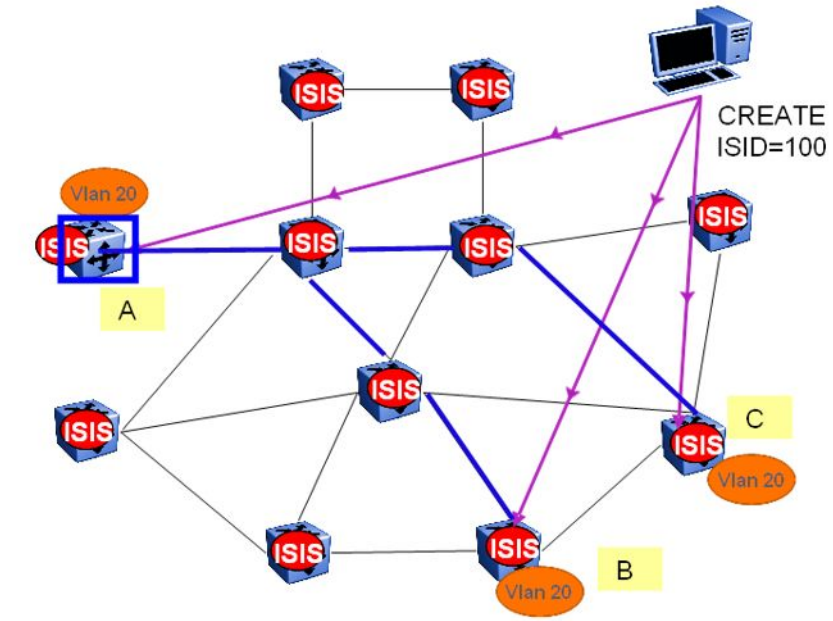

Figure 4: Shortest path tree for source node A

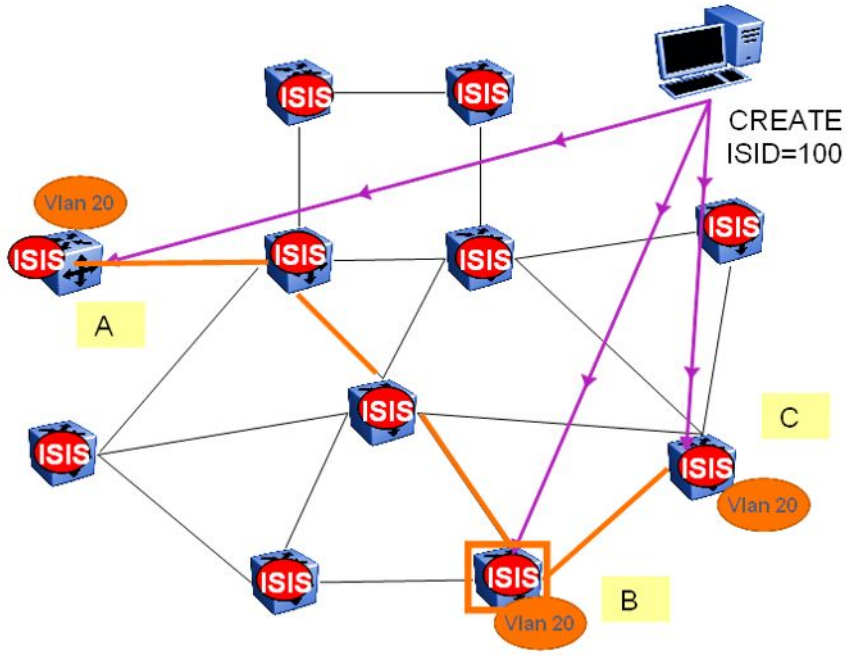

Figure 5: Shortest path tree for source node B

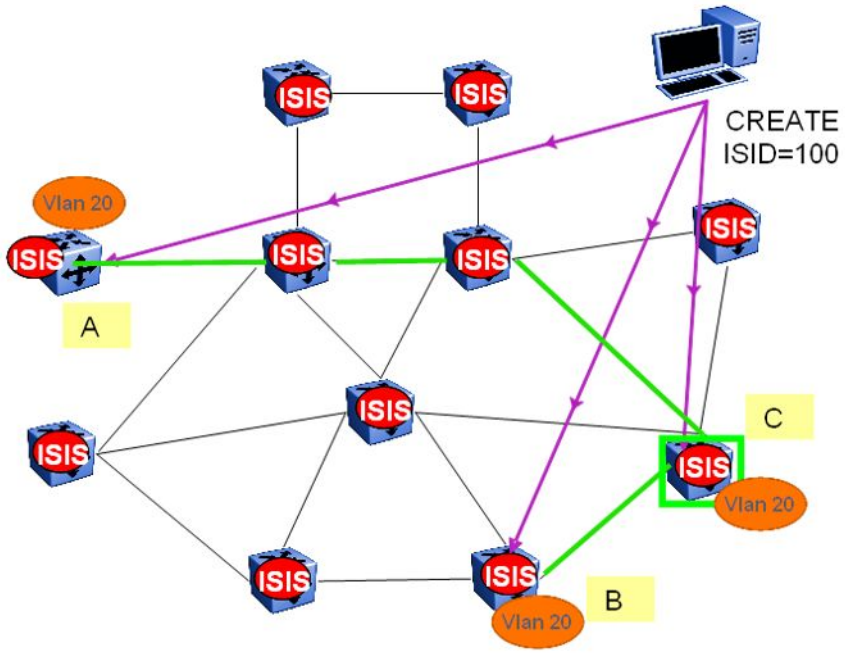

Figure 6: Shortest path tree for source node C

The paths between any two nodes are always the shortest paths. Also, the paths in either direction are congruent, thus a bidirectional communication stream can be monitored easily by mirroring ingress and egress on a link to a network analyzer.

VLAN traffic arriving on switch A and VLAN 20 is forwarded following the blue path, traffic arriving on switch B and VLAN 20 the orange path and on switch C VLAN 20 traffic is following the green path.

If the destination CMAC is unknown at the SPBM ingress node or the traffic is of type broadcast or multicast, then it is flooded to all members of the topology which spans VLAN 20. If the destination CMAC is already known, then the traffic is only forwarded as a unicast to the appropriate destination. In the SPBM domain, the traffic is switched on the BMAC header only. The bridge filtering database (FDB) at the VLAN to I-SID boundary (backbone edge bridge BEB), maintains a mapping between CMACs and corresponding BMACs.

For example, Switch B learns all CMACs which are on VLAN 20 connected to switch A with the BMAC of A in its FDB and the CMACs which are behind C are learnt with the BMAC of C.

SPBM L2 VSN

SPBM supports L2 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.

In the backbone VLAN (B-VLAN), Backbone Core Bridges (BCBs) forward the encapsulated traffic based on the BMAC-DA, using the shortest path topology learned using IS-IS.

The following figure shows a sample campus SPBM L2 VSN network.

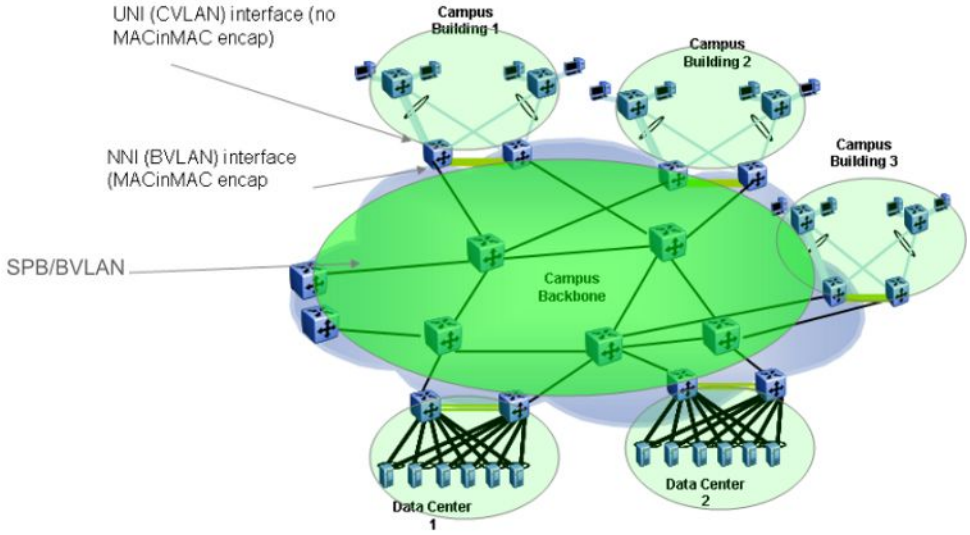

Figure 7: SPBM L2 VSN in a campus

One of the key advantages of the SPBM L2 VSN is that network virtualization provisioning is achieved by configuring only the edge of the network (BEBs). As a result, the intrusive core provisioning that other Layer 2 virtualization technologies require is not needed when new connectivity services are added to the SPBM network. For example, when new virtual server instances are created and need their own VLAN instances, they are provisioned at the network edge only and do not need to be configured throughout the rest of the network infrastructure.

Based on its I-SID scalability, this solution can scale much higher than any 802.1Q tagging based solution. Also, due to the fact that there is no need for Spanning Tree in the core, this solution does not need any core link provisioning for normal operation.

Redundant connectivity between the C-VLAN domain and the SPBM infrastructure can be achieved by operating two SPBM switches in Switch Clustering (SMLT) mode. This allows the dual homing of any traditional link aggregation capable device into an SPBM network.

SPBM Native IP shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM Native IP shortcuts.

Unlike L2 VSN, with SPBM Native IP shortcuts, no I-SID configuration is required. Instead, SPBM nodes propagate Layer 3 reachability as “leaf” information in the IS-IS LSPs using Extended IP reachability TLVs (TLV 135), which contain routing information such as neighbors and locally configured subnets. SPBM nodes receiving the reachability information can use this information to populate the routes to the announcing nodes. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes.

The following figure shows a network running SPBM Native IP shortcuts.

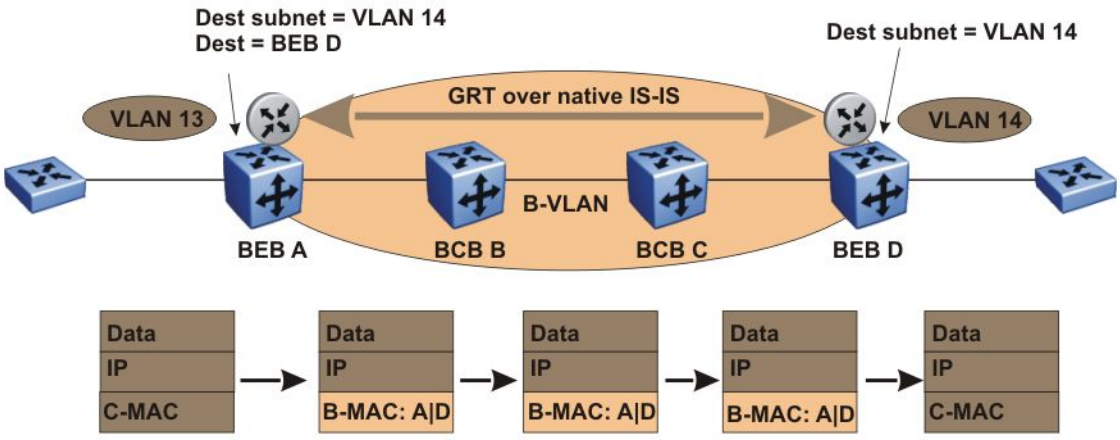

Figure 8: SPBM Native IP shortcuts

In this example, BEB A receives a packet with a destination IP address in the subnet of VLAN 14 and knows to forward the packet to BEB D based on the IP route propagation within IS-IS. After a route lookup, BEB A knows that BEB D is the destination for the subnet and constructs a new B-MAC header with destination B-MAC: D. BCBs B and C need only perform normal Ethernet switching to forward the packet to BEB D. A route lookup is only required once, at the source BEB, to identify BEB D as the node that is closest to the destination subnet.

In contrast to IP routing or MPLS, SPBM Native IP shortcuts provides a simpler method of forwarding IP packets in an Ethernet network using the preestablished Ethernet FIBs on the BEBs. SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet

switching based on the existing SPT. This allows for end to end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

In the above example, the SPBM nodes in the core that are not enabled with Native IP shortcuts can be involved in the forwarding of IP traffic. Since SPBM nodes only forward on the MAC addresses that comprise the B-MAC header, and since unknown TLVs in IS-IS are relayed to the next hop but ignored locally, SPBM nodes need not be aware of IP subnets to forward IP traffic.

With Native IP shortcuts, there is only one IP routing hop, as the SPBM backbone acts as a virtualized switching backplane.

The following figure shows a sample campus network implementing SPBM Native IP shortcuts.

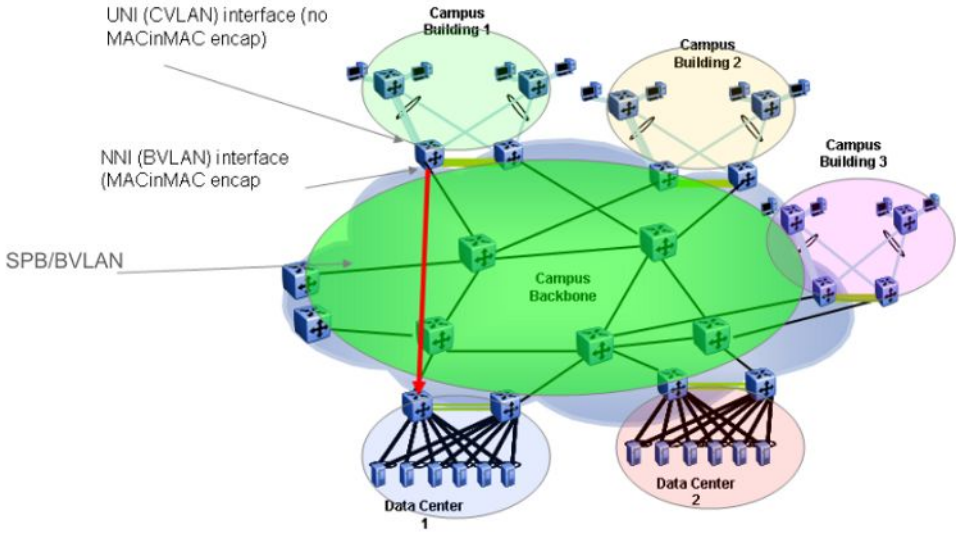

Figure 9: SPBM Native IP shortcuts in a campus

To enable IP shortcuts on the BEBs, you must configure a circuitless IP address (loopback address) and specify this adress as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct, static, OSPF, RIP, or BGP routes into IS-IS.

ECMP

It is possible for the same route to be announced by multiple BEBs, either because the Layer 2 LAN is connected to multiple BEBs for redundancy or because its segments are L2 bridged. In either case, the receiving BEB has to tie-break between the BEBs or add ECMP routes. The first criterion to tie-break is to compare the IS-IS cost to each BEB, for example, choose the route from a BEB that is less number of hops away. If that cost is the same, then compare the metric of the route in the IP reachability information TLV to break the tie. If that metric the same

as well, then these routes will be added in the Route Table Manager (RTM) as ECMP routes, pointing to different BEBs.

The above ECMP is based on SPBM running a single B-VLAN in the core and thus having a single path in the core to get to any other BEB. For SMLT on the access or for load-sharing purposes, the SPBM has to support two B-VLANs and a different tie-breaking rule for SPBM for each of them (higher node ID for one B-VLAN and the lower ID for the other BVID). Note that, this tie-breaking is for the nodes in the SPBM core. With two B-VLANs in the SPBM core, there will always be two paths to reach a particular SPBM node, one on each B-VLAN, therefore, any IP prefix received from a BEB will result in two ECMP paths. These paths might or might not be physically diverse.

If ECMP is enabled with Native IP Shortcuts, IP ECMP must also be enabled.

SPBM L3 VSN

The SPBM L3 VSN feature is a mechanism to provide IP connectivity over SPBM for VRFs. SPBM L3 VSN uses IS-IS to exchange the routing information for each VRF.

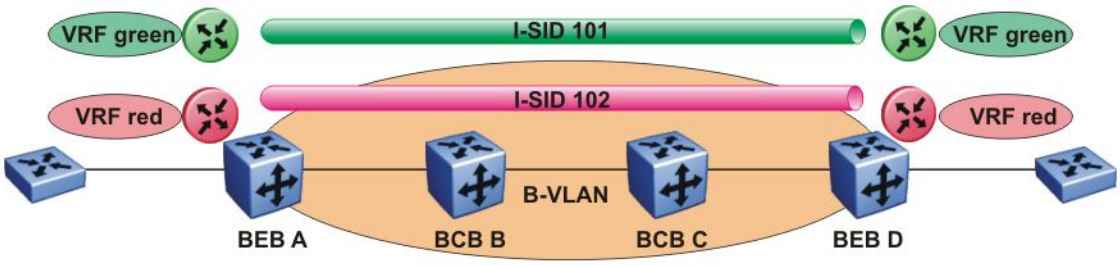

Figure 10: SPBM L3 VSN

In the preceding figure, the BEBs are connected over the SPBM cloud running IS-IS. VRF red and green are configured on the BEBs. VRF red on BEB A has to send and receive routes from VRF red on BEB D. Similar operations are required for VRF green on BEB A and BEB D.

IS-IS TLV 184 is used to advertise SPBM L3 VSN route information across the SPBM cloud. To associate advertised routes with the appropriate VRF, each VRF is associated with an I-SID. All VRFs in the network that share the same I-SID participate in the same VSN.

In this example, I-SID 101 is associated with VRF green and I-SID 102 is associated with VRF red. The I-SID is used to tie the advertised routes to a particular VRF. This identifier has to be the same on all edge nodes for a particular VRF, and has to be unique across all the VRFs on the same node

When IS-IS receives an update from an edge node, it looks for the L3 VSN TLV, and if one exists, it looks at the I-SID identifier. If that identifier is mapped to a local VRF, it extracts the IP routes and adds them to the RTM of that VRF; otherwise the TLV is ignored.

In the VRF, just like in the Network Routing Engine (NRE), the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

With SPBM L3 VSN, the packet forwarding works in a similar fashion as the Native IP Shortcuts on the NRE, with the difference that the encapsulation includes the I-SID to identify the VRF that the packet belongs to. The following figure shows the packet forwarding for VRF red.

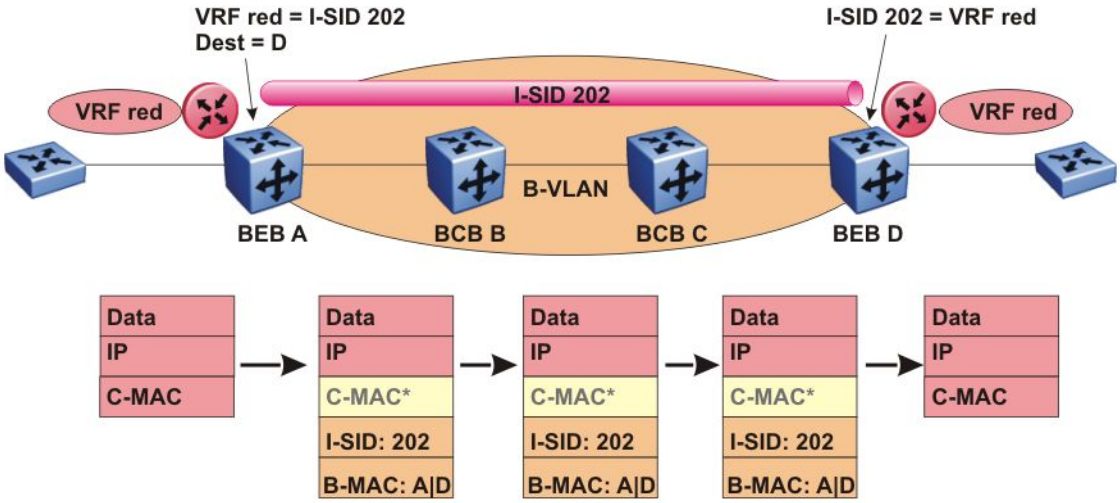

Figure 11: Packet forwarding in SPBM L3 VSN

When BEB A receives traffic from VRF red that must be forwarded to the far-end location, it performs a lookup and determines that VRF red is associated with ISID 202 and that BEB D is the destination for ISID 202. BEB A then encapsulates the IP data into a new B-MAC header, using destination B-MAC: D.

Note:

With SPBM L3 VSN, the CMAC header is all null. This header does not have any significance in the backbone. It is included to maintain the same 802.1ah format for ease of implementation.

At BEB D, the node strips off the B-MAC encapsulation, and performs a lookup to determine the destination for traffic with ISID 202. After identifying the destination as VRF red, the node forwards the packet to the destination VRF.

IP VPN Lite over SPBM

One of the other VPN models supported in the Avaya Ethernet Routing Switch 8800/8600 is the IP VPN Lite model. This is the traditional 2547 kind of VPN with BGP running between the NREs to exchange routes between the VRFs. The VRFs themselves are identified by the BGP communities configured for the VRFs. However, unlike the 2547 model, the Lite model does not use MPLS labels to identify the edge node or the VRF, nor does it use the MPLS transport. Instead it maps a service label, which is an IP address, per VRF and uses IP-in-IP encapsulation with the outer IP being the service label for the VRF.

For more information on IP VPN Lite, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP VPN* (NN46205–520).

In the Lite model, the IP-in-IP packet is routed in the core. However, with SPBM running in the core and the NRE Route Table Manager (RTM) exchanging the service label IPs, then the IP-in-IP packet can be switched in the SPBM core by simply doing a lookup of the service label and applying the B-MAC header corresponding to the particular IP service label.

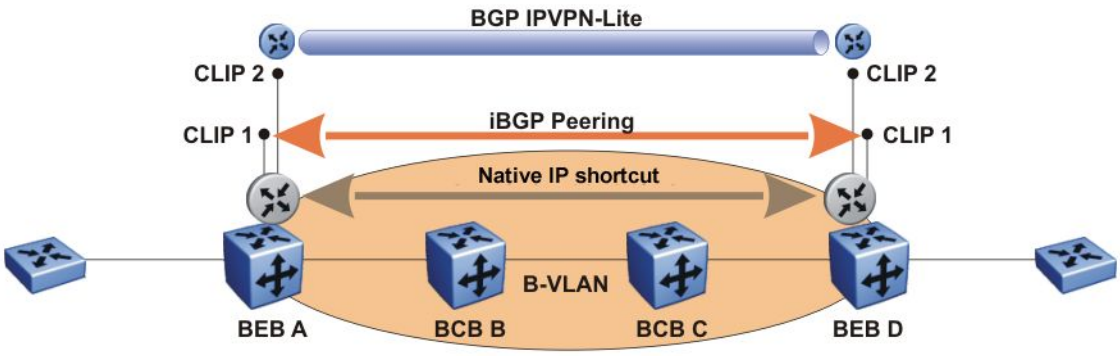

Figure 12: IP VPN Lite over SPBM

Note:

The header encapsulation in this mode is a regular Ethernet header with IP in IP, without the 802.1ah header.

In the above example, BGP configuration is only required on the BEB nodes. The BCB switches have no knowledge of any Layer 3 VSN IP addresses or routes, and just forward traffic based on the B-MAC header.

Inter-ISID routing

Inter-ISID routing with SPBM allows routing between Layer 2 VLANs with different I-SIDs.

The following figure shows a sample Inter-ISID routing topology.

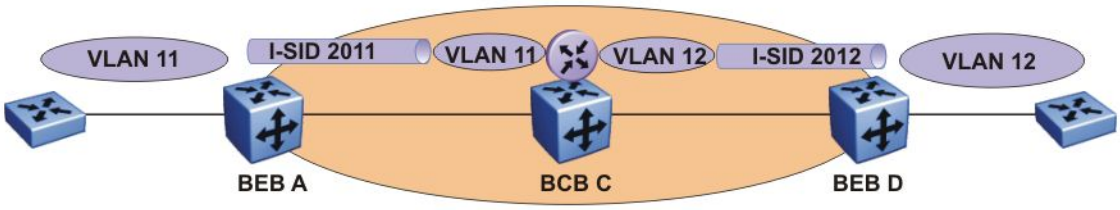

Figure 13: Inter-ISID routing

In this example, the C-VLANs are associated with I-SIDs on the BEBs using SPBM L2 VSN. With inter-ISID routing enabled, BCB C can transmit traffic between VLAN 11 (I-SID 2011) and VLAN 12 (I-SID 2012). The BEB switches can forward traffic between VLANs 11 and 12 on the VRF instance configured on the BCB.

While this example illustrates a VRF configured on a core BCB switch, inter-ISID can also be configured on the GRT.

IS-IS IP redistribution policies

When interconnecting an SPBM Core using Native IP shortcuts (or L3 VSNs) to existing networks running a routing protocol such as OSPF or RIP, a redundant configuration requires two ERS 8800 or 8600 routers.

- One router redistributes IP routes from RIP/OSPF into IS-IS (IP).
- The second router redistributes from IS-IS (IP) into RIP/OSPF.

The following figure illustrates this configuration.

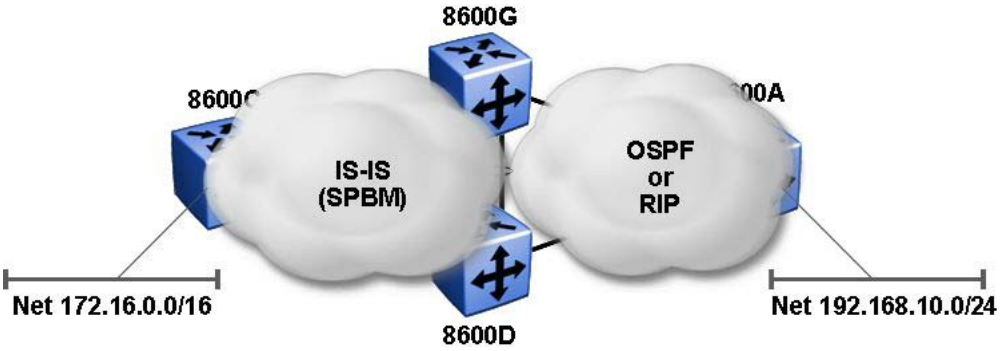

Figure 14: Redundant OSPF or RIP network

In this scenario it is necessary to take extra care when redistributing through both ERS 8600s. By default the preference value for IP routes generated by SPBM-IP (IS-IS) is 7. This is a higher preference than OSPF (20 for intra-area, 25 for inter-area, 120 for ext type1, 125 for ext type2) or RIP (100).

Important:

The lower numerical value determines the higher preference.

In the diagram above both nodes (8600G and 8600D) have an OSPF or a RIP route to 192.168.10.0/24 with the next-hop to 8600A.

As soon as the 8600G node redistributes that IP route into IS-IS, the 8600D node learns the same route through IS-IS from 8600G. (The 8600G node already has the route through OSPF or RIP). Because IS-IS has a higher preference, 8600D replaces its 192.168.10.0 OSPF route with an IS-IS one pointing at 8600G as the next-hop. The following figure illustrates this scenario.

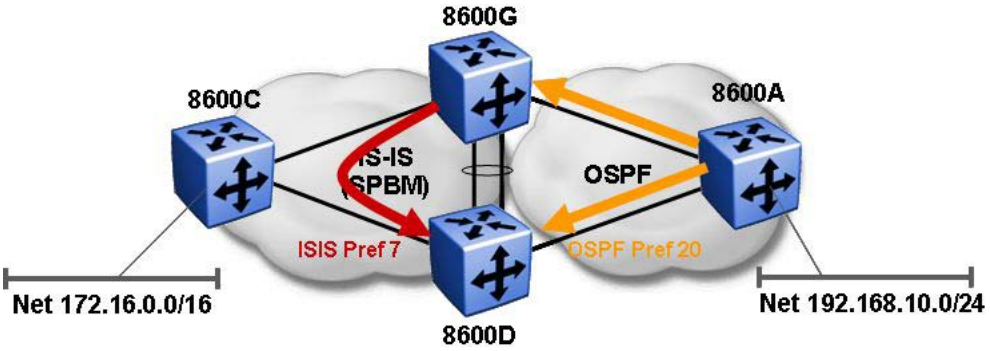

Figure 15: Redistributing routes into IS-IS

Clearly this is undesirable and care needs to be taken to ensure that the two redistributing nodes (8600G and 8600D) do not accept redistributed routes from each other. IS-IS accept policies are not available in the initial 7.1 release, but they will be available in a subsequent release. With IS-IS accept policies, you can associate an IS-IS accept policy on 8600D to reject all redistributed IP routes received from 8600G, and vice-versa.

An alternative way to solve the above problem with 7.1 functionality is to reverse the problem by lowering the SPBM-IP (IS-IS) preference by configuring it to a value greater than RIP (100) or OSPF (20,25,120,125). For example, use the following command to set a preference of 130:

```
ip route preference protocol spbm-level1 130
```

Or, in the case of an L3 VSN:

```
ip vrf <vrf-name>route preference protocol spbm-level1 130
```

Now that the OSPF or RIP routes have a higher preference than SPBM-IP (IS-IS), the above problem is temporarily solved. However, the same issue resurfaces when the IS-IS IP routes are redistributed into OSPF or RIP in the reverse direction as shown in the following figure for OSPF:

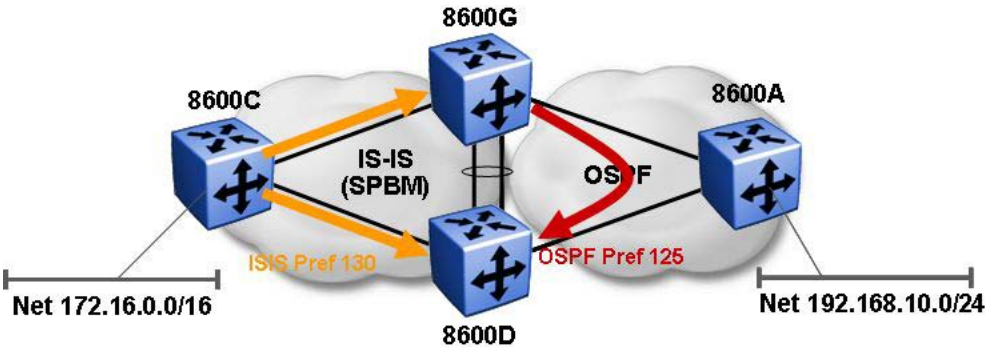

Figure 16: Redistributing routes into OSPF

In the above figure, both 8600G and 8600D have an IS-IS IP route for 172.16.0.0/16 with the next hop as 8600C. As soon as the 8600G node redistributes the IP route into OSPF, the 8600D node learns that same route through OSPF from 8600G. (The 8600G node already has the route through IS-IS).

Because OSPF has a higher preference, 8600D replaces its 172.16.0.0/16 IS-IS route with an OSPF one. (Note that the 172.16.0.0/16 route will be redistributed into OSPF as an AS external route, hence with preference 120 or 125 depending on whether type1 or type2 was used). In this case, however, we can leverage OSPF Accept policies, which can be configured to prevent 8600D from accepting any AS External (LSA5) routes from 8600G and vice-versa. The following is a sample configuration:

```
ip route-policy "reject" seq 1 create
ip route-policy "reject" seq 1 enable
ip route-policy "reject" seq 1 action deny
ip ospf accept adv-rtr <router-id of other router> create
ip ospf accept adv-rtr <router-id of other router> enable
ip ospf accept adv-rtr <router-id of other router> route-policy "reject"
 ip ospf accept adv-rtr <router-id of other router> apply
```

In the case of a RIP access network, the above is not possible because RIP has no concept of external routes and no equivalent of accept policies. However, if you assume that a RIP network acts as an access network to an SPBM core, then it is sufficient to ensure that when

IS-IS IP routes are redistributed into RIP they are aggregated into a single default route at the same time. The following figure and sample configuration example illustrates this scenario:

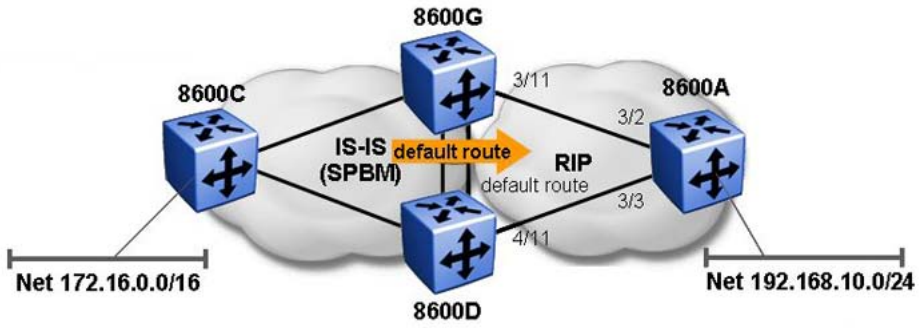

Figure 17: Redistributing routes into RIP

8600G

```
ip prefix-list "default" add-prefix 0.0.0.0/0 maskLenFrom 0 maskLenTo 32
ip route-policy "inject-default" seq 1 create
ip route-policy "inject-default" seq 1 enable
ip route-policy "inject-default" seq 1 match-network "default"
ip route-policy "inject-default" seq 1 set-injectlist "default"
ethernet 3/11 ip rip default-supply enable
ip rip redistribute isis create
ip rip redistribute isis route-policy "inject-default"
ip rip redistribute isis enable
ip rip redistribute isis apply
```

8600A

```
ethernet 3/2, 3/3 ip rip default-listen enable
```

8600D

```
ip prefix-list "default" add-prefix 0.0.0.0/0 maskLenFrom 0 maskLenTo 32
ip route-policy "inject-default" seq 1 create
ip route-policy "inject-default" seq 1 enable
ip route-policy "inject-default" seq 1 match-network "default"
ip route-policy "inject-default" seq 1 set-injectlist "default"
ethernet 4/11 ip rip default-supply enable
ip rip redistribute isis create
ip rip redistribute isis route-policy "inject-default"
ip rip redistribute isis enable
ip rip redistribute isis apply
```

You can control the propagation of the default route on the RIP network so that both 8600G and 8600D supply the default route on their relevant interfaces, and not accept it on the same interfaces. Likewise, 8600A will accept the default route on its interfaces to both 8600G and 8600D but it will not supply the default route back to them. This will prevent the default route advertised by 8600G from being installed by 8600D, and vice-versa.

The above example where IS-IS IP routes are aggregated into a single default route when redistributed into the RIP network can also be applied when redistributing IS-IS IP routes into OSPF if that OSPF network is an access network to an SPBM Core. In this case use the

following redistribution policy configuration as an example for injecting IS-IS IP routes into OSPF:

```
ip prefix-list "default" add-prefix 0.0.0.0/0 maskLenFrom 0 maskLenTo 32
ip route-policy "inject-default" seq 1 create
ip route-policy "inject-default" seq 1 enable
ip route-policy "inject-default" seq 1 match-network "default"
ip route-policy "inject-default" seq 1 set-injectlist "default"
ip ospf as-boundary-router enable
ip ospf redistribute isis create
ip ospf redistribute isis route-policy "inject-default"
ip ospf redistribute isis enable
ip ospf redistribute isis apply
```

Replication in SPBM backbone

In order to flood C-VLAN broadcast and multicast packets within an I-SID, the backbone needs to replicate packets. To achieve this, SPBM uses Ethernet-based L2 Multicast forwarding. Multicast addresses are built out of two pieces. The first portion contains the unique nickname that must be assigned to each SPBM node. The second portion of the Multicast MAC address is the I-SID converted to hex format.

For example, if a node has a nickname of 03.00.61 and I-SID of 100 (hex 0x064), its Multicast MAC address is 33:00:61:00:00:64 where the 03 is shifted by 4 to form 30:00:61:00:00:64 , and another 3 indicating multicast bit and local generated.

```
# show isis spbm multicast-fib
********************************************************************************
Command Execution Time: SUN SEP 26 20:35:36 2010 UTC
********************************************************************************
================================================================================
SPBM MULTICAST FIB ENTRY INFO
================================================================================
MCAST DA ISID BVLAN SYSID HOST-NAME OUTGOING
-INTERFACES
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Total number of SPBM MULTICAST FIB entries 3
--------------------------------------------------------------------------------
03:00:61:00:00:64 100 10 0080.2dc1.37ce 8800-1 4/7
03:00:61:00:00:c8 200 10 0080.2dc1.37ce 8800-1 4/2,4/
```

The multicast addresses in the core of the network apply to a per node, per service tree. This means if a node is provisioned with I-SID 200 in addition to I-SID 100, this will result in two separate multicast trees. Each tree will be sourced from this node, A, reaching every other node with the same service termination.

Keep in mind the multicast tree applies to traffic from the source node and only to the other nodes with I-SID endpoints on the same service. Each node with endpoints on the I-SID will have its own individual multicast tree.

When a node recognizes that it is in the path for the I-SID service it then builds the individual multicast MAC addresses.

Multicast, broadcast, or unknown unicast customer frame forwarding

When a known unicast flow occurs then the unicast nodal MAC addresses are used for forwarding through the SPBM core.

If an unknown unicast, multicast or broadcast customer frame enters the I-SID service, SPBM forwards the packet on the individual multicast tree created for the source node. In the core the data frame will have the source node's nodal B-MAC address and the destination will be the multicast MAC address.

SPBM MGID usage

The multicast group ID (MGID) is a hardware mechanism the switch uses to send data to several ports simultaneously. Instead of sending the data to a specific port number, the data is directed to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs. The system also reserves a small number of MGIDs.

SPBM also requires MGIDs for proper operation. When SPBM is enabled on the switch, the system reserves 519 MGIDs for SPBM operation. Therefore, the number of MGIDs on the system available for VLANs and IP multicast traffic is reduced by 519. To determine how many MGIDs are available, enter **show sys mgid-usage**.

```
ERS-8606:5# show sys mgid-usage
********************************************************************************
Command Execution Time: WED FEB 02 19:42:27 2011 UTC
********************************************************************************
Number of MGIDs used for VLANs : (3)
Number of MGIDs used for SPBM : (519)
Number of MGIDs used for multicast : (0)
Number of MGIDs remaining for VLANs : (1514)
Number of MGIDs remaining for multicast : (2048)
```

Before you enable SPBM on the switch, be sure that your network will not be adversely affected by this reduction in available MGIDs.

The Ethernet Routing Switch 8800/8600 supports a total of 4096 MGIDs, split between the system, VLAN, IPMC, and now SPBM. You can reserve MGIDs for IP Multicast (IPMC) traffic. You can reserve between 64 and 4084 MGIDs for IPMC. The default for IPMC is 2048. It is the responsibility of the network administrator to fully understand the network deployment

strategy. Please ensure that MGIDs are planned appropriately. If assistance is required, please contact your Avaya technical representative.

For information about reserving MGIDs for IPMC, see *Avaya Ethernet Routing Switch 8800/8600 Administration* (NN46205–605).

Chapter 3: IS-IS fundamentals

This chapter provides additional details about IS-IS operation and parameters.

IS-IS Overview

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).

IS-IS Operation

IS-IS operation is similar to Open Shortest Path First (OSPF). Both protocols divide large domains into smaller areas, and both use the shortest path first (SPF) algorithm and link state information to determine the best path to a destination. The following sections describe these concepts in detail.

IS-IS System Identifiers

The IS-IS system identifiers consist of three parts:

- Manual area — The manual area or area ID is up to 13 bytes long. The first byte of the area number (for example, 49) is the Authority and Format Indicator (AFI). The next bytes are the assigned domain (area) identifier, which is up to 12 bytes (for example, 49.0102.0304.0506.0708.0910.1112).
- System ID — The system ID is any 6 bytes that are unique in a given area or level. The system ID defaults to the node BMAC.
- NSEL — The last byte (00) is the n-selector. In the Avaya Ethernet Routing Switch 8800/8600 implementation, this part is automatically attached. There is no user input accepted.

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 prerequisites 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.

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 the domain's backbone, connecting to all the Level 1 areas.

Level 1 routers route only within their assigned area, and cannot route outside of that area. When the destination is to another area, they route toward the closest Level 2 router. Each Level 1 router in the area must have a unique system ID relative to the others.

Level 2 routers route between areas and toward other domains. They do not have any capabilities to route within an area. Each Level 2 router in a domain must have a unique system ID relative to the others.

Important:

The Avaya Ethernet Routing Switch 8800/8600 is a Level 1 router, which means it has only Level 1 links and can route within only one area.

IS-IS Information Exchange

IS-IS uses the following packet formats to exchange information within an area and between areas:

- Intermediate System to Intermediate System Hello (IIH) packets contain addresses of the interface over which the Hello is transmitted.

IS-IS broadcasts IIH packets to discover neighboring IS-IS routers, and to determine whether the neighbors are Level 1 or Level 2 routers. This is how IS-IS initializes and maintains adjacencies between neighboring routers.

- Link State Packets (LSP) contain information about the state of adjacencies or defined and distributed static routes. IS-IS exchanges this information with neighboring IS-IS routers by flooding LSPs throughout an area at periodic intervals. Every router in the domain has an identical link state database and each runs SPF to calculate routes.
- Complete Sequence Number Packets (CSNP) contain the most recent sequence numbers of all LSPs in the database. When all routers update their LSP database, synchronization is complete.
- Partial Sequence Number Packets (PSNP) are requests for missing LSPs. When a receiving router detects that it is missing an LSP, it sends a PSNP to the router that sent the CSNP.

Note:

If the switch has a large number of MACs, a burst of aging activity can cause IS-IS adjacencies to bounce. To prevent this from happening, stagger the aging times on a per VLAN basis and vary the number of seconds as much as possible.

The aging timer command is **`config vlan <vid> fdb-entry aging-time <seconds>`**, and the range of values for *seconds* is 10..1000000. The default is 300. For more information about configuring VLANs, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — VLANs and Spanning Tree* (NN46205-517).

IS-IS Designated Router Election

In IS-IS terminology, the designated router is called a Designated Intermediate System (DIS). 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.

Note:

Electing a DIS affects broadcast interfaces only. SPBM supports point-to-point links only, and they do not need DIS election.

Each link elects its own DIS, and that DIS is responsible for the synchronization of the link state database between all the linked routers at that level. The DIS multicasts CSNPs periodically, and the receiving routers use them to synchronize their LSP metric. The LSP metric represents the cost of the link, which IS-IS uses to determine the best path to the destination router. IS-IS assigns a metric to each interface and then calculates the routing path by adding the metric value for each interface.

You can configure the metric on each interface. Narrow type metrics support values from 0 to 63. Wide metrics accept values up to 16777215. The Avaya Ethernet Routing Switch 8800/8600 accepts wide metrics only.

IS-IS Parameters

The following sections describe the IS-IS parameters in more detail.

PSNP interval

When a router detects a missing LSP, the router sends a PSNP to the system that transmitted the Complete Sequence Number Packet (CSNP). The transmitting router then forwards the missing LSP to the requesting router. The PSNP interval sets the interval for how long a router waits before sending a PSNP.

You can change the Partial Sequence Number Packet (PSNP) interval rate. A longer interval reduces overhead, while a shorter interval speeds up convergence.

CSNP periodic and interval rate

CSNPs contain the most recent sequence numbers of all LSPs in the database. Routers usually send CSNPs periodically on all interfaces. Other routers use CSNPs to update and synchronize their LSP databases. When the routers in the domain receive all the LSPs, synchronization is complete.

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 link state packet (LSP) to reduce overhead or speed up convergence.

- The max-lsp-gen-interval is the maximum interval between generated LSPs. The default is 900 seconds with a range of 0 to 900.
- The min-lsp-gen-interval is the minimum amount of time between successive generations of LSP packets with the same LSP ID. The Avaya Ethernet Routing Switch 8800/8600 allows for the configuration of this parameter at the system level and applies it to Level 1 LSP generation. This interval prevents flapping of the interface. The default is 30 seconds with a range of 1 to 65535. The maximum range value of 65535 is implemented for future compatibility. Currently, the min-lsp-gen-interval cannot exceed the max-lsp-gen-interval. The maximum range value for the max-lsp-gen-interval is 900.

Important:

If you set the min-lsp-gen-interval to a high value, the IS-IS network will not converge. The max-lsp-gen-interval of 900 and minimum of 899 can be configured, but network convergence will not occur.

- The retransmit-lspint is the minimum amount of time between retransmission of an LSP. This defines how quickly the same LSP is resent. The Avaya Ethernet Routing Switch 8800/8600 allows the configuration of this parameter at the system level and applies it to Level 1 retransmission of LSPs. The default is 5 seconds with a range of 1 to 300.

Point-to-point mode

By default, IS-IS runs in point-to-point mode over all interfaces. This option cannot be modified.

When configuring point-to-point interfaces, be sure that both ends of the network connection are set to the same interface type.

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. Authentication checks that each Hello, PSNP, and CSNP packet contains authentication information on the originator or the contents of the packet. 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 an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum. There is an optional key ID.

Password considerations

The passwords for all authentications are saved as cleartext in the configuration file on the Avaya Ethernet Routing Switch 8800/8600. The passwords for simple and HMAC-MD5 are displayed in cleartext through the CLI. The HMAC-MD5 packet is encrypted when transmitted over the network.

If you change passwords for any authentication type, the change does not take effect until the LSP entry in the database is refreshed. There are two ways to refresh:

- Wait for the LSDB timer to expire, which is 1200 seconds.
- Disable and then enable IS-IS globally.

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

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. IS-IS Hello packets contain the IP address of the interface over which the Hello is transmitted. These packets are broadcast to discover the identities of neighboring IS-IS systems and to determine whether the neighbor is a Level 1 router.

You can configure the interface level Hello interval to change how often Hello packets are sent out from an interface level. The configured Hello interval is also used to calculate the DIS Hello interval, which is used when the router is a Designated Intermediate System (DIS). The dis-hello-interval equals the hello-interval, divided by three. For example, if the Hello interval is 9, then the DIS interval is 3.

Hello multiplier

You can configure the Hello multiplier to specify how long the Avaya Ethernet Routing Switch 8800/8600 must wait before considering a neighboring router down. By default, the hold time is three times the Hello interval. In other words, if the Avaya Ethernet Routing Switch 8800/8600 does not see a Hello for three consecutive Hello intervals, it determines that the neighboring router is not operating. For example, 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.

Important:

If the switch has a large number of MACs in the VLAN FDB entry table and the primary SF/CPU fails, it can take longer than 27 seconds for the secondary SF/CPU to become the new primary. In this scenario, the IS-IS adjacency goes down because the default hold time (27 seconds) is too short. To prevent this from happening, increase the Hello multiplier to allow more time for the HA failover.

This hello multiplier is also used to calculate the hold time for the DIS, which does not wait as long as other routers in the area. The dis-hold-time equals the hold time, divided by three. For example, if the hold time is 27, then the DIS hold time is 9.

Cost metric

You can configure the cost metric to overwrite the default metric value derived from reference bandwidth. By configuring the cost metric, you can specify a preferred path, which typically reflects the speed of the transmission media. 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. Wide metrics more accurately reflect the underlying interface speeds.

In this release, only the wide metric is supported.

The metric cost is defined per interface. The cost of the route equals the sum of all the costs assigned to each Level 1 interface along the route. IS-IS always chooses the lowest-cost route available.

The default value for wide metrics is 10.

For loopback interfaces, the default metric value is 0.

Important:

You must disable and then enable the interface for a change in the metric to take effect.

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.

Note that the IS-IS interface must be disabled before you change the interface type.

Important:

If IS-IS is globally disabled, then disabling or enabling at the interface level does not take effect. There is no informative error message indicating that IS-IS is disabled at the global level. However, the administrator can enable IS-IS on the interface level and save the configuration for future use.

Overload bit

If the overload bit parameter is set, the Avaya Ethernet Routing Switch 8800/8600 sets the overload bit in the LSP. The setting affects Level 1 LSPs. The **`overload`** parameter works in conjunction with the **`overload-on-startup`** parameter. When the **`overload-on-startup`**timer expires, the SPBM node clears the overload bit and re-advertises its LSP.

When an LSP with an overload bit is received, the Avaya Ethernet Routing Switch 8800/8600 ignores the LSP in its SPF calculation. By default, overload is set to false. If overload is set to true, the Avaya Ethernet Routing Switch 8800/8600 cannot be a transit node, but it can still receive traffic destined to the ERS.

Chapter 4: CFM fundamentals

The SPBM network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of ping and traceroute. To support troubleshooting of the SPBM cloud, the Ethernet Routing Switch 8800/8600 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. There are two types of MP:

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

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

On the Ethernet Routing Switch 8800/8600, CFM is implemented using the LBM and LTM features only to debug SPBM. CCM messages are not required or supported in the current release.

Maintenance Domain (MD)

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

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

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

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

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

Maintenance Association (MA)

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

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

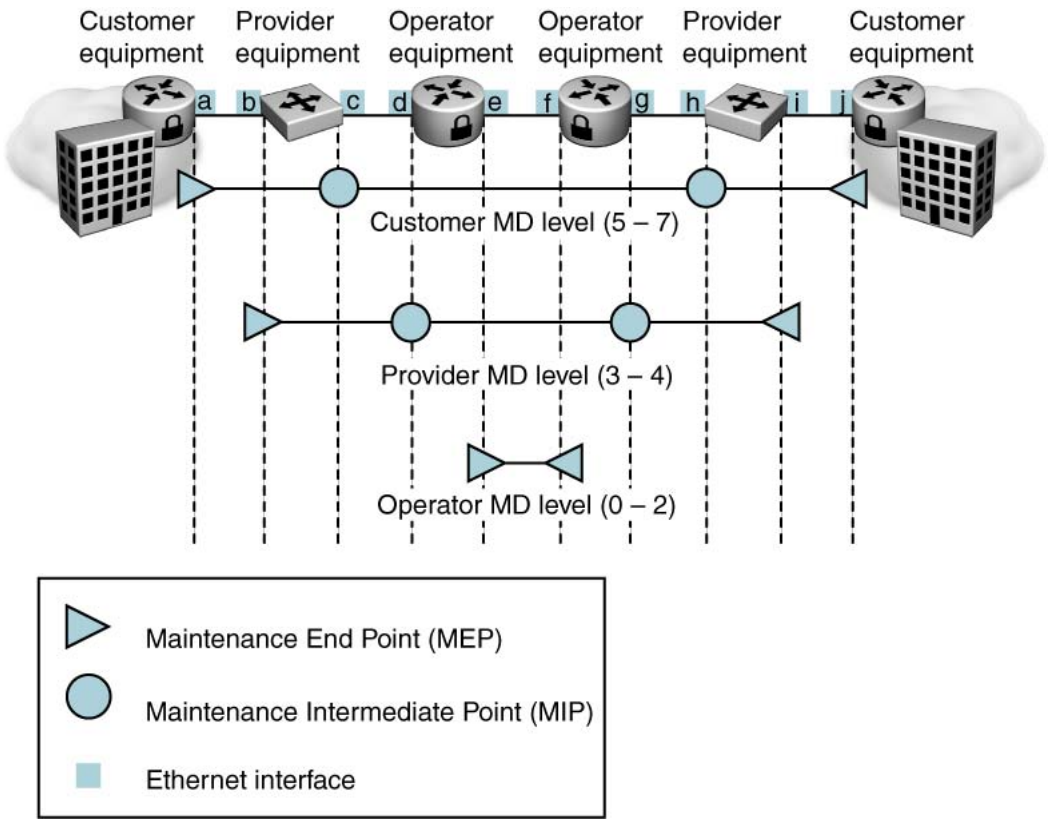

Maintenance association Endpoints (MEP)

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

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

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

The LBM packet is often compared to a ping. A MEP transmits the LBM packet. This packet can be addressed to another MEP or to an MP's MAC address; in the case of SPBM, this will be the SPBM system ID or its virtual SMLT MAC. Only the MP for which the packet is addressed will respond with an LBR message.

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

l2ping

The **`l2ping`** command is a proprietary command that allows a user to trigger an LBM message by specifying the B-VLAN and either the destination MAC address or node name. This provides a simpler command syntax than the standard LBM commands, which require the user to specify the MD, MA, and MEP ID information. The **`l2ping`** command provides a ping equivalent at layer 2 for use with nodes on the SPBM B-VLAN.

l2ping with IP address

The **`l2ping`** command also allows a user to specify an IP address as the destination address. In this case, the IP address must be on a node that has IP shortcuts enabled. The destination IP address is the IP source address that is configured for use with IP shortcuts.

When the **`l2ping`** command is executed with an IP address as the destination, the operation finds all the valid BVID/MAC combinations that provide valid paths to the destination. If ECMP is enabled, there can be multiple paths to the destination. In this case, l2ping is internally run for each of the VLAN/BMAC paths returned, and a summary of results is displayed. If ECMP is disabled, the results display for only one path.

Fault Isolation

Fault isolation is achieved through the use of Linktrace Messages (LTM). LTM is intercepted by all the MPs on the way to the destination MP. The Ethernet Routing Switch 8800/8600 supports two types of LTM.

The first, 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 Replay (LTR), and forwards the original LTM to the destination. The LTM is forwarded until it reaches its destination or the TTL value is decremented to 0. LTR is a unicast message addressed to the originating MEP.

The second, the proprietary LTM, is used to map the MAC addresses of the SPBM network; in this case the target MAC is not an MP, but rather an I-SID.

LTM Message - UNICAST

The link trace message (LTM) is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP which is the SPBM system id or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

- Trace the path to any given MAC address.
- DA is unicast
- LTM contains:
    - Time to live (TTL)
    - Transaction Identifier

  - Originator MAC address
  - Target MAC address
- CFM unaware entities forward the frame as is like any other data frame.
- MIP or MEP that is not on the path to the target discards the LTM and does not reply.
- MIP that is on the path to the target
  - Forwards the LTM after decrementing the TTL and replacing the SA with its own address.
  - Sends a reply (LTR) to the originator.
  - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- If the MIP or MEP is a target
  - Sends a 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.

l2trace

The **l2trace** command is a proprietary command that allows a user to trigger an LTM message by specifying the B-VLAN and either the destination MAC address or node name. This provides a simpler command syntax than the standard LTM commands, which require the user to specify the MD, MA, and MEP ID information. The **l2trace** command provides a traceroute equivalent at layer 2 for use with nodes on the SPBM B-VLAN.

l2trace with IP address

The **l2trace** command also allows a user to specify an IP address as the destination address. In this case, the IP address must be on a node that has IP shortcuts enabled. The destination IP address is the IP source address that is configured for use with IP shortcuts.

When the **l2trace** command is executed with an IP address as the destination, the operation finds all the valid BVID/MAC combinations that provide valid paths to the destination. If ECMP is enabled, there can be multiple paths to the destination. In this case, l2trace is internally run for each of the VLAN/BMAC paths returned, and a summary of the results is displayed. If ECMP is disabled, the results display for only one path.

l2tracetree

The **l2tracetree** command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and ISID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular ISID.

Maintenance domain Intermediate Points (MIP)

MIPs do not initialize any CFM messages. MIPs passively receive CFM messages, process the messages received and respond back to the originating MEP. By responding to received CFM messages, MIPs can support discovery of hop-by-hop path among MEPs, allow connection failures to be isolated to smaller segments of the network to help discover location of faults along the paths. MIPs can be created independent of MEPs. MIP functionality can be summarized as:

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

Nodal MPs

Nodal MPs will 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 will provide traceability and troubleshooting at the system level for a given B-VLAN. Each node (chassis) will have a given MAC address and communicate with other nodes. The SPBM instance MAC address will be used as the MAC address of the Nodal MP. The Nodal B-VLAN MPs will support eight levels of CFM and will be configured on a per B-VLAN basis. Virtual SMLT MAC addresses will also be able to respond for LTM and LBM.

Nodal B-VLAN MEPs

The Nodal B-VLAN MEPs will be created on the CP and will function as if they are connected to the virtual interface of the given B-VLAN. Because of this they will be supported for both port and MLT based B-VLANs. To support this behavior a MAC Entry will be added to the FDB

and a new CFM data-path table containing the B-VLAN and MP level will be added to direct CFM frames to the CP as required.

Nodal B-VLAN MIPs

The Nodal MIP is associated with a B-VLAN; VLAN and level are sufficient to specify the Nodal MIP entity. The Nodal MIP MAC address will be the SPBM system ID for the node on which it resides. If the fastpath sends a message to the CP, the MIP will respond if it is not the target and the MEP will respond if it is the target.

Nodal B-VLAN MEPs and MIPs with SPBM

When Nodal MEPs or MIPs are on SPBM B-VLANs the LTM code will use a unicast MAC DA. The LTM DA will be the same as the target MAC address, which will be the SPBM MAC address or the SMLT MAC address of the target node.

In this release it is a requirement to support SMLT interaction with SPBM. This will be accomplished by using two B-VIDs into the core from each pair of SMLT terminating nodes. Both nodes will advertise the Nodal B-MAC into the core on both B-VIDS. In addition each node will advertise the SMLT virtual B-MAC on one of the two B-VLANs.

The Nodal MEP and MIP will be expanded to respond to both the Nodal MAC address as will as the Virtual SMLT MAC address if both MACs are being advertised on its B-VLAN. In addition a source mode will be added to the LTM and LBM command to use either the Nodal MAC or the SMLT virtual MAC address as the source MAC in the packet.

Configuration considerations

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

- There is a limit of two nodal MEPs and two nodal MIPs on a single switch.
- All nodal MEPs and MIPs are restricted to SPBM BVIDs.
- The Maintenance level for MEPs and MIPs on a given BVID (in a network) must be configured to the same level for them to respond to a given CFM command.

Chapter 5: SPBM configuration examples

The following sections show configuration examples for the various SPBM deployment options.

Note:
For additional configuration examples, including SMLT configuration examples, see *Shortest Path Bridging (802.1aq) for ERS 8800/8600 Technical Configuration Guide* (NN48500-5xx).

Basic SPBM configuration example

The following figure shows a sample greenfield deployment for SPBM.

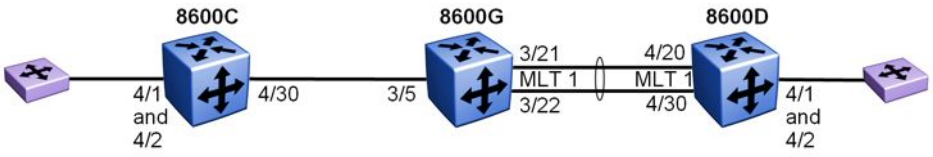

Figure 18: Greenfield SPBM deployment

Note the following:

- By removing the port from STG1 as shown in the configuration, it is also removed from VLAN 1. (This step is not necessary on MLTs, which get added to STGs if added to VLANs.) Otherwise default STG 1 (with VLAN 1) will be running in parallel on core links.
- For migration purposes, SPBM can coexist with existing SMLT or STP config

Ethernet and MLT configuration

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

8600C

```
ethernet 4/30 perform-tagging enable
stg 1 remove ports 4/30 (*)
```

8600G

```
ethernet 3/5 perform-tagging enable
stg 1 remove ports 3/5    (*)
```

```
mlt 1 create
mlt 1 add ports 3/21-3/22
mlt 1 perform-tagging enable
```

8600D

```
mlt 1 create
mlt 1 add ports 4/20,4/30
mlt 1 perform-tagging enable
```

IS-IS SPBM global configuration

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

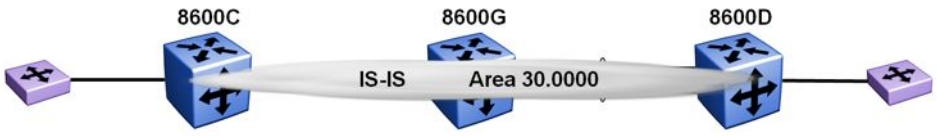

Figure 19: IS-IS SPBM global

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

Note that, if the CLI Prompt has already been set, the IS-IS system name (**`isis sys-name`** command) is automatically set to the same string. The following configuration includes the configuration for the IS-IS system name.

8600C

```
cli prompt 8600C
spbm enable
isis spbm 1 create
isis spbm 1 nick-name f.30.13
isis spbm 1 add-b-vid 4000
vlan 4000 create spbm-bvlan name "B-VLAN"
isis manual-area add 30.0000
isis sys-name 8600C
isis enable
```

8600G

```
cli prompt 8600G
spbm enable
isis spbm 1 create
isis spbm 1 nick-name f.30.10
isis spbm 1 add-b-vid 4000
vlan 4000 create spbm-bvlan name "B-VLAN"
isis manual-area add 30.0000
isis sys-name 8600G
isis enable
```

8600D

```
cli prompt 8600D
spbm enable
```

```
isis spbm 1 create
isis spbm 1 nick-name f.30.14
isis spbm 1 add-b-vid 4000
vlan 4000 create spbm-bvlan name "B-VLAN"
isis manual-area add 30.0000
isis sys-name 8600D
isis enable
```

IS-IS SPBM Interface Configuration

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

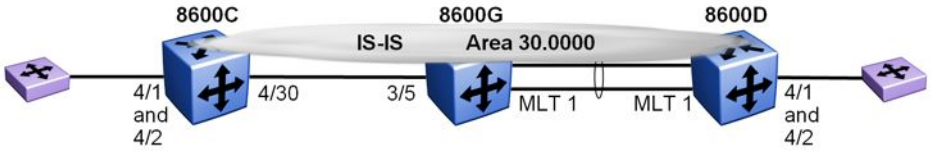

Figure 20: IS-IS SPBM interface

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

8600C

```
ethernet 4/30 isis create
ethernet 4/30 isis spbm 1 state enable
ethernet 4/30 isis enable
```

8600G

```
ethernet 3/5 isis create
ethernet 3/5 isis spbm 1 state enable
ethernet 3/5 isis enable
mlt 1 isis create
mlt 1 isis spbm 1 state enable
mlt 1 isis enable
```

8600D

```
mlt 1 isis create
mlt 1 isis spbm 1 state enable
mlt 1 isis enable
```

Verifying SPBM operations

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

Checking operation — 8600C

```
8600C:5# show isis interface
================================================================================
                                ISIS Interfaces
================================================================================
```

```
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
-------------------------------------------------------------------------------
Port4/30   pt-pt   Level 1   UP        UP         1      1       10

8600C:5# show isis adjacencies
===============================================================================
                               ISIS Adjacencies
===============================================================================
INTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME
-------------------------------------------------------------------------------
Port4/30  1 UP      1d 19:11:30 127       26 000e.6225.a3df    8600G

8600C:5# show isis spbm unicast-fib
===============================================================================
                          SPBM UNICAST FIB ENTRY INFO
===============================================================================
DESTINATION ADDRESS  BVLAN  SYSID               HOST-NAME         OUTGOING-INTERFACE
-------------------------------------------------------------------------------
00:0e:62:25:a3:df    4000   000e.6225.a3df       8600G               4/30
00:14:0d:a0:13:df    4000   0014.0da0.13df       8600D               4/30

8600C:5# show isis spbm unicast-tree 4000
Node:000e.6225.a3df.00 (8600G) -> ROOT
Node:0014.0da0.13df.00 (8600D) -> Node:000e.6225.a3df.00 (8600G) -> ROOT
```

Checking operation — 8600G

```
8600G:5# show isis interface
===============================================================================
                               ISIS Interfaces
===============================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
-------------------------------------------------------------------------------
Port3/5    pt-pt   Level 1   UP        UP         1      1       10
Mlt1       pt-pt   Level 1   UP        UP         1      1       10

8600G:5# show isis adjacencies
===============================================================================
                               ISIS Adjacencies
===============================================================================
INTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME
-------------------------------------------------------------------------------
Port3/5   1 UP      1d 19:19:52 127       26 0015.e89f.e3df    8600C
Mlt1      1 UP         04:57:34 127       20 0014.0da0.13df    8600D

8600G:5# show isis spbm unicast-fib
===============================================================================
                          SPBM UNICAST FIB ENTRY INFO
===============================================================================
DESTINATION ADDRESS  BVLAN  SYSID               HOST-NAME         OUTGOING-INTERFACE
-------------------------------------------------------------------------------
00:14:0d:a0:13:df    4000   0014.0da0.13df       8600D               MLT-1
00:15:e8:9f:e3:df    4000   0015.e89f.e3df       8600C               3/5

8600G:5# show isis spbm unicast-tree 4000
Node:0015.e89f.e3df.00 (8600C) -> ROOT
Node:0014.0da0.13df.00 (8600D) -> ROOT
```

Checking operation — 8600D

```
8600D:5# show isis interface
===============================================================================
                               ISIS Interfaces
===============================================================================
```

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

8600D:5# show isis adjacencies
===============================================================================
                               ISIS Adjacencies
===============================================================================
INTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME
-------------------------------------------------------------------------------
Mlt1      1 UP         05:03:59 127       21 000e.6225.a3df    8600G

8600D:5# show isis spbm unicast-fib
===============================================================================
                          SPBM UNICAST FIB ENTRY INFO
===============================================================================
DESTINATION ADDRESS  BVLAN  SYSID               HOST-NAME         OUTGOING-INTERFACE
-------------------------------------------------------------------------------
00:0e:62:25:a3:df    4000   000e.6225.a3df       8600G               MLT-1
00:15:e8:9f:e3:df    4000   0015.e89f.e3df       8600C               MLT-1

8600D:5# show isis spbm unicast-tree 4000
Node:000e.6225.a3df.00 (8600G) -> ROOT
Node:0015.e89f.e3df.00 (8600C) -> Node:000e.6225.a3df.00 (8600G) -> ROOT
```

CFM configuration

The following sections show the steps required to configure the CFM in this IS-IS SPBM network.

CFM — 8600C

```
cfm md "spbm" create
cfm md "spbm" ma "bvid" create
cfm md "spbm" ma "bvid" mep 1 create state enable
vlan 4000 add-nodal-mep spbm.bvid.1
vlan 4000 add-nodal-mip-level 4
```

CFM — 8600G

```
cfm md "spbm" create
cfm md "spbm" ma "bvid" create
cfm md "spbm" ma "bvid" mep 1 create state enable
vlan 4000 add-nodal-mep spbm.bvid.1
vlan 4000 add-nodal-mip-level 4
```

CFM — 8600D

```
cfm md "spbm" create
cfm md "spbm" ma "bvid" create
cfm md "spbm" ma "bvid" mep 1 create state enable
vlan 4000 add-nodal-mep spbm.bvid.1
vlan 4000 add-nodal-mip-level 4
```

CFM — checking operation on 8600C

The output below shows how to verify that the sample CFM configuration is operating properly.

```
8600C:5# l2ping 4000.8600D
```

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

----00:14:0d:a0:13: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 =  530/530/530.00/ 0.00

8600C:5# l2traceroute 4000.8600D

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

l2traceroute to 8600D  (00:14:0d:a0:13:df),  vlan 4000
0    8600C               (00:15:e8:9f:e3:df)
1    8600G               (00:0e:62:25:a3:df)
2    8600D               (00:14:0d:a0:13:df)
```

SPBM Native IP Shortcuts configuration example

The following figure shows a sample Native IP Shortcuts over SPBM deployment.

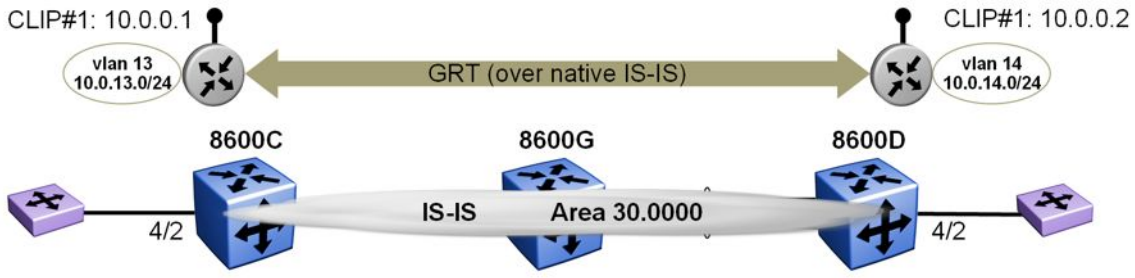

Figure 21: SPBM Native IP Shortcuts

The following sections show the steps required to configure the SPBM Native IP Shortcuts parameters in this example.

Note the following:

- IP IS-IS redistribution needs to be configured to inject GRT routes into IS-IS.
- In the displayed configuration, only direct routes are injected (the same configuration is possible for RIP, OSPF, BGP, and static routes).
- No IP address needs to be configured on 8600G.

The following sections show the steps required to configure the SPBM Native IP Shortcuts parameters in this example.

8600C

```
ip circuitless-ip-int 1 create 10.0.0.1/32
isis ip source-address 10.0.0.1
isis spbm 1 ip enable

vlan 13 create byport 1
vlan 13 ports add 4/2
vlan 13 ip create 10.0.13.1/24
```

```
ip isis redistribute direct create
ip isis redistribute direct enable
ip isis redistribute direct apply
```

8600D

```
ip circuitless-ip-int 1 create 10.0.0.2/32
isis ip source-address 10.0.0.2
isis spbm 1 ip enable

vlan 14 create byport 1
vlan 14 ports add 4/2
vlan 14 ip create 10.0.14.1/24

ip isis redistribute direct create
ip isis redistribute direct enable
ip isis redistribute direct apply
```

Verifying operation — 8600C

```
8600C:5# show isis spbm ip-unicast-fib
================================================================================
                          SPBM IP-UNICAST FIB ENTRY INFO
================================================================================
                                                 Outgoing     SPBM         Prefix
VRF     ISID    Destination    NH BEB    VLAN    Interface    Cost     Cost
--------------------------------------------------------------------------------
GRT     -       10.0.0.2/32    8600D     4000     4/30        20        1
GRT     -       10.0.14.0/24   8600D     4000     4/30        20        1

8600C:5# show ip route info
================================================================================
                            IP Route - GlobalRouter
================================================================================
                                                   NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.0.0.1       255.255.255.255 10.0.0.1             -    1    0     LOC  0   DB   0
10.0.0.2       255.255.255.255 00:14:0d:a0:13:df    Glob~ 20   4000  ISIS 0   IBS  7
10.0.13.0      255.255.255.0   10.0.13.1            -    1    13    LOC  0   DB   0
10.0.14.0      255.255.255.0   00:14:0d:a0:13:df    Glob~ 20   4000  ISIS 0   IBS  7
 IBS  7
```

Verifying operation — 8600D

```
8600D:5# show isis spbm ip-unicast-fib
================================================================================
                          SPBM IP-UNICAST FIB ENTRY INFO
================================================================================
                                                 Outgoing     SPBM         Prefix
VRF     ISID    Destination    NH BEB    VLAN    Interface    Cost     Cost
--------------------------------------------------------------------------------
GRT     -       10.0.0.1/32    8600C     4000     4/20        20        1
GRT     -       10.0.13.0/24   8600C     4000     4/20        20        1

8600D:5# show ip route inf
================================================================================
                            IP Route - GlobalRouter
================================================================================
                                                   NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.0.0.1       255.255.255.255 00:15:e8:9f:e3:df    Glob~ 20   4000  ISIS 0   IBS  7
10.0.0.2       255.255.255.255 10.0.0.2             -    1    0     LOC  0   DB   0
```

```
10.0.13.0       255.255.255.0   00:15:e8:9f:e3:df    Glob~ 20   4000  ISIS 0   IBS  7
10.0.14.0       255.255.255.0   10.0.14.1             -    1    14    LOC  0   DB   0
 DB    0
```

L2 VSN configuration example

The following figure shows a sample L2 VSN deployment.

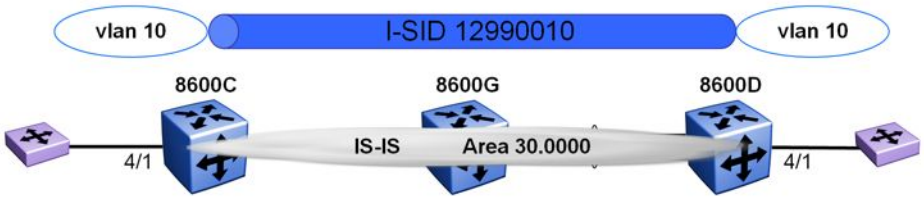

Figure 22: L2 VSN

The following sections show the steps required to configure the L2 VSN parameters in this example.

8600C

```
vlan 10 create byport 1
vlan 10 ports add 4/1
vlan 10 i-sid 12990010
```

8600D

```
vlan 10 create byport 1
vlan 10 ports add 4/1
vlan 10 i-sid 12990010
```

Verifying L2 VSN operation

The following sections show how to verify the L2 VSN operation in this example.

8600C

```
8600C:5# show isis spbm i-sid all
================================================================================
                                SPBM ISID INFO
================================================================================
ISID   SOURCE NAME    VLAN   SYSID               TYPE
--------------------------------------------------------------------------------
12990010 f.30.14        4000   0014.0da0.13df       discover
12990010 f.30.13        4000   0015.e89f.e3df       config

8600C:5# show isis spbm multicast-fib
================================================================================
                         SPBM MULTICAST FIB ENTRY INFO
================================================================================
```

```
MCAST DA           ISID   BVLAN  SYSID                HOST-NAME         OUTGOING-
INTERFACES
-------------------------------------------------------------------------------
f3:30:14:c6:36:3a  12990010 4000   0014.0da0.13df       8600D             4/1
f3:30:13:c6:36:3a  12990010 4000   0015.e89f.e3df       8600C             4/30,4/1
```

8600D

```
8600D:5# show isis spbm i-sid all
===============================================================================
                              SPBM ISID INFO
===============================================================================
ISID   SOURCE NAME    VLAN   SYSID               TYPE
-------------------------------------------------------------------------------
12990010 f.30.14        4000   0014.0da0.13df       config
12990010 f.30.13        4000   0015.e89f.e3df       discover

8600D:5# show isis spbm multicast-fib
===============================================================================
                         SPBM MULTICAST FIB ENTRY INFO
===============================================================================
MCAST DA           ISID   BVLAN  SYSID                HOST-NAME         OUTGOING-
INTERFACES
-------------------------------------------------------------------------------
f3:30:14:c6:36:3a  12990010 4000   0014.0da0.13df       8600D             MLT-1,4/1
f3:30:13:c6:36:3a  12990010 4000   0015.e89f.e3df       8600C             4/1
```

8600C — verifying with CFM

```
8600C:5# l2tracetree 4000.12990010

 Please wait for l2tracetree to complete or press any key to abort

l2tracetree to f3:30:13:c6:36:3a, vlan 4000 i-sid 12990010 nickname f.30.13 hops 64
1   8600C             00:15:e8:9f:e3:df -> 8600G             00:0e:62:25:a3:df
2   8600G             00:0e:62:25:a3:df -> 8600D             00:14:0d:a0:13:df
```

8600D — verifying with CFM

```
8600D:5# l2tracetree 4000.12990010

 Please wait for l2tracetree to complete or press any key to abort

l2tracetree to f3:30:14:c6:36:3a, vlan 4000 i-sid 12990010 nickname f.30.14 hops 64
1   8600D             00:14:0d:a0:13:df -> 8600G             00:0e:62:25:a3:df
2   8600G             00:0e:62:25:a3:df -> 8600C             00:15:e8:9f:e3:df
```

8600C — verifying FDB

```
8600C:5# show vlan inf fdb-entry 10
===============================================================================
                                  Vlan Fdb
===============================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
-------------------------------------------------------------------------------
10   learned    00:00:00:00:00:01  Port-4/1         false   1      false
10   learned    00:00:00:00:00:02  I-SID-12990010   false   1      false

8600C:5# show vlan inf remote-mac-table 10
===============================================================================
                            Vlan Remote Mac Table
===============================================================================
VLAN STATUS  MAC-ADDRESS        DEST-MAC           DEST-SYSID     DEST-SYSNAME
PORTS    SMLTREMOTE
```

```
--------------------------------------------------------------------------------
10   learned 00:00:00:00:00:02  00:14:0d:a0:13:df  0014.0da0.13df 8600D
4/30     false
--------------------------------------------------------------------------------
 Total number of VLAN Remote MAC entries 1
 * For each remote-mac entry, the first line is for SPBM primary
   b-vid, the second line is for SPBM secondary b-vid.
--------------------------------------------------------------------------------
```

8600D — verifying FDB

```
8600D:5# show vlan info fdb-entry 10
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
10   learned    00:00:00:00:00:01  I-SID-12990010   false   1      false
10   learned    00:00:00:00:00:02  Port-4/1         false   1      false

8600D:5# show vlan info remote-mac-table 10
================================================================================
                             Vlan Remote Mac Table
================================================================================
VLAN STATUS  MAC-ADDRESS        DEST-MAC           DEST-SYSID     DEST-SYSNAME
PORTS    SMLTREMOTE
--------------------------------------------------------------------------------
10   learned 00:00:00:00:00:01  00:15:e8:9f:e3:df  0015.e89f.e3df 8600C
MLT-1    false
--------------------------------------------------------------------------------
 Total number of VLAN Remote MAC entries 1
 * For each remote-mac entry, the first line is for SPBM primary
   b-vid, the second line is for SPBM secondary b-vid.
--------------------------------------------------------------------------------
```

L2 VSN example with VLAN ID translation

The following figure shows a sample L2 VSN deployment where the C- VLAN IDs are different at each end.

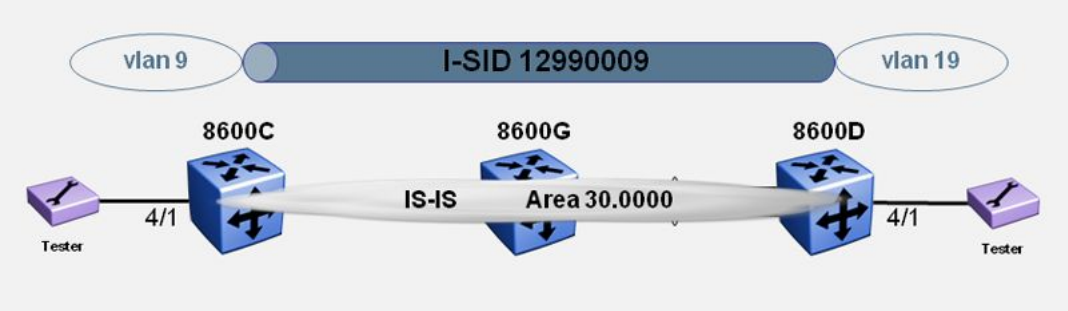

Figure 23: L2 VSN with different VLAN IDs

The following sections show the steps required to configure the L2 VSN parameters in this example.

8600C

```
vlan 9 create byport 1
vlan 9 ports add 4/1
vlan 9 i-sid 12990009
```

8600D

```
vlan 19 create byport 1
vlan 19 ports add 4/1
vlan 19 i-sid 12990009
```

L3 VSN configuration example

The following figure shows a sample L3 VSN deployment.

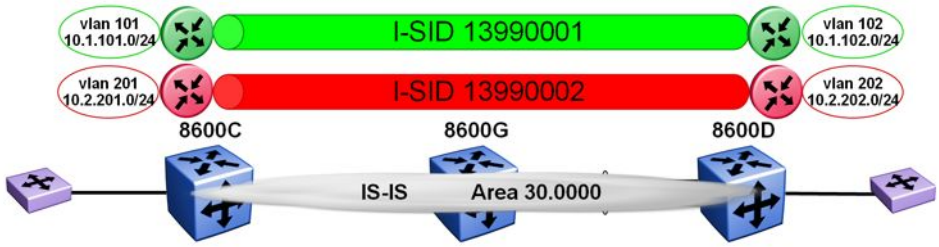

Figure 24: L3 VSN

The following sections show the steps required to configure the L3 VSN parameters in this example.

Note that IP ISIS redistribution needs to be configured to inject the VRF routes into IS-IS.

VRF green configuration

The following figure shows the green VRF in this L3 VSN example.

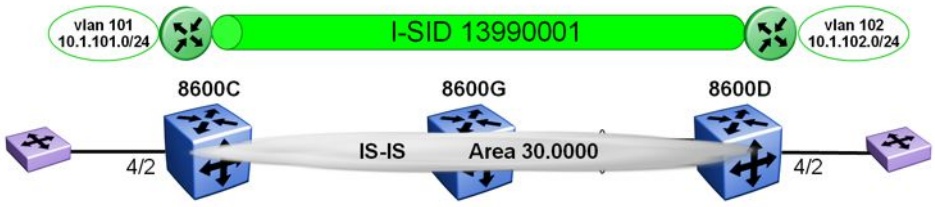

Figure 25: L3 VSN—VRF green

The following sections show the steps required to configure the green VRF parameters in this example.

VRF green – 8600C

```
ip vrf green create id 1
vlan 101 create byport 1
vlan 101 ports add 4/2
vlan 101 vrf green
vlan 101 ip create 10.1.101.1/24
ip vrf green ipvpn create
ip vrf green ipvpn i-sid 13990001
ip vrf green ipvpn enable
ip vrf green isis redistribute direct create
ip vrf green isis redistribute direct enable
ip vrf green isis redistribute direct apply
```

VRF green – 8600D

```
ip vrf green create id 1
vlan 102 create byport 1
vlan 102 ports add 4/2
vlan 102 vrf green
vlan 102 ip create 10.1.102.1/24
ip vrf green ipvpn create
ip vrf green ipvpn i-sid 13990001
ip vrf green ipvpn enable
ip vrf green isis redistribute direct create
ip vrf green isis redistribute direct enable
ip vrf green isis redistribute direct apply
```

VRF red configuration

The following figure shows the red VRF in this L3 VSN example.

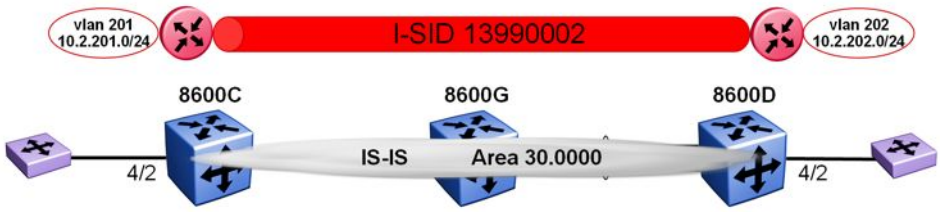

Figure 26: L3 VSN—VRF red

The following sections show the steps required to configure the red VRF parameters in this example.

VRF red – 8600C

```
ip vrf red create id 2
vlan 201 create byport 1
vlan 201 ports add 4/2
vlan 201 vrf red
vlan 201 ip create 10.2.201.1/24
ip vrf red ipvpn create
ip vrf red ipvpn i-sid 13990002
ip vrf red ipvpn enable
ip vrf red isis redistribute direct create
ip vrf red isis redistribute direct enable
ip vrf red isis redistribute direct apply
y
```

VRF red – 8600D

```
ip vrf red create id 2
vlan 202 create byport 1
vlan 202 ports add 4/2
vlan 202 vrf red
vlan 202 ip create 10.2.202.1/24
ip vrf red ipvpn create
ip vrf red ipvpn i-sid 13990002
ip vrf red ipvpn enable
ip vrf red isis redistribute direct create
ip vrf red isis redistribute direct enable
ip vrf red isis redistribute direct apply
```

Verifying L3 VSN operation

The following sections show the steps required to verify the L3 VSN configuration in this example.

8600C

```
8600C:5# show isis spbm ip-unicast-fib
================================================================================
                          SPBM IP-UNICAST FIB ENTRY INFO
================================================================================
                                                        Outgoing     SPBM        Prefix
VRF     ISID      Destination      NH BEB    VLAN     Interface     Cost      Cost
--------------------------------------------------------------------------------
GRT     -         10.0.0.2/32      8600D     4000      4/30         20         1
GRT     -         10.0.14.0/24     8600D     4000      4/30         20         1
```

```
green   13990001  10.1.102.0/24    8600D     4000      4/30         20         1
red     13990002  10.2.202.0/24    8600D     4000      4/30         20         1
```

8600D

```
8600D:5/config# show isis spbm ip-unicast-fib
================================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO
================================================================================
                                                      Outgoing     SPBM        Prefix
VRF     ISID      Destination      NH BEB    VLAN     Interface    Cost      Cost
--------------------------------------------------------------------------------
GRT     -         10.0.0.1/32      8600C     4000      4/20         20         1
GRT     -         10.0.13.0/24     8600C     4000      4/20         20         1
green   13990001  10.1.101.0/24    8600C     4000      4/20         20         1
red     13990002  10.2.201.0/24    8600C     4000      4/20         20         1
```

VRF green—8600C

```
8600C:5# show ip route info vrf green
================================================================================
                              IP Route - VRF green
================================================================================
                                                    NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.1.101.0      255.255.255.0   10.1.101.1          -    1    101   LOC  0   DB   0
10.1.102.0      255.255.255.0   00:14:0d:a0:13:df   Glob~ 20   4000  ISIS 0   IBSV 7
```

VRF green—8600D

```
8600D:5# show ip route info vrf green
================================================================================
                              IP Route - VRF green
================================================================================
                                                    NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.1.101.0      255.255.255.0   00:15:e8:9f:e3:df   Glob~ 20   4000  ISIS 0   IBSV 7
10.1.102.0      255.255.255.0   10.1.102.1          -    1    102   LOC  0   DB   0
```

VRF red—8600C

```
8600C:5# show ip route info vrf red
================================================================================
                               IP Route - VRF red
================================================================================
                                                    NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.2.201.0      255.255.255.0   10.2.201.1          -    1    201   LOC  0   DB   0
10.2.202.0      255.255.255.0   00:14:0d:a0:13:df   Glob~ 20   4000  ISIS 0   IBSV 7
```

VRF red—8600D

```
8600D:5# show ip route info vrf red
================================================================================
                               IP Route - VRF red
================================================================================
                                                    NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
```

```
10.2.201.0      255.255.255.0   00:15:e8:9f:e3:df    Glob~ 20   4000  ISIS 0   IBSV 7
10.2.202.0      255.255.255.0   10.2.202.1            -    1    202   LOC  0   DB   0
```

BGP over SPBM network configuration example

The following figure shows a sample BGP deployment over SPBM.

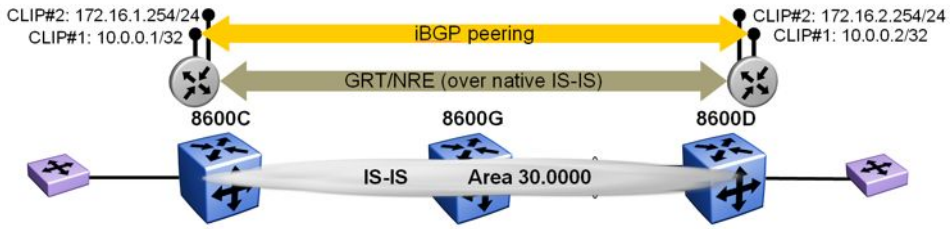

Figure 27: Global BGP configuration over SPBM network

The following sections show the steps required to configure the BGP parameters in this example.

Note that the BGP Router ID is set by configuring the OSPF router ID (**`ip ospf router-id`** command).

Also, with BGP, the L3 VPN routes will not be seen by the 8600G in the IS-IS LSDB.

8600C

```
ip circuitless-ip-int 2 create 172.16.1.254/24
ip circuitless-ip-int 2 ipvpn-lite-capability enable
ip ospf router-id 10.0.0.1
ip bgp auto-summary disable
ip bgp synchronization disable
ip bgp local-as 65000
ip bgp aggregation disable
ip bgp enable
ip bgp quick-start enable
ip bgp neighbor "10.0.0.2" create
ip bgp neighbor 10.0.0.2 remote-as 65000
ip bgp neighbor 10.0.0.2 update-source-interface 10.0.0.1 add
ip bgp neighbor 10.0.0.2 address-family vpnv4 enable
ip bgp neighbor 10.0.0.2 ipvpn-lite-capability enable
ip bgp neighbor 10.0.0.2 admin-state enable
```

8600D

```
ip circuitless-ip-int 2 create 172.16.2.254/24
ip circuitless-ip-int 2 ipvpn-lite-capability enable
ip ospf router-id 10.0.0.2
ip bgp auto-summary disable
ip bgp synchronization disable
ip bgp local-as 65000
ip bgp aggregation disable
ip bgp enable
ip bgp quick-start enable
ip bgp neighbor "10.0.0.1" create
ip bgp neighbor 10.0.0.1 remote-as 65000
```

```
ip bgp neighbor 10.0.0.1 update-source-interface 10.0.0.2 add
ip bgp neighbor 10.0.0.1 address-family vpnv4 enable
ip bgp neighbor 10.0.0.1 ipvpn-lite-capability enable
ip bgp neighbor 10.0.0.1 admin-state enable
```

Verifying BGP over SPBM operation

The following sections show how to verify BGP over SPBM operation in this example.

8600C

```
8600C:5# show ip bgp summary
================================================================================
                            BGP Summary - GlobalRouter
================================================================================
                                     BGP version - 4
                                        local-as - 65000
                                      Identifier - 10.0.0.1
                                  Decision state - Idle
                    The total number of routes is   0
BGP NEIGHBOR INFO :
   NEIGHBOR       RMTAS     STATE    HLDTM  KPALV  HLDCFG KPCFG   WGHT  CONRTY ADVINT
--------------------------------------------------------------------------------
10.0.0.2          65000 Established  180    60     180    60     100    120    5
Total bgp neighbors: 1
```

8600D

```
8600D:5# show ip bgp summary
================================================================================
                            BGP Summary - GlobalRouter
================================================================================
                                     BGP version - 4
                                        local-as - 65000
                                      Identifier - 10.0.0.2
                                  Decision state - Idle
                    The total number of routes is   0
BGP NEIGHBOR INFO :
   NEIGHBOR       RMTAS     STATE    HLDTM  KPALV  HLDCFG KPCFG   WGHT  CONRTY ADVINT
--------------------------------------------------------------------------------
10.0.0.1          65000 Established  180    60     180    60     100    120    5
Total bgp neighbors: 1
```

IP VPN-Lite L3 VSN over IS-IS configuration example

The following figure shows a sample IP VPN-Lite L3 VSN deployment.

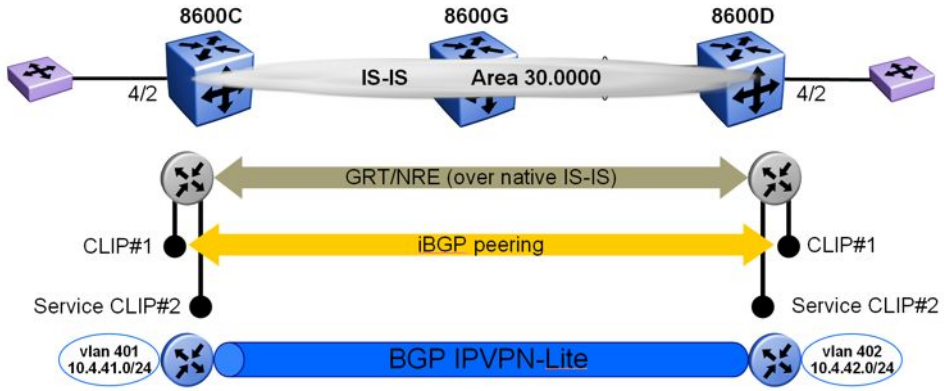

Figure 28: IP VPN-Lite L3 VSN over IS-IS

The following sections show the steps required to configure the IP VPN—Lite L3 VSN parameters in this example.

8600C

```
ip vrf blue create id 4
vlan 401 create byport 1
vlan 401 ports add 4/2
vlan 401 vrf blue
vlan 401 ip create 10.4.41.1/24
ip vrf blue ipvpn create
ip vrf blue ipvpn rd 172.16.1.4:4
ip vrf blue ipvpn rt add import 65000:60004
ip vrf blue ipvpn rt add export 65000:60004
ip vrf blue ipvpn enable
ip vrf blue bgp auto-summary disable
```

8600D

```
ip vrf blue create id 4
vlan 402 create byport 1
vlan 402 ports add 4/2
vlan 402 vrf blue
vlan 402 ip create 10.4.42.1/24
ip vrf blue ipvpn create
ip vrf blue ipvpn rd 172.16.2.4:4
ip vrf blue ipvpn rt add import 65000:60004
ip vrf blue ipvpn rt add export 65000:60004
ip vrf blue ipvpn enable
ip vrf blue bgp auto-summary disable
```

Verifying IP VPN-Lite L3 VSN operation

The following sections show how to verify IP VPN-Lite L3 VSN operation in this example.

8600C

```
8600C:5# show ip route info vrf blue
================================================================================
                              IP Route - VRF blue
================================================================================
```

```
                                                    NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
-------------------------------------------------------------------------------------
10.4.41.0      255.255.255.0   10.4.41.1            -    1    401   LOC  0   DB   0
10.4.42.0      255.255.255.0   00:14:0d:a0:13:df   Glob~ 0    4000  BGP  0   IBSV 175
```

8600D

```
8600D:5# show ip route info vrf blue
=============================================================================
                            IP Route - VRF blue
=============================================================================
                                                    NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
-------------------------------------------------------------------------------------
10.4.41.0      255.255.255.0   00:15:e8:9f:e3:df   Glob~ 0    4000  BGP  0   IBSV 175
10.4.42.0      255.255.255.0   10.4.42.1            -    1    402   LOC  0   DB   0
```

Inter-ISID Routing configuration example

The following figure shows a sample Inter-ISID deployment.

Figure 29: Inter-ISID routing configuration

The following sections show the steps required to configure the Inter-ISID parameters in this example.

Note that the IP interfaces are configured where the routing instance exists, namely, on 8600G.

8600C

```
vlan 11 create byport 1
vlan 11 ports add 4/2
vlan 11 i-sid 12990011
```

8600G

```
ip vrf isid create id 100
vlan 11 create byport 1
vlan 11 vrf isid
vlan 11 i-sid 12990011
vlan 11 ip create 10.100.11.1/24
vlan 12 create byport 1
vlan 12 vrf isid
vlan 12 i-sid 12990012
vlan 12 ip create 10.100.12.1/24
```

8600D

```
vlan 12 create byport 1
vlan 12 ports add 4/2
vlan 12 i-sid 12990012
```

Verifying Inter-ISID Routing operation

The following sections show how to verify Inter-ISID Routing operation in this example. .

8600G

```
8600G:5# show ip route info vrf isid
================================================================================
                              IP Route - VRF isid
================================================================================
                                                      NH          INTER
DST             MASK            NEXT              VRF   COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.100.11.0    255.255.255.0   10.100.11.1          -    1    11    LOC  0   DB   0
10.100.12.0    255.255.255.0   10.100.12.1          -    1    12    LOC  0   DB   0

8600G:5# show ip arp info vrf isid
================================================================================
                              IP Arp - VRF isid
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT   TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
10.100.11.1     00:0e:62:25:a2:00  11       -        LOCAL   2160
10.100.11.255   ff:ff:ff:ff:ff:ff  11       -        LOCAL   2160
10.100.12.1     00:0e:62:25:a2:01  12       -        LOCAL   2160
10.100.12.255   ff:ff:ff:ff:ff:ff  12       -        LOCAL   2160
10.100.11.10    00:00:00:00:01:02  11    I-SID-12990011  DYNAMIC 2159
10.100.12.10    00:00:00:00:02:02  12    I-SID-12990012  DYNAMIC 2160
```

8600G

```
8600G:5# show vlan info fdb-entry 11
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
11   learned    00:00:00:00:01:02  I-SID-12990011   false   1      false
11   self       00:0e:62:25:a2:00  Port-cpp         false   1      false

8600G:5# show vlan info fdb-entry 12
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
12   learned    00:00:00:00:02:02  I-SID-12990012   false   1      false
12   self       00:0e:62:25:a2:01  Port-cpp         false   1      false
 1      false
```

8600C

```
8600C:5# show vlan info fdb-entry 11
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
11   learned    00:00:00:00:01:02  Port-4/2         false   1      false
11   learned    00:0e:62:25:a2:00  I-SID-12990011   false   1      false
```

8600D

```
8600D:5# show vlan info fdb-entry 12
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
12   learned    00:00:00:00:02:02  Port-4/2         false   1      false
12   learned    00:0e:62:25:a2:01  I-SID-12990012   false   1      false
```

Chapter 6: Configuring SPBM using the CLI

This chapter describes how to configure SPBM using the CLI.

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.

Procedure steps

1. Enable SPBM globally:

   ```
   config spbm {enable|disable}
   ```

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

   ```
   config isis spbm <inst> {create|delete}
   ```

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

   ```
   config vlan <vlan-id> create spbm-bvlan [name <value>]
   [ color <value>]
   ```

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

   ```
   config isis spbm <inst> add-b-vid {<vlan-id [-vlan-id][,...]}
   [primary <vid>]
   ```

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

   ```
   config isis spbm <inst> nick-name <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 (using the `config isis system-id <system-id>` command) . This helps to recognize source and destination addresses for troubleshooting purposes.

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

   ```
   config isis manual-area {add <area> | delete <area>}
   ```

7. On the ports or MLTs that are going to link to the SPBM network, enable tagging:

```
config {ethernet <slot/port> | mlt <id>} perform-tagging
enable
```

8. Create an IS-IS interface on the SPBM ports or MLTs:

```
config {ethernet <slot/port> | mlt <id>} isis {create |
delete}
```

9. On the IS-IS interfaces, enable the SPBM instance:

```
config {ethernet <slot/port> | mlt <id>} isis spbm <inst>
state {enable|disable}
```

10. Enable the IS-IS interface on the ports or MLTs:

```
config {ethernet <slot/port> | mlt <id>} isis {enable |
disable}
```

11. Enable IS-IS globally:

```
config isis {enable | disable}
```

12. To display the SPBM configurations, enter:

```
show isis spbm info
```

13. To display the global IS-IS configuration, enter:

```
show isis info
```

14. To display the interface IS-IS configuration, enter:

```
show isis interface
```

Important:

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

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

Variable definitions

Use the data in the following table to configure the SPBM and IS-IS parameters.

Variable	Value
config spbm {enable|disable}	Enables or disables SPBM globally on the switch.

Variable	Value
config isis spbm <*inst*> {create| delete}	Creates or deletes the specified SPBM instance. In this release, only one SPBM instance is supported.
config vlan <*vlan-id*> create spbm-bvlan [name <*value*>] [color <*value*>]	Creates an SPBM Backbone VLAN (B-VLAN). You can optionally specify a name and color for the SPBM B-VLAN. name <value> specifies the name of VLAN {string length 0..64}. color <value> specifies the color of the VLAN {0..32}.
config isis spbm <*inst*> add-b-vid {<*vlan-id [-vlan-id][,...]*} [primary <*vid*>]	Specifies the SPBM B-VLANs to add to the SPBM instance.
config isis spbm <*inst*> nick-name <*nick-name*>	Specifies a nickname for the SPBM instance globally. Value is 2.5 bytes in the format <x.xx.xx>.
config isis manual-area {add <*area*> | delete <*area*>}	Adds or deletes the specified IS-IS manual area. In this release, only one manual area is supported. <area> is 1-13 bytes in the format <xx.xxxx.xxxx...xxxx>.
config isis {enable | disable}	Enables or disables IS-IS globally on the switch.
config {ethernet <*slot/port*> | mlt <*id*>} isis {create | delete}	Creates or deletes an IS-IS circuit/interface on the specified port or MLT.
config {ethernet <*slot/port*> | mlt <*id*>} isis spbm <*inst*> state {enable|disable}	Enables or disables the SPBM instance on the specified port or MLT.
config {ethernet <*slot/port*> | mlt <*id*>} isis {enable | disable}	Enables or disables the IS-IS circuit/interface on the specified port or MLT.

Configuring SMLT parameters for SPBM

Use the following procedure to configure the required SMLT parameters to allow SPBM to interoperate with SMLT on the switch.

Note:

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

Procedure steps

1. Disable IS-IS on the switch:

```
config isis disable
```

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

   ```
   config isis spbm <inst> smlt peer-system-id <system-id>
   ```

3. Configure the virtual-bmac, which is shared and advertised by both peers.

   ```
   config isis spbm <inst> smlt virtual-bmac <virtual-mac>
   ```

4. Enable IS-IS on the switch:

   ```
   config isis enable
   ```

5. To display the SPBM SMLT configuration, enter:

   ```
   config isis spbm <inst> smlt info
   ```

Variable definitions

Use the data in the following table to configure the SMLT parameters for SPBM.

Variable	Value
smlt-peer-system-id<*system-id*>	Specifies the IST peer system ID address in the xxxx.xxxx.xxxx format.
smlt-virtual-bmac <*virtual-mac*>	Specifies a virtual MAC address that can be used by both peers.
smlt info	Displays SPBM SMLT info.

Configuring SPBM L2 VSN

After you have configured the SPBM infrastructure, you can enable SPBM L2 VSN using the following procedure.

Procedure steps

Map a customer VLAN (CVLAN) to an instance identifier (ISID):

```
config vlan <vlan-id> i-sid <isid>
```

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.

Variable definitions

Variable	Value
config vlan <*vlan-id*> i-sid <*isid*>	Specifies the customer VLAN (CVLAN) to associate with the ISID.

Configuring SPBM Native IP shortcuts

After you have configured the SPBM infrastructure, you can enable SPBM Native IP shortcuts to advertise IP routes across the SPBM network using the following procedure.

Procedure steps

1. Create a CLIP interface to use as the source address for SPBM Native IP shortcuts:

   ```
   config ip circuitless-ip-int <id> create <ipaddr/mask>
   ```

2. Specify the CLIP interface as the source address for SPBM Native IP shortcuts:

   ```
   config isis ip source-address <A.B.C.D>
   ```

3. Configure SPBM Native IP shortcuts:

   ```
   config isis spbm <inst> ip {enable | disable}
   ```

4. To view the status of SPBM Native IP shortcuts on the switch, enter:

   ```
   config isis spbm <inst> info
   ```

5. To identify routes on the local switch to be announced into the SPBM network, enter:

   ```
   config ip isis redistribute {direct | bgp | ospf | rip |
   static} create
   ```

   ```
   config ip isis redistribute {direct | bgp | ospf | rip |
   static} enable
   ```

   ```
   config ip isis redistribute {direct | bgp | ospf | rip |
   static} apply
   ```

6. To display the redistribution configuration, enter:

   ```
   show ip isis redistribute info
   ```

Variable definitions

Use the data in the following table to configure the SPBM Native IP shortcuts parameters.

Variable	Value
circuitless-ip-int <*id*> create <*ipaddr/mask*>	Specifies the IP address and ID (1–256) for the circuitless IP interface.

Variable	Value
ip source-address <*A.B.C.D*>	Specifies the CLIP interface to use as the source address for SBPM Native IP shortcuts.
spbm <*inst*>	Specifies the SPBM instance ID.
ip {enable |disable}	Enables or disables SPBM IP shortcut state.
info	Displays SPBM configuration information.

Configuring SPBM L3 VSN

After you have configured the SPBM infrastructure, you can enable SPBM L3 VSN to advertise IP routes across the SPBM network from one VRF to another using the following procedure.

Prerequisites

You must configure a VRF and IP VPN instance on the switch. For more information, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP VPN* (NN46205–520).

Procedure steps

1. To enable SPBM L3 VSN, enter:

   ```
   config ip vrf <vrf-name> ipvpn i-sid <isid>
   ```

2. To identify routes on the local switch to be announced into the SPBM network, enter:

   ```
   config ip vrf <vrf-name> isis redistribute {direct | bgp |
   ospf | rip | static} create

   config ip vrf <vrf-name> isis redistribute {direct | bgp |
   ospf | rip | static} enable

   config ip vrf <vrf-name> isis redistribute {direct | bgp |
   ospf | rip | static} apply
   ```

3. To display the redistribution configuration, enter:

   ```
   show ip isis redistribute info [vrf <vrf-name>] [vrfids
   <vrfids>]
   ```

Variable definitions

Use the data in the following table to configure the SPBM L3 VSN parameters.

Variable	Value
<*vrf-name*>	Specifies the VRF name.
<*isid*>	Specifies the ISID to associate with the VRF.

Configuring optional SPBM parameters

Use the following procedure to configure optional SPBM parameters.

Procedure steps

1. Configure the SPBM ethertype:

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

2. Configure the optional LSDB trap global parameter. To configure this parameter, you must globally disable IS-IS on the switch:

   ```
   config isis disable

   config isis spbm <inst> lsdb-trap {enable | disable}

   config isis enable
   ```

3. To view the configured global SPBM parameters, enter

   ```
   config isis spbm <inst> info
   ```

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

   ```
   config {ethernet <slot/port> | mlt <mltid>} isis disable

   config {ethernet <slot/port> | mlt <mltid>} isis spbm <inst>
   [interface-type <if-type>] [l1-metric <cost>]

   config {ethernet <slot/port> | mlt <mltid>} isis enable
   ```

5. To view the configured interface SPBM parameters, enter

   ```
   config {ethernet <slot/port> | mlt <mltid>} isis spbm <inst>
   info
   ```

Variable definitions

Use the data in the following table to configure the optional SPBM parameters.

Variable	Value
config spbm ethertype {0x8100 | 0x88a8}	Configures the SPBM ethertype. The default value is 0x8100.
config isis disable	Globally disables IS-IS on the switch.
spbm *<inst>*	Specifies the SPBM instance ID.
lsdb-trap {enable | disable}	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable.
config isis enable	Globally enables IS-IS on the switch.

Variable	Value
info	Displays SPBM configuration information.
isis disable	Disables IS-IS on the interface.
interface-type *<if-type>*	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT. In this release, only the pt-pt interface type is supported.
l1-metric <*cost*>	Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10. Range is 1–16777215.
isis enable	Enables IS-IS on the interface.

Configuring optional IS-IS global parameters

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

Procedure steps

Configure optional IS-IS global parameters:

```
config isis
[csnp-interval <csnp-interval>]
[info]
[is-type {l1 | l12}]
[ip source—address <value>]
[max-lsp-gen-interval <max-lsp-interval>]
[metric {narrow | wide}]
[min-lsp-gen-interval <min-lsp-interval>]
[overload {true | false}]
[overload-on-startup <overload-on-startup>]
[psnp-interval <psnp-interval>]
[retransmit-lspint <retransmit-lspint>]
[spf-delay <spf-delay-time>]
[system-id <system id>]
```

```
[sys-name <sys-name>]
```

Important:

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

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

Variable definitions

Use the data in the following table to configure the optional IS-IS global parameters.

Variable	Value
csnp-interval <*csnp-interval*>	Specifies the CSNP interval in seconds. This is a system level parameter that applies for L1 CSNP generation on all interfaces. Default value is 10
info	Displays IS-IS configuration information.
is-type {l1 | l12}	Sets the router type globally: • l1: Level-1 router type • l12: Not valid in the current release The default value is l1.
ip source-address <*value*>	Indicates the IP source address used for SPBM Native IP shortcuts.
max-lsp-gen-interval <*max-lsp-interval*>	Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate system. The value must be greater than any value configured for min-lsp-gen-interval. Default value is 900 seconds. Range is 0–900.
metric {narrow | wide}	Specifies the IS-IS metric type. Only wide is supported in this release.
min-lsp-gen-interval <*min-lsp-interval*>	Specifies the minimum time between successive generation of LSPs with the same LSPID. This a system level parameter that applies to L1 LSP generation.

Variable	Value
	Default value is 30 seconds.
overload {true | false}	This sets or clears the overload condition. Default value is false.
overload-on-startup <*overload-on-startup*>	This sets the IS-IS overload-on-startup value in seconds. The overload-on-startup value is used as a timer to control when to send out LSPs with the overload bit cleared after IS-IS startup. The default value is 20.
psnp-interval <*psnp-interval*>	Specifies the PSNP interval in seconds. This is a system level parameter that applies for L1 PSNP generation on all interfaces. Default value is 2.
retransmit-lspint <*retransmit-lspint*>	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. Default value is 5 seconds.
spf-delay <*spf-delay-time*>	Specifies the SPF delay in milliseconds. This value is used to pace successive SPF runs. The timer prevents two SPF runs from being scheduled very closely. The default value is 100 milliseconds.
sys-name <*sys-name*>	Specifies a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system name comes from the host name configured at the system level.
system-id <*system id*>	Specifies the IS-IS system ID for the switch.

Configuring optional IS-IS interface parameters

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

Procedure steps

Configure optional IS-IS interface parameters:

```
config {ethernet <slot/port> | mlt <mltid>} isis
[hello-auth type [key <value>] [key-id <value>]]
[info]
```

```
[l1-hello-interval <seconds>]
[l1-hello-multiplier <seconds>]
[l1-dr-priority <integer>]
```

Variable definitions

Use the data in the following table to configure the optional IS-IS interface parameters.

Variable	Value
config {ethernet <*slot/port*> | mlt <*mltid*>} isis	Specifies the port or MLT for which to configure IS-IS parameters.
hello-auth type [key <*value*>] [key-id <*value*>]	Specifies the authentication type used for IS-IS hello packets on the interface. <type> can be one of the following: • none • simple: if selected, you can also specify a key value • hmac-md5: if selected, you can also specify a key value and key-id
info	Displays IS-IS interface parameters.
l1-dr-priority <*integer*>	Configures the level 1 IS-IS designated router priority to the specified value. Default value is 64.
l1-hello-interval <*seconds*>	Configures the level 1 hello interval. Default value is 9 seconds.
l1-hello-multiplier <*seconds*>	Configures the level 1 hello multiplier. Default value is 3 seconds.

Chapter 7: Displaying SPBM and IS-IS using the CLI

This chapter describes how to display SPBM and IS-IS parameters using the CLI.

Displaying global SPBM parameters

Use the following procedure to display global SPBM parameters.

Procedure steps

1. To display the SPBM configuration, enter:

   ```
   show isis spbm info
   ```

2. To display all SPBM information, enter:

   ```
   show isis spbm show-all [file <value>]
   ```

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

   ```
   show vlan info all
   ```

Variable definitions

Use the data in the following table to use the global SPBM show commands.

Variable	Value
[file <*value*>]	Displays the SPBM configuration from the previously saved configuration file specified.

Job aid

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

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

Displaying CVLAN ISID information

Use the following procedure to display CVLAN ISID information.

Procedure steps

1. To display the CVLAN to ISID associations:

   ```
   show vlan info i-sid
   ```

2. You can also use the following command to view the ISID of the CVLAN for entries learned from the SPBM L2 VSN. The ISID of the CVLAN is displayed under the attribute "INTERFACE".

   ```
   show vlan info fdb-entry
   ```

3. To display the VLAN remote MAC table for a CVLAN, enter:

   ```
   show vlan info remote-mac-table <vid>
   ```

Variable definitions

Use the data in the following table to use the CVLAN ISID show commands.

Variable	Value
{all|config|discover}	• all: displays all ISID entries • config: displays configured ISID entries • discover: displayes discovered ISID entries
[vlan <*vid*>]	Displays information for the specified SPBM VLAN.
[id <*isid*>]	Displays information for the specified ISID.
[nick-name <*nickname*>]	Displays information for the specified nickname.

Job aid

The following sections describe the fields in the outputs for the CVLAN ISID show commands.

show vlan info i-sid

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

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

show vlan info remote-mac-table

The following table describes the fields in the output for the **`show vlan info remote-mac-table`** command.

Parameter	Description
VLAN	Indicates the VLAN for the remote MAC entry.
STATUS	Indicates the status of the remote MAC entry.
MAC-ADDRESS	Indicates the MAC address of the remote MAC entry.
DEST-MAC	Indicates the MAC address of the destination remote MAC.
BVLAN	Indicates the B-VLAN of the remote MAC entry.
DEST-SYSNAME	Indicates the Destination System Name of the remote MAC entry.
PORTS	Indicates the ports associated with the remote MAC entry.
SMLTREMOTE	Indicates the MAC that is learned locally or learned in the IST peer and synched via IST.

Displaying the multicast FIB, unicast FIB, and unicast tree

Use the following procedure to display SPBM multicast FIB, unicast FIB, and the unicast tree.

Procedure steps

1. To display the SPBM multicast FIB, enter:

   ```
   show isis spbm multicast-fib [vlan <value>] [i-sid <value>]
   [nick-name <value>] [summary]
   ```

2. To display the SPBM unicast FIB, enter:

   ```
   show isis spbm unicast-fib [b-mac <value>] [vlan <value>]
   [summary]
   ```

3. To display the SPBM unicast tree, enter:

   ```
   show isis spbm unicast-tree <vlan> [destination <value>]
   ```

4. To display the IS-IS SPBM multicast-FIB calculation results by I-SID, enter:

   ```
   show isis spbm i-sid {all|config|discover} [vlan <value>] [id
   <isid>] [nick-name <value>]
   ```

Variable definitions

Use the data in the following table to use the SPBM show commands.

Variable	Value
[vlan <*value*>]	Displays the FIB information for the specified VLAN.
[i-sid <*value*>]	Displays the FIB information for the specified ISID.
[nick-name <*value*>]	Displays the FIB information for the specified nickname.
[summary]	Displays a summary of the FIB.
[b-mac <*value*>]	Displays the FIB for the specified BMAC.
[destination <*value*>]	Displays the unicast tree for the specified destination.
{all|config|discover}	Specifies the multicast-FIB calculation results to display: • all: displays results for all ISID entries • config: displays results for configured ISID entries • discover: displays results for discovered ISID entries
[id <*isid*>]	Displays calculation results for the specified ISID.

Job aid

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

show isis spbm multicast-fib

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

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

show isis spbm unicast-fib

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

Parameter	Description
DESTINATION ADDRESS	Indicates the destination MAC Address of the unicast FIB entry.
BVLAN	Indicates the B-VLAN of the unicast FIB entry.
SYSID	Indicates the system identifier of the unicast FIB entry.
HOST-NAME	Indicates the 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.

show isis spbm i-sid

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

Parameter	Description
ISID	Indicates the IS-IS SPBM I-SID identifier.
SOURCE NAME	Indicates the nickname of the node where this I-SID was configured or discovered.

Parameter	Description
	Note: SOURCE NAME is equivalent to nickname.
VLAN	Indicates the B-VLAN where this I-SID was configured or discovered.
SYSID	Indicates the system identifier.
TYPE	Indicates the SPBM I-SID type; either configured or discovered.

Displaying global IS-IS parameters

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

Procedure steps

1. Display IS-IS configuration information:

   ```
   show isis info
   ```

2. Display the IS-IS system-id:

   ```
   show isis system-id
   ```

3. Display IS-IS net info:

   ```
   show isis net
   ```

Job aid

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

show isis info

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

Parameter	Description
Adminstate	Indicates the administrative state of the router.
RouterType	Indicates whether a router is Level 1 or 1/2.
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.

Parameter	Description
Overload-on-startup	Indicates the IS-IS overload-on-startup value in seconds.
Overload	Indicates if there is an overload condition.
Csnp Interval	Indicates the interval between CSNP updates in seconds.
PSNP Interval	Indicates the interval between PSNP updates in seconds.
Rxmt LSP Interval	Indicates the received LSP time interval.
spf-delay	Indicates the interval between successive SPF runs in milliseconds.
Router Name	Indicates the IS-IS name of the router.
ip source-address	Indicates the IP source address used for SPBM Native IP shortcuts.
Num of Interfaces	Indicates the number of interfaces on the router.
Num of Area Addresses	Indicates the number of area addresses on the router.

show isis system-id

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

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

Displaying IS-IS areas

Use the following procedure to display IS-IS areas.

Procedure steps

1. Display IS-IS areas:

   ```
   show isis area
   ```

2. Display the IS-IS manual areas:

   ```
   show isis manual-area
   ```

Job aid

show isis area

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

Parameter	Description
ADDRESS	Shows the area address defined. A Level 1 or Level 2 Intermediate System computes the values of area addresses (the set of area addresses for this Level 1 area), by forming the union of the sets of manual area addresses reported in the area address field of all Level 1 LSPs. The area addresses that are computed are those with LSP number zero in the local Intermediate System's link state database from all ISs which are reachable through Level 1 routing.

show isis manual-area

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.

Displaying IS-IS interface parameters

Displaying IS-IS interface parameters

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

Procedure steps

1. Display IS-IS interface configuration and status parameters (including adjacencies):

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

2. Display IS-IS interface authentication configuration:

```
show isis int-auth
```

3. Display IS-IS interface timers:

```
show isis int-timers
```

4. Display IS-IS circuit level parameters:

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

Variable definitions

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

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

Job aid

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

show isis interface

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

Parameter	Description
IFIDX	Indicates the interface index for the Ethernet or MLT interface.
TYPE	Indicates the type of interface configured (in this release, only pt-pt is supported).
LEVEL	Indicates the level of the IS-IS interface (Level 1 [default] or Level 2).
OP-STATE	Shows the physical connection state of the interface.
ADM-STATE	Shows the configured state of the interface.
ADJ	Shows how many adjacencies are learned through the interface.
UP-ADJ	Shows how many adjacencies are active through the interface.
SPBM-L1-Metric	Indicates the SPBM instance l1-metric on the IS-IS interface.

show isis int-auth

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

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

show isis int-timers

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

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

show isis int-ckt-level

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

Parameter	Description
IFIDX	Shows the interface index for the VLAN or MLT interface.
LEVEL	Shows the level of the IS-IS interface (Level 1 [default] or Level 2).
DIS	Shows the circuit's Designated Intermediate System (DIS).
CKTID	Shows the circuit number of the interface.

Displaying IS-IS LSDB and adjacencies

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

Procedure steps

1. Display the IS-IS LSDB:

   ```
   show isis lsdb [level <value>] [sysid <sys-id>] [lspid <lsp-id>] [tlv <value>] [detail]
   ```

2. Display IS-IS adjacencies:

   ```
   show isis adjacencies
   ```

Variable definitions

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

Variable	Value
[level <*value*>]	Displays the LSDB for the specified level: l1, l2, or l12.
[sysid <*sys-id*>]	Displays the LSDB for the specified system ID.
[lspid <*lsp-id*>]	Displays the LSDB for the specified LSP ID.
[tlv <*value*>]	Displays the LSDB by TLV type: 1–184.
[detail]	Displays detailed information.

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 (Level 1, Level 2, or Level 1/2) of the external router.
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

Parameter	Description
	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]). It 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 router's system ID.
HOST-NAME	Indicates the hostname listed in the LSP, or the system name if host name is not configured.

Displaying IS-IS statistics and counters

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

Procedures steps

1. Display IS-IS system statistics:

```
show isis stats
```

2. Display IS-IS interface counters:

```
show isis int-counters
```

3. Display IS-IS L1 control packet counters:

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

4. Display IS-IS L2 control packet counters:

```
show isis int-l2-cntl-pkts
```

Job aid

show isis stats

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

Parameter	Description
LEVEL	Shows the level of the IS-IS interface (Level 1 [default] or Level 2).
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 Ethernet Routing Switch 8800/8600 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 [default] or Level 2).

Parameter	Description
AUTH FAILS	Shows the number of times authentication has failed per interface.
ADJ CHANGES	Shows the number of times the adjacencies have changed.
INIT FAILS	Shows the number of times the adjacency has failed to establish.
REJ ADJ	Shows the number of times the adjacency was rejected by another router.
ID LEN	Shows the ID field length mismatches.
MAX AREA	Shows the maximum area address mismatches.
LAN DIS CHANGES	Shows the number of times the DIS has changed.

show isis int-l1-cntl-pkts

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

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

show isis int-l2-cntl-pkts

The following table describes the fields in the output for the **`show isis int-l2-cntl-pkts`** command.

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

Chapter 8: Configuring SPBM using the ACLI

This chapter describes how to configure SPBM using the ACLI.

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.

Procedure steps

1. From the global configuration mode, enable SPBM globally:

   ```
   [no] [default] spbm
   ```

2. From the global configuration mode, enter the router IS-IS configuration mode:

   ```
   router isis
   ```

3. From the router IS-IS configuration mode, create the SPBM instance (in this release, only one SPBM instance is supported):

   ```
   [no] [default] spbm <inst>
   ```

4. Exit router IS-IS configuration mode:

   ```
   exit
   ```

5. From the global configuration mode, create the SPBM backbone VLAN (B-VLAN):

   ```
   vlan create <vlan-id> type spbm-bvlan
   ```

6. From the global configuration mode, enter the router IS-IS configuration mode:

   ```
   router isis
   ```

7. From the router IS-IS configuration mode, add the SPBM B-VLAN to the SPBM instance:

   ```
   [no] spbm <inst> b-vid {<vlan-id [-vlan-id][,...]} [primary
   <vid>]
   ```

8. From the router IS-IS configuration mode, configure the system nickname (2.5 bytes in the format <x.xx.xx>):

   ```
   [no] [default] spbm <inst> nick-name <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 (using the `config isis system-id <system-id>` command) . This helps to recognize source and destination addresses for troubleshooting purposes.

9. From the router IS-IS configuration mode, configure an IS-IS manual area (1-13 bytes in the format <xx.xxxx.xxxx...xxxx>. In this release, only one manual area is supported.):

   ```
   [no] manual-area <area>
   ```

10. Exit router IS-IS configuration mode:

    ```
    exit
    ```

11. From the global configuration mode, specify the ports or MLTs that are going to link to the SPBM network:

    ```
    interface {GigabitEthernet <slot/port> | mlt <mltid> }
    ```

12. From the interface configuration mode, enable tagging on the selected ports or MLTs:

    ```
    [no] [default] encapsulation dot1q
    ```

13. From the interface configuration mode, create an IS-IS circuit/interface on the selected ports or MLTs:

    ```
    [no] isis
    ```

14. From the interface configuration mode, enable the SPBM instance on the IS-IS interfaces:

    ```
    [no] [default] isis spbm <inst>
    ```

15. From the interface configuration mode, enable the IS-IS circuit/interface on the selected ports or MLTs:

    ```
    [no] [default] isis enable
    ```

16. Exit interface configuration mode:

    ```
    exit
    ```

17. From the global configuration mode, enable IS-IS globally:

    ```
    [no] [default] router isis enable
    ```

18. To display the SPBM configurations, enter:

    ```
    show isis spbm
    ```

19. To display the global IS-IS configuration, enter:

    ```
    show isis
    ```

20. To display the interface IS-IS configuration, enter:

```
show isis interface
```

Important:

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

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

Variable definitions

Use the data in the following table to configure the SPBM and IS-IS parameters.

Variable	Value
[no] [default] spbm	Enables or disables SPBM globally on the switch. Use the no or default options to disable SPBM globally.
router isis	Sets the configuration mode to config-isis for access to IS-IS configuration commands.
[no] [default] spbm <*inst*>	Creates the SPBM instance. In this release, only one SPBM instance is supported. Use the no or default options to delete the specified instance.
[no] vlan create <*vlan-id*> type spbm-bvlan	Creates an SPBM Backbone VLAN (B-VLAN). You can optionally specify a name and color for the SPBM B-VLAN. • name <value> specifies the name of VLAN {string length 0..64}. • color <value> specifies the color of the VLAN {0..32}. Use the no option to delete the specified SPBM VLAN.
[no] spbm <*inst*> b-vid {<*vlan-id [-vlan-id][,...]*>} [primary <*vid*>]	Specifies the SPBM B-VLANs to add to the SPBM instance. Use the no option to remove the specified B-VLAN from the SPBM instance.
[no] [default] spbm <*inst*> nick-name <*nick-name*>	Specifies a nickname for the SPBM instance globally. Value is 2.5 bytes in the format <x.xx.xx>. Use the no or default options to delete the configured nickname.

Variable	Value
[no] manual-area <*area*>	Adds or deletes the specified IS-IS manual area. <area> is 1-13 bytes in the format <xx.xxxx.xxxx...xxxx>. Use the no option to delete the manual area
[no] [default] encapsulation dot1q	Enables or disables tagging on the specified interface. Use the no or default options to disable tagging on the specified interface.
[no] isis	Creates or deletes an IS-IS circuit/interface on the specified port or MLT. Use the no option to delete the IS-IS interface.
[no] [default] isis spbm <*inst*>	Enables or disables the SPBM instance on the specified port or MLT. Use the no or default options to disable SPBM on the specified interface.
[no] [default] isis enable	Enables or disables the IS-IS circuit/interface on the specified port or MLT. Use the no or default option to disable ISIS on the specified interface.
[no] [default] router isis enable	Enables IS-IS globally on the switch. Use the no or default options to disable IS-IS on the switch.

Configuring SMLT parameters for SPBM

Use the following procedure to configure the required SMLT parameters to allow SPBM to interoperate with SMLT on the switch.

Note:

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

Procedure steps

1. From the global configuration mode, disable IS-IS on the switch:

   ```
   no router isis enable
   ```

2. From the global configuration mode, enter the router IS-IS configuration mode:

   ```
   router isis
   ```

3. From the router IS-IS configuration mode, specify the BMAC of the IST peer, so that if it goes down, the local peer can take over forwarding for the failed peer.

```
spbm <inst> smlt-peer-bmac <mac-addr>
```

4. From the router IS-IS configuration mode, configure the virtual-bmac, which is shared and advertised by both peers.

   ```
   spbm <inst> smlt-virtual-bmac <virtual-mac>
   ```

5. From the global configuration mode, enable IS-IS on the switch:

   ```
   router isis enable
   ```

6. To display the SPBM SMLT configuration, enter:

   ```
   show isis spbm
   ```

Variable definitions

Use the data in the following table to configure the SMLT parameters for SPBM.

Variable	Value
smlt-peer-bmac <*mac-addr*>	Specifies the IST peer BMAC address.
smlt-virtual-bmac <*virtual-mac*>	Specifies a virtual MAC address that can be used by both peers.
show isis spbm	Displays SPBM SMLT info.

Configuring SPBM L2 VSN

After you have configured the SPBM infrastructure, you can enable SPBM L2 VSN using the following procedure.

Procedure steps

From the global configuration mode, map a customer VLAN (CVLAN) to an instance identifier (ISID):

```
[no] [default] vlan i-sid <vlan-id> <isid>
```

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.

Variable definitions

Variable	Value
[no] [default] vlan i-sid <*vlan-id*> <*isid*>	Specifies the customer VLAN (CVLAN) to associate with the ISID. Use the no or default options to remove the ISID from the the specified VLAN.

Configuring SPBM Native IP shortcuts

After you have configured the SPBM infrastructure, you can enable SPBM Native IP shortcuts to advertise IP routes across the SPBM network using the following procedure.

Procedure steps

1. From the global configuration mode, create a CLIP interface to use as the source address for SPBM Native IP shortcuts:

   ```
   interface loopback <1-256>
   ```

2. In the loopback interface configuration mode, configure an IP address for the CLIP interface:

   ```
   [no] ip address <1-256> <A.B.C.D/X>
   ```

3. Exit the loopback interface configuration mode:

   ```
   exit
   ```

4. From the global configuration mode, enter router IS-IS configuration mode:

   ```
   router isis
   ```

5. From the router IS-IS configuration mode, specify the CLIP interface as the source address for SPBM Native IP shortcuts:

   ```
   ip-source-address <A.B.C.D>
   ```

6. From the router IS-IS configuration mode, configure SPBM Native IP shortcuts:

   ```
   [no] [default] spbm <inst> ip enable
   ```

7. To view the status of SPBM Native IP shortcuts on the switch, enter:

   ```
   show isis spbm
   ```

8. From the router IS-IS configuration mode, identify routes on the local switch to be announced into the SPBM network:

   ```
   [no] [default] redistribute {direct | bgp | ospf | rip |
   static}
   ```

```
[no] [default] redistribute {direct | bgp | ospf | rip |
static} enable
```

9. From the router IS-IS (or global) configuration mode, apply the configured redistribution:

```
isis apply redistribute {direct | bgp | ospf | rip | static}
```

Variable definitions

Use the data in the following table to configure the SPBM Native IP shortcuts parameters.

Variable	Value
interface loopback <*1-256*>	Specifies an ID for the CLIP interface.
[no] ip address <*1-256*> <*A.B.C.D/X*>	Specifies an IP address and mask for the CLIP interface. Use the no option to delete the specified IP address.
ip-source-address <*A.B.C.D*>	Specifies the CLIP interface to use as the source address for SBPM Native IP shortcuts.
[no] [default] spbm <*inst*> ip enable	Enables or disables SPBM IP shortcut state. Use the no or default options to disable SPBM Native IP shortcuts.
[no] [default] redistribute {direct | bgp | ospf | rip | static}	Configures the redistribution of the specified protocol into the SPBM network. Use the no or default options to delete the configuration.
[no] [default] redistribute {direct | bgp | ospf | rip | static} enable	Enables the redistribution of the specified protocol into the SPBM network. Use the no or default options to disable the redistribution.
isis apply redistribute {direct | bgp | ospf | rip | static}	Applies the redistribution of the specified protocol into the SPBM network.

Configuring SPBM L3 VSN

After you have configured the SPBM infrastructure, you can enable SPBM L3 VSN to advertise IP routes across the SPBM network from one VRF to another using the following procedure.

Prerequisites

You must configure a VRF and IP VPN instance on the switch. For more information, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP VPN* (NN46205–520).

Procedure steps

1. From the global configuration mode, specify the VRF to configure:

   ```
   router vrf <vrf-name>
   ```

2. From VRF configuration mode, configure SPBM L3 VSN:

   ```
   [no] [default] i-sid <isid>
   ```

3. From VRF configuration mode, identify routes on the local switch to be announced into the SPBM network:

   ```
   [no] [default] isis redistribute {direct | bgp | ospf | rip | static}
   ```

   ```
   [no] [default] isis redistribute {direct | bgp | ospf | rip | static} enable
   ```

4. Exit the VRF configuration mode:

   ```
   exit
   ```

5. From the global configuration mode, apply the configured redistribution:

   ```
   isis apply redistribute {direct | bgp | ospf | rip | static} vrf <vrf-name>
   ```

6. To display the redistribution configuration, enter:

   ```
   show ip isis redistribute [vrf <vrf-name>] [vrfids <vrfids>]
   ```

Variable definitions

Use the data in the following table to configure the SPBM L3 VSN parameters.

Variable	Value
<*vrf-name*>	Specifies the VRF name.
[no] [default] i-sid <*isid*>	Assigns an ISID to the VRF being configured. Use the no or default option to remove the ISID to VRF allocation for this VRF.
[no] [default] isis redistribute {direct | bgp | ospf | rip | static}	Configures the redistribution of the specified protocol into the SPBM network. Use the no or default options to delete the configuration.
[no] [default] isis redistribute {direct | bgp | ospf | rip | static} enable	Enables the redistribution of the specified protocol into the SPBM network. Use the no or default options to disable the redistribution.

Variable	Value
isis apply redistribute {direct | bgp | ospf | rip | static} vrf <vrf-name>	Applies the redistribution of the specified protocol into the SPBM network on the specified VRF.

Configuring optional SPBM parameters

Use the following procedure to configure optional SPBM parameters.

Procedure steps

1. From the global configuration mode, configure the SPBM ethertype:

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

2. Enter the router IS-IS configuration mode:

   ```
   router isis
   ```

3. From the router IS-IS configuration mode, configure the optional LSDB trap global parameter. To configure this parameter, you must globally disable IS-IS on the switch:

   ```
   no router isis enable

   [no] [default] spbm <inst> lsdb-trap enable

   router isis enable
   ```

4. Exit router IS-IS configuration mode:

   ```
   exit
   ```

5. From the global configuration mode, specify an SPBM interface to configure:

   ```
   interface {GigabitEthernet <slot/port> | mlt <mltid> }
   ```

6. From the interface configuration mode, configure the optional SPBM interface parameters. To configure these parameters, you must disable IS-IS on the interface:

   ```
   no isis enable

   [[no] [default] isis spbm <inst> interface-type <if-type>]

   [[no] [default] isis spbm <inst> l1-metric <cost>]

   isis enable
   ```

Variable definitions

Use the data in the following table to configure the optional SPBM parameters.

Variable	Value
spbm ethertype {0x8100 | 0x88a8}	Configures the SPBM ethertype. The default value is 0x8100.
no router isis enable	Disables ISIS globally on the switch.
spbm <*inst*>	Specifies the SPBM instance ID.
[no] [default] spbm <*inst*> lsdb-trap enable	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable. Use the no or default options to disable LSDB traps.
router isis enable	Enables ISIS globally on the switch.
no isis enable	Disables IS-IS on the interface.
[no] [default] isis spbm <*inst*> interface-type <*if-type*>	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT. In this release, only the pt-pt interface type is supported. Use the no or default options to set this parameter to the default value of pt-pt.
[no] [default] isis spbm <*inst*> l1-metric <*cost*>	Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10. Range is 1–16777215. Use the no or default options to set this parameter to the default value of 10.
isis enable	Enables IS-IS on the interface.

Configuring optional IS-IS global parameters

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

Procedure steps

1. Enter the router IS-IS configuration mode:

   ```
   router isis
   ```

2. From the router IS-IS configuration mode, configure optional IS-IS global parameters:

   ```
   [[no] [default] csnp-interval <csnp-interval>]

   [[no] [default] is-type {l1}]

   [[no] [default] max-lsp-gen-interval <max-lsp-interval>]

   [[no] [default] metric {wide}]
   ```

```
[[no] [default] min-lsp-gen-interval <min-lsp-interval>]

[[no] [default] overload]

[[no] [default] overload-on-startup <overload-on-startup>]

[[no] [default] psnp-interval <psnp-interval>]

[[no] [default] retransmit-lsp-interval <retransmit-lsp-
interval>]

[[no] [default] spf-delay <spf-delay-time>]

[[no] [default] sys-name <sys-name>]

[[no] [default] system-id <system id>]
```

Important:

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

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

Variable definitions

Use the data in the following table to configure the optional IS-IS global parameters.

Variable	Value
[no] [default] csnp-interval <*csnp-interval*>	Specifies the CSNP interval in seconds. This is a system level parameter that applies for L1 CSNP generation on all interfaces. Default value is 10. Range is 1–600. Use the no or default options to set this parameter to the default value of 10.
[no] [default] is-type {l1}	Sets the router type globally: • l1: Level-1 router type • l12: Not valid in the current release. The default value is l1. Use the no or default options to set this parameter to the default value of l1.

Variable	Value
[no] [default] max-lsp-gen-interval <*max-lsp-interval*>	Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate system. The value must be greater than any value configured for min-lsp-gen-interval. Default value is 900 seconds. Range is 0–900. Use the no or default options to set this parameter to the default value of 900.
[no] [default] metric {wide}	Specifies the IS-IS metric type. Only wide is supported in this release. Use the no or default options to set this parameter to the default value of wide.
[no] [default] min-lsp-gen-interval <*min-lsp-interval*>	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. Default value is 30 seconds. Range is 1–65535. Use the no or default options to set this parameter to the default value of 30.
[no] overload	Sets or clears the overload condition. Default value is false. Use the no or default options to set this parameter to the default value of false.
[no] [default] overload-on-startup <*overload-on-startup*>	Sets the IS-IS overload-on-startup value in seconds. The overload-on-startup value is used as a timer to control when to send out LSPs with the overload bit cleared after IS-IS startup. The default value is 20. Range is 15–3600. Use the no or default options to set this parameter to the default value of 20.
[no] [default] psnp-interval <*psnp-interval*>	Specifies the PSNP interval in seconds. This is a system level parameter that applies for L1 PSNP generation on all interfaces. Default value is 2. Range is 1–120. Use the no or default options to set this parameter to the default value of 2.
[no] [default] retransmit-lsp-interval <*retransmit-lsp-interval*>	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. Default value is 5 seconds. Use the no or default options to set this parameter to the default value of 5.
[no] [default] spf-delay <*spf-delay-time*>	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.

Variable	Value
	The default value is 100 milliseconds. Use the no or default options to set this parameter to the default value of 100 milliseconds.
[no] [default] sys-name <*sys-name*>	Specifies a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system name comes from the host name configured at the system level. Use the no or default options to set this parameter to the default value (host name).
[system-id <*system id*>]	Specifies the IS-IS system ID for the switch. Use the no or default options to set this parameter to the default value (node BMAC).

Configuring optional IS-IS interface parameters

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

Procedure steps

1. From the global configuration mode, specify an IS-IS interface to configure:

   ```
   interface {GigabitEthernet <slot/port> | mlt <mltid>}
   ```

2. From the interface configuration mode, configure optional IS-IS interface parameters:

   ```
   [no] [default] isis
   [hello-auth type <type> [key <key>] [key-id <key-id>]]
   [l1-dr-priority <l1-dr-priority>]
   [l1-hello-interval <l1-hello-interval>]
   [l1-hello-multiplier <l1-hello-multiplier>]
   ```

Variable definitions

Use the data in the following table to configure the optional IS-IS interface parameters.

Variable	Value
hello-auth type <*type*> [key <*key*>] [key-id <*key-id*>]	Specifies the authentication type used for IS-IS hello packets on the interface. <type> can be one of the following: • none • simple: if selected, you can also specify a key value • hmac-md5: if selected, you can also specify a key value and key-id Use the no or default options to set the hello-auth type to none.
l1-dr-priority <*l1-dr-priority*>	Configures the level 1 IS-IS designated router priority to the specified value. Default value is 64. Use the no or default options to set this parameter to the default value of 64.
l1-hello-interval <*l1-hello-interval*>	Configures the level 1 hello interval. Default value is 9 seconds. Use the no or default options to set this parameter to the default value of 9 seconds.
l1-hello-multiplier <*l1-hello-multiplier*>	Configures the level 1 hello multiplier. Default value is 3 seconds. Use the no or default options to set this parameter to the default value of 3 seconds. .

Chapter 9: Displaying SPBM and IS-IS using the ACLI

This chapter describes how to display SPBM and IS-IS parameters using the ACLI.

Displaying global SPBM parameters

Use the following procedure to display global SPBM parameters.

Procedure steps

1. To display the SPBM configuration, enter:

   ```
   show isis spbm
   ```

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

   ```
   show vlan basic
   ```

Job aid

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

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

Parameter	Description
SMLT-SPLIT-BEB	Specifies whether the switch is the primary or secondary IST peer.
SMLT-VIRTUAL-MAC	Specifies a virtual MAC address that can be used by both peers.
SMLT-PEER-BMAC	Specifies the IST peer BMAC address.

Displaying CVLAN ISID information

Use the following procedure to display CVLAN ISID information.

Procedure steps

1. To display the CVLAN to ISID associations:

   ```
   show vlan i-sid [<vid>]
   ```

2. To display the IS-IS SPBM multicast-FIB calculation results by I-SID, enter:

   ```
   show isis spbm i-sid {all|config|discover} [vlan <vid>] [id <isid>] [nick-name <nickname>]
   ```

3. You can also use the following command to view the ISID of the CVLAN for entries learned from the SPBM L2 VSN. The ISID of the CVLAN is displayed under the attribute "INTERFACE".

   ```
   show vlan mac-address-entry
   ```

4. To display the VLAN remote MAC table for a CVLAN, enter:

   ```
   show vlan remote-mac-table <vid>
   ```

Variable definitions

Use the data in the following table to use the CVLAN ISID show commands.

Variable	Value
{all|config|discover}	• all: displays all ISID entries • config: displays configured ISID entries • discover: displayes discovered ISID entries
[vlan <*vid*>]	Displays ISID information for the specified SPBM VLAN.
[id <*isid*>]	Displays ISID information for the specified ISID.
[nick-name <*nickname*>]	Displays ISID information for the specified nickname.

Job aid

The following sections describe the fields in the outputs for the CVLAN ISID show commands.

show vlan i-sid

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

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

show isis spbm i-sid

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

Parameter	Description
ISID	Indicates the IS-IS SPBM I-SID identifier.
SOURCE NAME	Indicates the nickname of the node where this I-SID was configured or discovered. **Note:** SOURCE NAME is equivalent to nickname.
VLAN	Indicates the B-VLAN where this I-SID was configured or discovered.
SYSID	Indicates the system identifier.
TYPE	Indicates the SPBM I-SID type; either configured or discovered.

show vlan remote-mac-table

The following table describes the fields in the output for the **`show vlan remote-mac-table`** command.

Parameter	Description
VLAN	Indicates the VLAN ID for this MAC address.
STATUS SMLTREMOTE	Indicates the status of this entry: • other • invalid • learned • self • mgmt

Parameter	Description
MAC-ADDRESS	Indicates the customer MAC address for which the bridge has forwarding and/or filtering information
DEST-MAC	Indicates the provider MAC address for which the bridge has forwarding and/or filtering information.
BVLAN	Indicates the B-VLAN ID for this MAC address.
DEST-SYSNAME	Indicates the system name of the node where the MAC address entry comes from.
PORTS	Either displays the value '0', or indicates the port on which a frame came from.

Displaying the multicast FIB, unicast FIB, and unicast tree

Use the following procedure to display SPBM multicast FIB, unicast FIB, and the unicast tree.

Procedure steps

1. To display the SPBM multicast FIB, enter:

   ```
   show isis spbm multicast-fib [vlan <value>] [i-sid <value>]
   [nick-name <value>] [summary]
   ```

2. To display the SPBM unicast FIB, enter:

   ```
   show isis spbm unicast-fib [b-mac <value>] [vlan <value>]
   [summary]
   ```

3. To display the SPBM unicast tree, enter:

   ```
   show isis spbm unicast-tree <vlan> [destination <value>]
   ```

Variable definitions

Use the data in the following table to use the SPBM show commands.

Variable	Value
[vlan <*value*>]	Displays the FIB for the specified SPBM VLAN.
[i-sid <*value*>]	Displays the FIB for the specified ISID.
[nick-name <*value*>]	Displays the FIB for the specified nickname.
summary	Displays a summary of the FIB.
[b-mac <*value*>]	Displays the FIB for the specified BMAC.
[destination <*value*>]	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 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 global IS-IS parameters

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

Procedure steps

1. Display IS-IS configuration information:

   ```
   show isis
   ```

2. Display the IS-IS system-id:

   ```
   show isis system-id
   ```

3. Display IS-IS net info:

   ```
   show isis net
   ```

Job aid

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

show isis

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

Parameter	Description
Adminstate	Indicates the administrative state of the router.
RouterType	Indicates the router Level: l1, l2, or l1/2.
System ID	Indicates the system ID.
Max LSP Gen Interval	Indicates the maximum time between LSP updates in seconds.
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 IS-IS overload-on-startup value in seconds.
Overload	Indicates if there is an overload condition.
Csnp Interval	Indicates the interval between CSNP updates in seconds.
PSNP Interval	Indicates the interval between PSNP updates in seconds.
Rxmt LSP Interval	Indicates the received LSP time interval.
spf-delay	Indicates the interval between successive SPF runs in milliseconds.
Router Name	Indicates the IS-IS name of the router.
ip source-address	Indicates the IP source address used for SPBM Native IP shortcuts.
Num of Interfaces	Indicates the number of interfaces on the router.
Num of Area Addresses	Indicates the number of area addresses on the router.

show isis system-id

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

Parameter	Description
SYSTEM-ID	Shows the system ID. Output from this show command come from the global IS-IS configuration of the system ID. There is one system ID configured. The system ID is up to 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 come from the global IS-IS configuration of the manual area and the configuration of the system ID. There is only one manual areasdefined and only one system ID. The manual area is from 1-13 bytes in length. The system ID is up to 6 bytes in length.

Displaying IS-IS areas

Use the following procedure to display IS-IS areas.

Procedure steps

Display the IS-IS manual areas:

```
show isis manual-area
```

Job aid

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

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

Displaying IS-IS interface parameters

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

Procedure steps

1. Display IS-IS interface configuration and status parameters (including adjacencies):

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

2. Display IS-IS interface authentication configuration:

   ```
   show isis int-auth
   ```

3. Display IS-IS interface timers:

   ```
   show isis int-timers
   ```

4. Display IS-IS circuit level parameters:

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

Variable definitions

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

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

Job aid

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

show isis interface

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

Parameter	Description
IFIDX	Indicates the interface index for the Ethernet or MLT interface.
TYPE	Indicates the type of interface configured (in this release, only pt-pt is supported).
LEVEL	Indicates the level of the IS-IS interface (Level 1 [default] or Level 2).
OP-STATE	Shows the physical connection state of the interface.
ADM-STATE	Shows the configured state of the interface.

Parameter	Description
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 l1-metric on the IS-IS interface.

show isis int-auth

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

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

show isis int-timers

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

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

show isis int-ckt-level

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

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

Displaying IS-IS LSDB and adjacencies

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

Procedure steps

1. Display the IS-IS LSDB:

   ```
   show isis lsdb [level <value>] [sysid <sys-id>] [lspid <lsp-id>] [tlv <value>] [detail]
   ```

2. Display IS-IS adjacencies:

   ```
   show isis adjacencies
   ```

Variable definitions

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

Variable	Value
[level <*value*>]	Displays the LSDB for the specified level: l1, l2, or l12.
[sysid <*sys-id*>]	Displays the LSDB for the specified system ID.
[lspid <*value*>]	Displays the LSDB for the specified LSP ID.
[tlv <*value*>]	Displays the LSDB by TLV type: 1–184.
[detail]	Displays detailed information.

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]). It 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 router's system ID.
HOST-NAME	Indicates the hostname listed in the LSP, or the system name if host name is not configured.

Displaying IS-IS statistics and counters

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

1. Display IS-IS system statistics:

   ```
   show isis statistics
   ```

2. Display IS-IS interface counters:

   ```
   show isis int-counters
   ```

3. Display IS-IS L1 control packet counters:

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

4. Display IS-IS L2 control packet counters:

   ```
   show isis int-l2-cntl-pkts
   ```

Job aid

show isis statistics

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

Parameter	Description
LEVEL	Shows the level of the IS-IS interface (Level 1 [default] or Level 2).
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 Ethernet Routing Switch 8800/8600 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 [default] or Level 2).
AUTH FAILS	Shows the number of times authentication has failed per interface.
ADJ CHANGES	Shows the number of times the adjacencies have changed.
INIT FAILS	Shows the number of times the adjacency has failed to establish.
REJ ADJ	Shows the number of times the adjacency was rejected by another router.
ID LEN	Shows the ID field length mismatches.
MAX AREA	Shows the maximum area address mismatches.
LAN DIS CHANGES	Shows the number of times the DIS has changed.

show isis int-l1-cntl-pkts

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

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

show isis int-l2-cntl-pkts

The following table describes the fields in the output for the **show isis int-l2-cntl-pkts** command.

Parameter	Description
IFIDX	Shows the interface index for the VLAN or Ethernet interface.
DIRECTION	Shows the packet flow (Trnasmitted or Received).

Parameter	Description
HELLO	Shows the amount of interface-level Hello packets.
LSP	Shows the amount of LSP packets.
CSNP	Shows the amount of CSNPs.
PSNP	Shows the amount of PSNPs.

Chapter 10: Configuring SPBM using EDM

This chapter describes how to configure SPBM and IS-IS using 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.

Procedure steps

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

19. Specify the Manual Area Address (valid value is 1–13 bytes in the format <xx.xxxx.xxxx...xxxx>).
20. Click **Insert**
21. 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.

22. In the AdminState field, click **on**, and click **Apply**.
23. Under the IS-IS tab, click the **Interfaces** tab.
24. Click **Insert** to create an IS-IS circuit.
25. In the **Index** field, specify the an index value for the IS-IS circuit.
26. In the **IfIndex** field, specify the port or MLT on which to create the IS-IS circuit.
27. Click **Insert**.
28. Select the newly created IS-IS circuit entry, and click **SPBM**.
29. In the **Interfaces SPBM** tab, click **Insert**.
30. In the **State** field, select **enable**.
31. Click **Insert** to enable the SPBM instance on the IS-IS circuit.
32. Under the IS-IS tab, click the **Interfaces** tab.
33. In the **AdminState** field for the IS-IS circuity entry, select **on** to enable the IS-IS circuit.
34. Click **Apply**.

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

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

Configuring SMLT parameters for SPBM

Use the following procedure to configure the required SMLT parameters to allow SPBM to interoperate with SMLT on the switch.

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

Procedure steps

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

Configuring SPBM L2 VSN

After you have configured the SPBM infrastructure, you can enable SPBM L2 VSN using the following procedure.

Procedure steps

1. From the navigation tree, choose **Configuration** > **VLAN** > **VLANs**.
2. Click the **Advanced** tab.
3. To map a CVLAN to an instance identifier (ISID), in the **I-sid** field, specify the ISID to associate with the specified VLAN.
4. 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.

Configuring SPBM Native IP shortcuts

After you have configured the SPBM infrastructure, you can enable SPBM Native IP shortcuts to advertise IP routes across the SPBM network using the following procedure.

Procedure steps

1. From the navigation tree, choose: **Configuration** > **IP** > **IP**.
2. Click the **Circuitless IP** tab
3. Click **Insert**.
4. In the **Interface** box, assign a CLIP interface number.
5. In the **Ip Address** box, type the IP address.
6. In the **Net Mask** box, type the network mask address.
7. Click **Insert**.
8. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
9. From the **Globals** tab, in the **IpSourceAddress** field, specify the CLIP interface to use as the source address for SBPm Native IP shortcuts.
10. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
11. Click the **SPBM** tab.
12. In the **IpShortcut** field, select **enable**.
13. Click **Apply**.
14. From the navigation tree, choose **Configuration** > **IP** > **Policy**.
15. Click the **Route Redistribution** tab.
16. Click **Insert** to identify routes on the local switch to be announced into the SPBM network.
17. Using the fields provided, specify the source protocols to redistribute into ISIS. In the **Protocol** field, ensure to specify **isis** as the destination protocol.
18. Click **Insert**.

Configuring SPBM L3 VSN

After you have configured the SPBM infrastructure, you can enable SPBM L3 VSN to advertise IP routes across the SPBM network from one VRF to another using the following procedure.

Prerequisites

You must configure a VRF and IP VPN instance on the switch. For more information, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP VPN* (NN46205–520).

Procedure steps

1. From the navigation tree, choose: **Configuration** > **IP** > **IP-VPN**.
2. Click the **VPN** tab.
3. To create an IP VPN instance, click **Insert**.
4. Click the ellipsis button (**...**), select a VRF to associate with the IP VPN, and click **OK**.
5. Click **Insert**.
6. In the **Enable** column, select **enable** to enable the IP VPN on the VRF.
7. In the **IsidNumber** column, specify an ISID to associate with the VPN.
8. Click **Apply**.
9. From the navigation tree, choose: **Configuration** > **IP** > **Policy**.
10. To identify routes on the local switch to be announced into the SPBM network, click the **Route Redistribution** tab.
11. Click **Insert**.
12. In the **DstVrfId** box, click the ellipsis button (**...**), select the destination VRF ID and click **Ok**.
13. In the **Protocol** box, click **isis** as the route destination.
14. In the **SrcVrfId** box, click (**...**) button, select the source VRF ID and click **Ok**.
15. In the **RouteSource** box, click the source protocol.
16. In the **Enable** box, click **enable**.
17. In the **RoutePolicy** box, click the ellipsis (**...**) button, choose the route policy to apply to the redistributed routes and click **Ok**.
18. Configure the other parameters as required.
19. Click **Insert**.
20. To apply the redistribution configuration, click the **Applying Policy** tab.
21. Select **RedistributeApply**, and then click **Apply**.

Enabling or disabling SPBM at the global level

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

Procedure steps

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

Variable definitions

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

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

Configuring SPBM parameters

Use the following procedure to configure SPBM global parameters.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
2. Click the **SPBM** tab.
3. To create an SPBM instance, click **Insert**.
4. Configure the SPBM parameters.
5. Click **Apply**.

Variable definitions

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

Variable	Value
Id	Specifies the SPBM instance ID. In this release, only one SPBM instance is supported.
NodeNickName	Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Specifies the primary SPBM B-VLANs to add to the SPBM instance.
Vlans	Specifies the SPBM B-VLANs to add to the SPBM instance.
LsdbTrap	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable.
IpShortcut	Enables or disables SPBM IP shortcut state.
SmltSplitBEB	Specifies whether the switch is the primary or secondary IST peer.
SmltBmac	Specifies a virtual MAC address that can be used by both peers.
SmltPeerBmac	Specifies the IST peer BMAC address.

Configuring interface SPBM parameters

Use the following procedure to configure SPBM interface parameters.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
2. Click the **Interfaces SPBM** tab.
3. Configure the SPBM interface parameters.
4. Click **Apply**.

Variable definitions

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

Variable	Value
Index	Specifies an Index value for the SPBM interface.
State	Specifies whether the SPBM interface is enabled or disabled.
Type	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT:

Variable	Value
	ptpt or bcast. In this release, only the ptpt interface type is supported.
L1Metric	Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10. Range is 1–16777215.

Configuring IS-IS global parameters

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

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
2. From the **Globals** tab, configure the global IS-IS parameters.
3. Click **Apply**.

Variable definitions

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

Variable	Value
AdminState	Specifies the global status of IS-IS on the switch: on or off.
LevelType	Sets the router type globally: • level1: Level-1 router type • level1and2: Level_L1_L2 router type. Not supported in this release. The default value is level1.
SystemId	Specifies the IS-IS system ID for the switch. Valid value is a 6–byte value in the format <xxxx.xxxx.xxxx>. **Important:** After you have configured the SPBM nickname and enabled IS-IS, if you require a change of the system ID, you must also change the nickname. However, for naming convention purposes or configuration purposes, you may not want to change the nickname. To maintain the same nickname with a different system ID, perform the following steps: 1. Disable IS-IS. 2. Change the system ID.

Variable	Value
	3. Change the nickname to a temporary one. 4. Enable IS-IS. 5. Disable IS-IS. 6. Change the nickname to the original nickname. 7. Enable IS-IS.
MaxLspGenInt	Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate system. The value must be greater than any value configured for RxmtLspInt. Default value is 900 seconds. Range is 0–900.
CsnpInt	Specifies the CSNP interval in seconds. This is a system level parameter that applies for L1 CSNP generation on all interfaces. Default value is 10. Range is 1–600.
RxmtLspInt	Specifies the minimum time between retransmission of an LSP. This defines how fast the switch resends the same LSP. This is a system level parameter that applies for L1 retransmission of LSPs. Default value is 5 seconds. Range is 1–300 seconds.
PSPNInterval	Specifies the PSNP interval in seconds. This is a system level parameter that applies for L1 PSNP generation on all interfaces. Default value is 2. Range is 1–120.
SpfDelay	Specifies the SPF delay in milliseconds. This value is used to pace successive SPF runs. The timer prevents two SPF runs from being scheduled very closely. The default value is 100 milliseconds. Range is 0–5000.
HostName	Specifies a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system name comes from the host name configured at the system level.
IPSourceAddressType	Specifies the Address type of the source address for SPBM Native IP shortcuts. In the current release, this value can only be ipv4.
IpSourceAddress	Specifies IP source address for SPBM Native IP shortcuts.

Configuring system-level IS-IS parameters

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

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **System Level** tab.
3. Configure the IS-IS system level parameters.
4. Click **Apply**.

Variable definitions

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

Variable	Value
Index	Specifies the level: l1 or l2. In this release, only l1 is supported.
State	Specifies the state of the level on the system.
MinLSPGenInt	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. Default value is 30 seconds. Range is 1–65535.
SetOverload	Sets or clears the overload condition. Default value is false. Possible values are true or false.
SetOverloadUntil	Sets the ISIS overload-on-startup value in seconds. The overload-on-startup value is used as a timer to control when to send out LSPs with the overload bit cleared after ISIS startup. The default value is 20. Range is 15–3600.
MetricStyle	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.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **Interfaces** tab.
3. Configure the IS-IS interface parameters.
4. Click **Apply**.

Variable definitions

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

Variable	Value
Index	The identifier of this circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
IfIndex	Specifies the interface on which the circuit is configured (port or MLT).
Type	Specifies the ISIS circuit type. In this release, only the ptToPt interface type is supported.
AdminState	Specifies the administrative state of the circuit: on or off.
OperState	Specifies the operational state of the circuit.
AuthType	Specifies the authentication type: • none • simple • hmac-md5
AuthKey	Specifies the authentication key, in 0 to 16 characters.
KeyId	Specifies the authentication key ID (1–255).
LevelType	Specifies the router type globally: • level1: Level-1 router type • level 1and2: Level_L1_L2 router type. Not supported in this release. The default value is level1.
NumAdj	Specifies the number of adjacencies on this circuit.
NumUpAdj	Specifies the number of adjacencies that are up.

Configuring IS-IS interface level parameters

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

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **Interfaces Level** tab.
3. Configure the IS-IS interface level parameters.
4. Click **Apply**.

Variable definitions

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

Variable	Value
Index	The identifier of this circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
Level	Specifies the router type globally: • l1: Level-1 router type • l12: Level_L1_L2 router type. Not supported in this release. The default value is l1.
ISPriority	Specifies an integer sub-range for IS-IS priority. Range is 0–127.
HelloMultiplier	Configures the level 1 hello multiplier. Default value is 3 seconds. Range is 1–600.
HelloTimer	Configures the level 1 hello interval. Default value is 9 seconds. Range is 1–600.
DRHelloTimer	Indicates the level 1 IS-IS designated router priority. Default value is 3. Range is 0–127.

Configuring SPBM on an interface

Use the following procedure to configure SPBM on an interface.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **Interfaces** tab.
3. Click the **SPBM** button.
4. In the Interfaces SPBM tab, click **Insert**.
5. Click **Insert**.

Variable definitions

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

Variable	Value
Index	Specifies an Index value for the SPBM interface.
SpbmId	Specifies the SPBM instance ID.
State	Specifies whether the SPBM interface is enabled or disabled.
Type	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT. In this release, only the pt-pt interface type is supported.
L1Metric	Configures the SPBM instance l1-metric on the IS-IS interface located on the specified port or MLT. The default value is 10. Range is 1–16777215.

Configuring an IS-IS Manual Area

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

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **Manual Area** tab.
3. Click **Insert**.
4. Specify an Area Address in the **AreaAddr** field, and click **Insert**.

Variable definitions

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

Variable	Value
AreaAddr	Specifies the IS-IS manual area. Valid value is 1-13 bytes in the format <xx.xxxx.xxxx...xxxx>. In this release, only one manual area is supported.

Configuring IS-IS redistribution

Use this procedure to configure IS-IS redistribution.

Procedure steps

1. From the navigation tree, select **Configuration** > **IP** > **IS-IS**.
2. Click the **Redistribute** tab.
3. Click **Insert**.
4. Complete the fields as required.
5. Click **Insert**.

Variable definitions

Use the data in the following table to configure the IS-IS redistribution.

Variable	Value
DstVrfId	Specifies the destination VRF ID used in the redistribution.
Protocol	Specifies the protocols that receive the redistributed routes (isis).
SrcVrfId	Specifies the source VRF ID used in the redistribution.
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables or disables a redistribution entry. The default is disable.
RoutePolicy	Specifies the route policy to be used for the detailed redistribution of external routes from a specified source into the IS-IS domain. The default is none.
Metric	Specifies the metric for the redistributed route. The value can be a range between 0 to 65535. The default value is 0. Avaya recommends that you use a value that is consistent with the destination protocol.

Variable	Value
MetricType	Specifies the metric type (applicable to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
Subnets	Indicates whether the subnets are advertised individually or suppressed (applicable to OSPF only).

Chapter 11: Displaying SPBM and IS-IS using EDM

This chapter describes how to display SPBM and IS-IS parameters using EDM.

Displaying SPBM and IS-IS summary information

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

Procedure steps

1. From the navigation tree, select **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **Protocol Summary** tab.

Variable definitions

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

Variable	Value
Globals ISIS	
AdminState	Indicates the global status of IS-IS on the switch.
SystemId	Indicates the IS-IS system ID for the switch. Valid value is a 6–byte value in the format <xxxx.xxxx.xxxx>
HostName	Indicates a name for the system. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system name comes from the host name configured at the system level.
Globals SPBM	
GobalEnable	Indicates whether SPBM is enabled or disabled at the global level.
NodeNickName	Indicates the nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
SmltSplitBEB	Indicates whether the switch is the primary or secondary IST peer.

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

Displaying the remote MAC table for a C-VLAN

Use the following procedure to view a the remote MAC table for a C-VLAN.

Procedure steps

1. From the navigation tree, select **Configuration** > **VLAN** > **VLANs** .
2. Highlight a C-VLAN, and click the **Bridge** button.
3. Under the Bridge tab, click the **Remote MAC** tab.

Variable definitions

Use the data in the following table to use the Remote MAC tab.

Variable	Value
VlanId	Indicates the VLAN ID for this MAC address.
Addr	Indicates the customer MAC address for which the bridge has forwarding and/or filtering information
DestAddr	Indicates the provider MAC address for which the bridge has forwarding and/or filtering information.
PrimaryBVlanId	Indicates the primary B-VLAN ID for this MAC address.
PrimaryDestSysName	Indicates the primary system name of the node where the MAC address entry comes from.
PrimaryPort	Either displays the value '0', or indicates the primary port on which a frame came from.
SecondaryBVlanId	Indicates the secondary B-VLAN ID for this MAC address
SecondaryDestSysName	Indicates the secondary system name of the node where the MAC address entry comes from.
SecondaryPort	Either displays the value '0', or indicates the secondary port on which a frame came from.
SmltRemote	Indicates the MAC address entry for the remote IST peer.
Status	Indicates the status of this entry: • other • invalid • learned

Variable	Value
	• self • mgmt

Displaying L1 Area information

Use the following procedure to display L1 area information.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**.
2. Click the **L1 Area** tab.

Variable definitions

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

Variable	Value
AreaAddr	Specifies an area address reported in a Level 1 LSP generated or received by this Intermediate System.

Displaying LSP summary information

Use the following procedure to display LSP summary information.

Procedure steps

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

Variable definitions

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

Variable	Value
Level	Specifies the level at which this LSP appears.
ID	Specifies the 8 byte LSP ID, consisting of the SystemID, Circuit ID, and Fragment Number.
Seq	Specifies the sequence number for this LSP.

Variable	Value
Checksum	Specifies the 16 bit Fletcher Checksum for this LSP.
LifetimeRemain	The remaining lifetime in seconds for this LSP.
HostName	The hostname listed in LSP, or the system name if host name is not configured.

Displaying IS-IS adjacencies

Use the following procedure to display IS-IS adjacency information

Procedure steps

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

Variable definitions

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

Variable	Value
Index	A unique value identifying the IS adjacency from all other such adjacencies on this circuit. This value is automatically assigned by the system when the adjacency is created.
AdjIfIndex	Specifies the IS-IS interface on which the adjacency is found.
Usage	Specifies how the adjacency is used. On a point-to-point link, this can be level1and2. But on a LAN, the usage is level1 on the adjacency between peers at L1, and level2 for the adjacency between peers at L2
State	Specifies the state of the adjacency: • down • initializing • up • failed
LastUpTime	Indicates when the adjacency most recently entered the state 'up', measured in hundredths of a second since the last re-initialization of the network management subsystem. Displays 0 if the adjacency has never been in state 'up'.

Variable	Value
NeighPriority	Specifies the priority of the neighboring Intermediate System for becoming the Designated Intermediate System.
HoldTimer	Specifies the holding time in seconds for this adjacency. This value is based on received IIH PDUs and the elapsed time since receipt.
NeighSysID	Specifies the system ID of the neighboring Intermediate System.
AdjHostName	Specifies the hostname listed in the LSP, or the system name if host name is not configured.

Displaying the SPBM I-SID information

Use the following procedure to display the SPBM I-SID information.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
2. Click the **I-SID** tab.

Variable definitions

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

Variable	Value
SysId	Indicates the system identifier.
Vlan	Indicates the B-VLAN where this I-SID was configured or discovered.
Isid	Indicates the IS-IS SPBM I-SID identifier.
NickName	Indicates the nickname of the node where this I-SID was configured or discovered.
HostName	Indicates the host name listed in the LSP, or the system name if the host name is not configured.
Type	Indicates the SPBM I-SID type; either configured or discovered.

Displaying SPBM nicknames

Use the following procedure to display SPBM nicknames.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
2. Click the **Nick Names** tab.

Variable definitions

Use the data in the following table to use the Nick Names tab.

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

Displaying the IP Unicast FIB

Use the following procedure to display the IP Unicast FIB.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
2. Click the **IP Unicast FIB** tab.

Variable definitions

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

Variable	Value
VrfId	Specifies the VRF ID of the IP unicast FIB entry, 0 indicates NRE.

Variable	Value
DestinationIpAddrType	Specifies the address type of the destination IP Address.
DestinationIpAddr	Specifies the destination IP Address of the IP unicast FIB entry.
DestinationMask	Specifies the destination IP mask of the IP unicast FIB entry
NextHopBmac	Specifies the nexthop BMAC of the IP unicast FIB entry.
Vlan	Specifies the VLAN of the IP unicast FIB entry.
Isid	Specifies the ISID of the IP unicast FIB entry.
NextHopName	Specifies the nexthop hostname of the IP unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the IP unicast FIB entry.

Displaying the Unicast FIB

Use the following procedure to display the Unicast FIB.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
2. Click the **Unicast FIB** tab.

Variable definitions

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

Variable	Value
SysId	Specifies the system ID of the node where the unicast FIB entry originated.
Vlan	Specifies the VLAN of the unicast FIB entry.
DestinationMacAddr	Specifies the destination MAC Address of the unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the unicast FIB entry.
HostName	Specifies the host name of the node where unicast FIB entry originated.

Variable	Value
Cost	Specifies the cost of the unicast FIB entry.

Displaying the Multicast FIB

Use the following procedure to display the Multicast FIB.

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **SPBM**.
2. Click the **Multicast FIB** tab.

Variable definitions

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

Variable	Value
SysId	System ID of the node where the multicast FIB entry originated.
Vlan	VLAN of the multicast FIB entry.
McastDestMacAddr	Multicast destination MAC Address of the multicast FIB entry
Isid	ISID of the multicast FIB entry.
OutgoingPorts	Nni port of the multicast FIB entry.
HostName	Host name of the node where the multicast FIB entry originated.

Displaying IS-IS system statistics

Use the following procedure to display IS-IS system statistics.

Procedure steps

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

Variable definitions

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

Variable	Value
CorrLSPs	Indicates the number of corrupted in-memory LSPs detected. LSPs received from the wire with a bad checksum are silently dropped and not counted.
AuthFails	Indicates the number of authentication key failures recognized by this Intermediate System.
LSPDbaseOloads	Indicates the number of times the LSP database has become overloaded.
ManAddrDropFromAreas	Indicates the number of times a manual address has been dropped from the area.
AttmptToExMaxSeqNums	Indicates the number of times the IS has attempted to exceed the maximum sequence number.
SeqNumSkips	Indicates the number of times a sequence number skip has occurred.
OwnLSPPurges	Indicates the number of times a zero-aged copy of the system's own LSP is received from some other node.
IDFieldLenMismatches	Indicates the number of times a PDU is received with a different value for ID field length to that of the receiving system.
PartChanges	Indicates partition changes.

Displaying IS-IS interface counters

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

Procedure steps

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

Variable definitions

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

Variable	Value
Index	A unique value identifying the IS-IS interface.
AdjChanges	The number of times an adjacency state change has occurred on this circuit.
InitFails	The number of times initialization of this circuit has failed. This counts events such as PPP NCP failures.

Variable	Value
	Failures to form an adjacency are counted by isisCircRejAdjs
RejAdjs	The number of times an adjacency has been rejected on this circuit
IDFieldLenMismatches	The number of times an IS-IS control PDU with an ID field length different to that for this system has been received.
MaxAreaAddrMismatches	The number of times an IS-IS control PDU with a max area address field different to that for this system has been received
AuthFails	The number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	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 steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **Stats**.
2. Click the **Interface Control Packets** tab.

Variable definitions

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

Variable	Value
Index	A unique value identifying the IS-IS interface.
Direction	Indicates whether the switch is sending or receiving the PDUs.
Hello	Indicates the number of IS-IS Hellos frames seen in this direction at this level.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.

Variable	Value
CSNP	Indicates the number of IS-IS CSNP frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS PSNP frames seen in this direction at this level.

Graphing IS-IS interface counters

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

Procedure steps

1. From the navigation tree, choose **Configuration** > **IS-IS** > **IS-IS**
2. Click the **Interfaces** tab.
3. Select an existing interface.
4. Click the **Graph** button.

Variable definitions

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

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

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

Graphing IS-IS interface sending control packet statistics

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

Procedure steps

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

Variable definitions

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

Variable	Value
Hello	Indicates the number of IS-IS Hello (IIH) PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS CSNP frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS PSNPs seen in this direction at this level.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.

Variable	Value
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Graphing IS-IS interface receiving control packet statistics

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

Procedure steps

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

Variable definitions

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

Variable	Value
Hello	Indicates the number of IS-IS Hello PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS CSNP frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS PSNPs seen in this direction at this level.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Chapter 12: Configuring CFM using the CLI

This chapter describes how to configure Connectivity Fault Management (CFM) using the CLI.

Important:
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 Ethertype

Use this procedure to configure the CFM Ethertype.

Procedure steps

1. Configure the CFM Ethertype:

   ```
   config cfm ethertype <ethertype>
   ```

2. Display the CFM configuration:

   ```
   config cfm info
   ```

Variable definitions

Use the data in the following table to configure the Ethertype parameters.

Variable	Value
<*ethertype*>	Specifies the CFM Ethertype. Range is 0x1 – 0xffff. Default value is 0x8902.

Configuring CFM MD

Use this procedure to configure the CFM MD.

Procedure steps

1. Create the CFM MD:

```
config cfm md <md-name> create [index <value>] [maint-level
<value>]
```

2. Display the CFM MD configuration:

```
config cfm md <md-name> info
```

3. To display the CFM MD configuration, you can also enter:

```
show cfm md info [<md-name>]
```

4. To delete the CFM MD, enter:

```
config cfm md <md-name> delete
```

Variable definitions

Use the data in the following table to configure the MD parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
[index <*value*>]	Specifies a maintenance domain entry index. Range is 1–2147483647.
[maint-level <*value*>]	Specifies the MD maintenance level. Range is 0–7.

Configuring CFM MA

Use this procedure to configure the CFM MA.

Prerequisites

You must configure a CFM MD.

Procedure steps

1. Create the CFM MA:

```
config cfm md <md-name> ma <ma-name> create [index <value>]
```

2. Display the CFM MA configuration:

```
config cfm md <md-name> ma <ma-name> info
```

3. To display the CFM MA configuration, you can also enter:

```
show cfm ma info [md <value>] [ma <value>]
```

4. To delete the CFM MA, enter:

```
config cfm md <md-name> ma <ma-name> delete
```

Variable definitions

Use the data in the following table to configure the MA parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
<*ma-name*>	Specifies the MA name in a string of 1–22 characters.
[index <*value*>]	Specifies a maintenance association entry index. Range is 1–2147483647.

Configuring CFM MEP

Use this procedure to configure the CFM MEP.

Prerequisites

You must configure a CFM MD and MA.

Procedure steps

1. Create the CFM MEP:

   ```
   config cfm md <md-name> ma <ma-name> mep <mepID> create
   [state {enable|disable}]
   ```

2. Enable the CFM MEP

   ```
   config cfm md <md-name> ma <ma-name> mep <mepID> state
   {enable | disable}
   ```

3. To display the CFM MEP configuration, enter:

   ```
   show cfm mep info [md <value>] [ma <value>] [mep-id <value>]
   [detail]
   ```

Variable definitions

Use the data in the following table to configure the MEP parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
<*ma-name*>	Specifies the MA name in a string of 1–22 characters.
<*mepID*>	Specifies the MEP ID. Range is 1–8191.

Variable	Value
state {enable | disable}	Enables or disables the MEP.

Assigning a CFM nodal MEP to an SPBM B-VLAN

Use this procedure to assign a CFM nodal MEP to an SPBM B-VLAN.

Prerequisites

You must configure a CFM MD, MA, and MEP.

Procedure steps

1. Add nodal MEPs to the B-VLAN:

   ```
   config vlan <vid> add-nodal-mep {<mep> [,<mep>] [,...]}
   ```

2. Add nodal MIP level to the B-VLAN:

   ```
   add-nodal-mip-level {<level> [,<level>] [,...]}
   ```

Variable definitions

Use the data in the following table to configure the nodal MEP and MIP level.

Variable	Value
{*<mep> [,<mep>] [,...]*}	Specifies the nodal MEPs to add to the VLAN. <mep> = <mdName.maName.mepId>, for example md10.ma20.30
{*<level> [,<level>] [,...]*}	Specifies a MIP level within the range of 0-7.

Triggering a loopback test (LBM)

Use this procedure to trigger a loopback test.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

1. To trigger the loopback test, enter:

```
config cfm md <md-name> ma <ma-name> mep <mepID> lbm
<rmepMac> [-c <value>] [-p <value>] [-d <value>] [-fs
<value>] [-t <value>] [-i <value>] [-f <value>] [-sm <value>]
```

2. Alternately, you can trigger the loopback using the following command:

```
lbm {<md-name>.<ma-name>.<mepID>.<rmepMac>} [-c <value>] [-p
<value>] [-d <value>] [-fs <value>] [-t <value>] [-i <value>]
[-f <value>] [-sm <value>]
```

Variable definitions

Use the data in the following table to configure the loopback parameters.

Variable	Value
<md-name>	Specifies the MD name in a string of 1–22 characters.
<ma-name>	Specifies the MA name in a string of 1–22 characters.
<mepID>	Specifies the MEP ID. Range is 1–8191.
lbm *<rmepMac>*	Specifies the MAC to reach the MEP/MIP. Range is 00:00:00:00:00:00 — FF:FF:FF:FF:FF:FF.
[-c <*value*>]	Specifies the burst-count. Range is 1 – 200.
[-p <*value*>]	Specifies the priority. Range is 0 – 7.
[-d <*value*>]	Specifies the size. Range is 0 – 400.
[-fs <*value*>]	Specifies the frame-size. Range is 64–500.
[-t <*value*>]	Specifies the timeout interval in seconds. Range is 1 – 10.
[-i <*value*>]	Specifies the interframe interval in msecs. Range is 0 – 1000}
[-f <*value*>]	Specifies the testfill pattern. Range is 1 – 4: • 1: allZero • 2: allZeroCrc • 3: pseudoRandomBitSequence • 4: pseduoRandomBitSequenceCrc
[-sm <*value*>]	Specifies the source mode. Range is 1 – 2: • 1: nodal • 2: smltVirtual

Triggering linktrace (LTM)

Use the following procedure to trigger a linktrace.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

1. To trigger the linktrace, enter:

   ```
   config cfm md <md-name> ma <ma-name> mep <mepID> ltm
   <rmepMac> [-t < value>] [-p < value>] [-sm < value>] [-
   detail]
   ```

2. Alternately, you can trigger the linktrace using the following command:

   ```
   lbm {<md-name>.<ma-name>.<mepID>.<rmepMac>} [-t < value>] [-
   p < value>] [-sm < value>] [-detail]
   ```

Variable definitions

Use the data in the following table to configure the linktrace parameters.

Variable	Value
ltm *<rmepMac>*	Specifies the taret MAC to reach the MEP. Range is 00:00:00:00:00:00 — FF:FF:FF:FF:FF:FF.
[-t <*value*>]	Specifies the ttl value. Range is 1 – 255.
[-p <*value*>]	Specifies the priority. Range is 0–7.
[-sm <*value*>]	Specifies the source mode. Range is 1 – 2: • 1: nodal • 2: smltVirtual
[-detail]	Displays expanded linktrace results, including TLVs.

Triggering an L2ping

Use this procedure to trigger an L2ping.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger an L2ping, enter the following command:

```
l2ping {<vlan.RouterNodeName> | <vlan.SystemIdMac> |
<ipaddress>} [burst-count <value>] [data-tlv-size <value>]
[frame-size <value>] [testfill-pattern <value>] [prior ity <
value>] [time-out <value>] [source-mode <value>] [vrf <value>]
```

Variable definitions

Use the data in the following table to configure the L2ping parameters.

Variable	Value
{<*vlan.RouterNodeName*> | <*vlan.SystemIdMac*> | <*ipaddress*>}	Specifies the destination for the L2ping: • <vlan.RouterNodeName> = <1–4094>.<string> • <vlan.SystemIDMac> = <1–4094>.<XX:XX:XX:XX:XX:XX> • <ipaddress> = <A.B.C.D>
[burst-count <*value*>]	Specifies the burst count. Range is 1–200.
[data-tlv-size <*value*>]	Specifies the data TLV size. Range is 0–400.
[frame-size <*value*>]	Specifies the frame size. Range is 64–500.
[testfill-pattern <*value*>]	Specifies the testfill pattern. Range is 1–4: • 1: allZero • 2: allZeroCrc • 3: pseudoRandomBitSequence • 4: pseudoRandomBitSequenceCrc
[priority <*value*>]	Specifies the priority. Range is 0–7.
[time-out <*value*>]	Specifies the interval in seconds. Range is 1–10.
[source-mode <*value*>]	Specifies the source mode. Range is 1–2. • 1: nodal • 2: smltVirtual
[vrf <*value*>]	Specifies the VRF name.

Triggering an L2traceroute

Use this procedure to trigger an L2traceroute.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger an L2traceroute, enter the following command:

```
l2traceroute {<vlan.RouterNodeName> | <vlan.SystemIdMac> |
<ipaddress>} [ttl <value>] [priority <value>] [source-mode
<value>] [vrf <value>]
```

Variable definitions

Use the data in the following table to configure the L2traceroute parameters.

Variable	Value
{<*vlan.RouterNodeName*> | <*vlan.SystemIdMac*> | <*ipaddress*>}	Specifies the destination for the L2traceroute: • <vlan.RouterNodeName> = <1–4094>.<string> • <vlan.SystemIDMac> = <1–4094>.<XX:XX:XX:XX:XX:XX> • <ipaddress> = <A.B.C.D>
[ttl <*value*>]	Specifies the TTL value. Range is 1–255.
[priority <*value*>]	Specifies the priority. Range is 0–7.
[source-mode <*value*>]	Specifies the source mode. Range is 1–2. • 1: nodal • 2: smltVirtual
[vrf <*value*>]	Specifies the VRF name.

Triggering an L2tracetree

Use this procedure to trigger an L2tracetree.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger an L2tracetree, enter the following command:

```
l2tracetree {<vlan.isid | vlan.isid.RouterNodeName |
vlan.isid.SystemIdMac>} [ttl <value>] [priority <value>]
[source-mode <value>]
```

Variable definitions

Use the data in the following table to configure the L2tracetree parameters.

Variable	Value
{<*vlan.isid* | *vlan.isid.RouterNodeName* | *vlan.isid.SystemIdMac*>}	{{1..4094}.{1..16777215}} | {{1..4094}.{1..16777215}.str} | {{1..4094}.{1..16777215}.{00:00:00:00:00:00...FF:FF:FF:FF:FF:FF}}
[ttl <*value*>]	Specifies the TTL value. Range is 1–255.
[priority <*value*>]	Specifies the priority value. Range is 0–7.
[source-mode <*value*>]	Specifies the source mode. Range is 1–2: • 1: nodal • 2: smltVirtual

CFM Sample output

The following sections show some sample CFM output.

L2ping can use the system ID or the router name. The example below shows a case where the VLAN and MAC are given.

show isis adjacencies

```
*peter-4:5# show isis adjacencies
********************************************************************************
Command Execution Time: FRI SEP 17 06:50:30 2010 UTC
********************************************************************************
 ===============================================================================
                                ISIS Adjacencies
===============================================================================I
NTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME
-------------------------------------------------------------------------------
```

```
Port2/3  1 UP         00:37:37 127       19 0014.0dbf.a3df    ERS-LEE1
Port2/19 1 UP      1d 05:09:16 127       21 0014.0da2.b3df    ERS-MONTI0
--------------------------------------------------------------------------------
2 out of 2 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

l2ping

```
*peter-4:5# l2ping 500.00.14.0d.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
```

l2ping

```
*peter-4:5# l2ping 500.ERS-MONTI0

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

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

l2traceroute

```
*peter-4:5# l2traceroute 500.ERS-MONTI0

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

l2traceroute to ERS-MONTI0  (00:14:0d:a2:b3:df),  vlan 500
0    ERS-PETER4         (00:15:9b:11:33:df)
1    ERS-MONTI0          (00:14:0d:a2:b3:df)
```

l2traceroute

```
*peter-4:5# l2tracetree 500.1

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10 hops 64
1   ERS-PETER4        00:15:9b:11:33:df -> ERS-MONTI0        00:14:0d:a2:b3:df
2   ERS-MONTI0        00:14:0d:a2:b3:df -> ERS-LEE2          00:15:e8:b8:a3:df
```

L2ping and L2traceroute can also be used with an IP address. The following outputs show examples using an IP address.

l2ping

```
*peter-4:5# l2ping 10.1.1.1

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

L2 PING  Statistics : IP 10.1.1.1, paths found 1, replies 1
================================================================================
TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                              PKTS  PKTS  LOSS     MIN/MAX/AVE (us)
```

```
================================================================================
500  ERS-SHAMIM        (00:1a:8f:08:53:df)  1     0     100.00%  0/0/0.00
```

l2traceroute

```
*peter-4:5# l2traceroute 10.1.1.1

Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 10.1.1.1, paths found 1
================================================================
ERS-SHAMIM  (00:1a:8f:08:53:df),  vlan 500
0    ERS-PETER4          (00:15:9b:11:33:df)
1    ERS-MONTI0          (00:14:0d:a2:b3:df)
```

Chapter 13: Configuring CFM using the ACLI

This chapter describes how to configure Connectivity Fault Management (CFM) using the ACLI.

Important:
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 ethertype

Use this procedure to configure the CFM ethertype.

Procedure steps

1. From the Global Configuration mode, configure the CFM ethertype:

   ```
   cfm ethertype <ethertype>
   ```

2. Display the CFM ethertype configuration:

   ```
   show cfm ethertype
   ```

Variable definitions

Use the data in the following table to configure the ethertype parameters.

Variable	Value
<ethertype>	Specifies the CFM ethertype. Range is 0x1 – 0xffff. Default is 0x8902.

Configuring CFM MD

Use this procedure to configure the CFM MD.

Procedure steps

1. From the Global Configuration mode, create the CFM MD:

   ```
   cfm maintenance-domain <md-name> create [index <value>]
   [maint-level <value>] [level <value>]
   ```

2. Display the CFM MD configuration:

   ```
   show cfm maintenance-domain
   ```

3. To delete the CFM MD, enter:

   ```
   no cfm maintenance-domain <md-name>
   ```

Variable definitions

Use the data in the following table to configure the MD parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
[index <*value*>]	Specifies a maintenance domain entry index. Range is 1–2147483647.
[maint-level <*value*>]	Specifies the MD maintenance level when creating the MD. Range is 0–7.
[level <*value*>]	Modifies the MD maintenance level for an existing MD. Range is 0–7.

Configuring CFM MA

Use this procedure to configure the CFM MA.

Prerequisites

You must configure a CFM MD.

Procedure steps

1. From the Global Configuration mode, create the CFM MA:

   ```
   cfm maintenance-association <md-name> <ma-name> [index
   <value>]
   ```

2. Display the CFM MA configuration:

   ```
   show cfm maintenance-association
   ```

3. To delete the CFM MA, enter:

```
no cfm maintenance-association <md-name> <ma-name>
```

Variable definitions

Use the data in the following table to configure the MA parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
<*ma-name*>	Specifies the MA name in a string of 1–22 characters.
[index <*value*>]	Specifies a maintenance association entry index. Range is 1–2147483647.

Configuring CFM MEP

Use this procedure to configure the CFM MEP.

Prerequisites

You must configure a CFM MD and MA.

Procedure steps

1. From the Global Configuration mode, create the CFM MEP:

   ```
   cfm maintenance-endpoint <md-name> <ma-name> <mepID> [state
   {enable|disable}]
   ```

2. To enable or disable an existing CFM MEP, enter

   ```
   [no] cfm maintenance-endpoint <md-name> <ma-name> <mepID>
   enable
   ```

3. To display the CFM MEP configuration, enter:

   ```
   show cfm maintenance-endpoint
   ```

4. To delete an existing CFM MEP, enter:

   ```
   no cfm maintenance-endpoint <md-name> <ma-name> <mepID>
   ```

Variable definitions

Use the data in the following table to configure the MEP parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.

Variable	Value
<*ma-name*>	Specifies the MA name in a string of 1–22 characters.
<*mepID*>	Specifies the MEP ID. Range is 1–8191.
state {enable | disable}	Enables or disables the MEP when creating the MEP.
enable	Enables an existing MEP. Use this parameter with the `no` option to disable an existing MEP.

Assigning a CFM nodal MEP to an SPBM B-VLAN

Use this procedure to assign a nodal MEP to an SPBM B-VLAN.

Prerequisites

You must configure a CFM MD, MA, and MEP.

Procedure steps

1. Add nodal MEPs to the B-VLAN:

   ```
   vlan nodal-mep <vid> <md-name> <ma-name> <mepID>
   ```

2. To display the nodal MEP configuration, enter:

   ```
   show vlan nodal-mep <vid>
   ```

3. Add nodal MIP level to the B-VLAN:

   ```
   vlan nodal-mip-level <vid> {<level>[,<level>][,...]}
   ```

4. To display the nodal MIP level configuration, enter:

   ```
   show vlan nodal-mip-level <vid>
   ```

Variable definitions

Use the data in the following table to configure the nodal MEP parameters.

Variable	Value
<*vid*>	Specifies the VLAN ID.
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
<*ma-name*>	Specifies the MA name in a string of 1–22 characters.
<*mepID*>	Specifies the nodal MEPs to add to the VLAN. Range is 1–8191.

Variable	Value
{*<level>[,<level>][,...]*}	Specifies a MIP level list, within the range of 0-7.

Triggering a loopback test (LBM)

Use this procedure to trigger a loopback test.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger the loopback test, enter:

```
loopback <md-name> <ma-name> <mepID> <rmepMac> [burst-count
<value>] [data-tlv-size <value>] [frame-size <value>] [priority
<value>] [source-mode <value>] [testfill-pattern <value>]
[time-out <value>]
```

Variable definitions

Use the data in the following table to configure the loopback parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
<*ma-name*>	Specifies the MA name in a string of 1–22 characters.
<*mepID*>	Specifies the MEP ID. Range is 1–8191.
<*rmepMac*>	Specifies the remote MAC address to reach the MEP/MIP. Range is 00:00:00:00:00:00 — FF:FF:FF:FF:FF:FF.
[burst-count <*value*>]	Specifies the burst-count. Range is 1 – 200.
[data-tlv-size <*value*>]	Specifies the data TLV size. Range is 0 – 400.
[frame-size <*value*>]	Specifies the frame-size. Range is 64–500.
[priority <*value*>]	Specifies the priority. Range is 0 – 7.

Variable	Value
[source-mode <*value*>]	Specifies the source mode. Range is 1 – 2: • nodal • smltVirtual
[testfill-pattern <*value*>]	Specifies the testfill pattern. Range is 1–4: • 1: all-zero • 2: all-zero-crc • 3: pseudo-random-bit-sequence • 4: pseduo-random-bit-sequence-crc
[time-out <*value*>]	Specifies the timeout interval in seconds. Range is 1 – 10.

Triggering linktrace (LTM)

Use the following procedure to trigger a linktrace.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger the linktrace, enter:

```
linktrace <md-name> <ma-name> <mepID> <rmepMac> [detail]
[priority <value>] [source-mode <value>] [ttl-value <value>]
```

Variable definitions

Use the data in the following table to configure the linktrace parameters.

Variable	Value
<*md-name*>	Specifies the MD name in a string of 1–22 characters.
<*ma-name*>	Specifies the MA name in a string of 1–22 characters.
<*mepID*>	Specifies the MEP ID. Range is 1–8191.

Variable	Value
<*rmepMac*>	Specifies the target MAC address to reach the MEP. Range is 00:00:00:00:00:00 — FF:FF:FF:FF:FF:FF.
[detail]	Displays linktrace result details.
[priority <*value*>]	Specifies the priority. Range is 0–7.
[source-mode <*value*>]	Specifies the source mode. Range is 1 – 2: • nodal • smltVirtual
[ttl-value <*value*>]	Specifies the ttl value. Range is 1 – 255.

Triggering an L2 ping

Use this procedure to trigger an L2 ping.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger an L2 ping, enter the following command from user EXEC mode:

```
l2 ping {vlan <vlan> routernodename <RouterNodeName> | vlan
<vlan> systemidmac <SystemIdMac> | ip-address <ipaddress>}
[burst-count <value>] [data-tlv-size <value>] [frame-size
<value>] [testfill-pattern <value>] [prior ity < value>] [time-
out <value>] [source-mode <value>] [vrf <value>]
```

Variable definitions

Use the data in the following table to configure the L2 ping parameters.

Variable	Value
{vlan <*vid*> routernodename <*RouterNodeName*>} (vlan <*vid*> systemidmac <*SystemIdMac*>} {ip-address <*ipaddress*>}	Specifies the destination for the L2 ping: • <vid> = <1–4094> • <.RouterNodeName> = Word <0–255> • <SystemIDMac> = <XX:XX:XX:XX:XX:XX> • <ipaddress> = <A.B.C.D>

Variable	Value
[burst-count <*value*>]	Specifies the burst count. Range is 1–200.
[data-tlv-size <*value*>]	Specifies the data TLV size. Range is 0–400.
[frame-size <*value*>]	Specifies the frame size. Range is 64–500.
[testfill-pattern <*value*>]	Specifies the testfill pattern. Range is 1–4: • 1: all-zero • 2: all-zero-crc • 3: pseudo-random-bit-sequence • 4: pseudo-random-bit-sequence-crc
[priority <*value*>]	Specifies the priority. Range is 0–7.
[time-out <*value*>]	Specifies the interval in seconds. Range is 1–10.
[source-mode <*value*>]	Specifies the source mode. Range is 1–2. • nodal • smltVirtual
[vrf <*value*>]	Specifies the VRF name.

Triggering an L2 traceroute

Use this procedure to trigger an L2 traceroute.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger an L2 traceroute, enter the following command from user EXEC mode:

```
l2 traceroute {<vlan <vlan> routernodename <RouterNodeName> |
<vlan <vlan> systemidmac <SystemIdMac> | ip-address
<ipaddress>} [priority <value>] [source-mode <value>] [ttl
<value>] [vrf <value>]
```

Variable definitions

Use the data in the following table to configure the L2 traceroute parameters.

Variable	Value
{vlan <*vid*> routernodename <*RouterNodeName*>} (vlan <*vid*> systemidmac <*SystemIdMac*>} {ip-address <*ipaddress*>}	Specifies the destination for the L2 traceroute: • <vid> = <1–4094> • <RouterNodeName> = Word <0–255> • <SystemIDMac> = <XX:XX:XX:XX:XX:XX> • <ipaddress> = <A.B.C.D>
[ttl-value]<*value*>	Specifies the TTL value. Range is 1–255.
[priority <*value*>]	Specifies the priority. Range is 0–7.
[source-mode <*value*>]	Specifies the source mode. Range is 1–2. • nodal • smltVirtual
[vrf <*value*>]	Specifies the VRF name.

Triggering an L2 tracetree

Use this procedure to trigger an L2 tracetree.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

To trigger an L2 tracetree, enter the following command from user EXEC mode:

```
l2 tracetree {<vlan> <isid> routernodename <RouterNodeName> |
<vlan> <isid> systemidmac <SystemIdMac>} [priority <value>]
[source-mode <value>] [ttl—value <value>]
```

Variable definitions

Use the data in the following table to configure the L2 tracetree parameters.

Variable	Value
{<*vlan*> <*isid*> *routernodename* <*RouterNodeName*> | <*vlan*> <*isid*> *systemidmac* <*SystemIdMac*>}	{{1..4094}.{1..16777215}} | {{1..4094}.{1..16777215}.str} | {{1..4094}.{1..16777215}.{00:00:00:00:00:00...FF:FF:FF:FF:FF:FF}}

Variable	Value
[ttl-value<*value*>]	Specifies the TTL value. Range is 1–255.
[priority <*value*>]	Specifies the priority value. Range is 0–7.
[source-mode <*value*>]	Specifies the source mode. Range is 1–2: • 1: nodal • 2: smltVirtual

CFM sample output

The following sections show sample CFM output.

L2ping can use the system ID or the router name. The example below shows a case where the VLAN and MAC are given.

show isis adjacencies

```
*peter-4:5# show isis adjacencies
********************************************************************************
Command Execution Time: FRI SEP 17 06:50:30 2010 UTC
********************************************************************************
 ================================================================================
                                ISIS Adjacencies
================================================================================I
NTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME
--------------------------------------------------------------------------------
Port2/3  1 UP         00:37:37 127       19 0014.0dbf.a3df    ERS-LEE1
Port2/19 1 UP      1d 05:09:16 127       21 0014.0da2.b3df    ERS-MONTI0
--------------------------------------------------------------------------------
2 out of 2 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

l2 ping

```
*peter-4:5# l2 ping vlan 500 systemidmac 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
```

l2 ping

```
*peter-4:5# l2 ping vlan 500 routernodename ERS-MONTI0


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

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

l2 traceroute

```
*peter-4:5# l2 traceroute 500 routernodename ERS-MONTI0


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

l2traceroute to ERS-MONTI0  (00:14:0d:a2:b3:df),  vlan 500
0    ERS-PETER4         (00:15:9b:11:33:df)
1    ERS-MONTI0          (00:14:0d:a2:b3:df)
```

l2 tracetree

```
*peter-4:5# l2 tracetree 500 1

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10 hops 64
1   ERS-PETER4        00:15:9b:11:33:df -> ERS-MONTI0        00:14:0d:a2:b3:df
2   ERS-MONTI0        00:14:0d:a2:b3:df -> ERS-LEE2          00:15:e8:b8:a3:df
```

L2ping and L2traceroute can also be used with an IP address. The following outputs show examples using an IP address.

l2 ping

```
*peter-4:5# l2 ping ip-address 10.1.1.1


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

L2 PING  Statistics : IP 10.1.1.1, paths found 1, replies 1
================================================================================
TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                                PKTS  PKTS  LOSS     MIN/MAX/AVE (us)
================================================================================
500  ERS-SHAMIM        (00:1a:8f:08:53:df)  1     0     100.00%  0/0/0.00
```

l2 traceroute

```
*peter-4:5# l2 traceroute ip-address 10.1.1.1


Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 10.1.1.1, paths found 1
================================================================
ERS-SHAMIM  (00:1a:8f:08:53:df),  vlan 500
0    ERS-PETER4          (00:15:9b:11:33:df)
1    ERS-MONTI0          (00:14:0d:a2:b3:df)
```

Chapter 14: Configuring CFM using EDM

This chapter describes how to configure Configuration Fault Management (CFM) using EDM.

Important:
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 the CFM ethertype

Use this procedure to configure the CFM ethertype.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **CFM**.
2. In the **Global** tab, specify the ethertype.
3. Click **Apply**.

Variable definitions

Use the data in the following table to configure the ethertype.

Variable	Value
`EtherType`	Specifies the CFM ethertype. Range is 0x1 – 0xffff. The default value is 0x8902.

Configuring CFM MD

Use this procedure to configure a CFM MD.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **CFM**.
2. Click the **MD** tab.
3. Click **Insert**.

4. In the fields provided, specify an index value, name, and level for the MD.
5. Click **Insert**.

Variable definitions

Use the data in the following table to configure the MD parameters.

Variable	Value
Index	Specifies a maintenance domain entry index. Range is 1–2147483647.
Name	Specifies the MD name in a string of 1–22 characters.
NumOfMa	Indicates the number of MAs that belong to this maintenance domain.
Level	Specifies the MD maintenance level. Range is 0–7.
NumOfMip	Indicates the number of MIPs that belong to this maintenance domain
Type	Indicates the type of domain.

Configuring CFM MA

Use this procedure to configure a CFM MA.

Prerequisites

You must configure a CFM MD.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **CFM**.
2. Click the **MD** tab.
3. Highlight an existing MD, and click the **MaintenanceAssociation** button.
4. In the **MA** tab, click **Insert**.
5. In the fields provided, specify an index value and name for the MA.
6. Click **Insert**.

Variable definitions

Use the data in the following table to configure the MA parameters.

Variable	Value
DomainIndex	Specifies the maintenance domain entry index. Range is 1–2147483647.
AssociationIndex	Specifies a maintenance association entry index. Range is 1–2147483647.
DomainName	Specifies the MD name in a string of 1–22 characters.
AssociationName	Specifies the MA name in a string of 1–22 characters.
NumOfMep	Indicates the number of MEPs that belong to this maintenance association.

Configuring CFM MEP

Use this procedure to configure the CFM MEP.

Prerequisites

You must configure a CFM MD and MA.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **CFM**.
2. Click the **MD** tab.
3. Highlight an existing MD, and click the **MaintenanceAssociation** button.
4. In the **MA** tab, highlight an existing MA, and click the **MaintenanceEndpoint** button.
5. Click **Insert**.
6. In the fields provided, specify the ID and the administrative state of the MEP.
7. Click **Insert**.

Variable definitions

Use the data in the following table to configure the MEP parameters.

Variable	Value
DomainIndex	Specifies the MD index.
AssociationIndex	Specifies the MA index.
Id	Specifies the MEP ID. Range is 1–8191.

Variable	Value
DomainName	Specifies the MD name in a string of 1–22 characters.
AssociationName	Specifies the MA name in a string of 1–22 characters.
AdminState	Specifies the administrative state of the MEP.
MepType	Specifies the MEP type: • trunk • sg • endpt • vlan • port • endptClient • nodal • remotetrunk • remotesg • remoteendpt • remoteVlan • remotePort • remoteEndptClient
ServiceDescription	Specifies the service to which this MEP is assigned.

Configuring CFM nodal MEP

Use this procedure to configure the CFM nodal MEP.

Prerequisites

You must configure a CFM MD, MA, and MEP.

Procedure steps

1. From the navigation tree, select **Configuration** > **VLAN** > **VLANs**.
2. Click the **Advanced** tab.
3. Select an SPBM VLAN.
4. Click the **Nodal** button.

5. In the **NodalMepList** field, specify the nodal MEPs to add to the VLAN.
6. In the **NodalMIPLevelList** field, specify a MIP level list (up to 8 levels within the range of 0-7).
7. Click **Apply**.

Variable definitions

Use the data in the following table to configure the nodal MEP parameters.

Variable	Value
NodalMepList	Specifies the nodal MEPs to add to the VLAN, in the format <mdName.maName.mepId>, for example md10.ma20.30.
NumOfNodalMep	Indicates the number of nodal MEPs assigned to this VLAN
NodalMipLevelList	Specifies a MIP level list, up to 8 levels within the range of 0-7.
NumOfNodalMipLevel	Indicates the number of nodal MIP levels assigned to this VLAN that allows MIP functionality to be enabled on a per level per VLAN basis.

Configuring L2 ping

Use this procedure to configure an L2 ping.

Prerequisites

On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **L2Ping/ L2Trace Route**.
2. From the **L2Ping** tab, configure the L2 ping properties.
3. To initiate an L2 ping, highlight an entry and click the **Start** button.
4. To stop the L2 ping, click the **Stop** button.

Variable definitions

Use the data in the following table to configure the L2 ping parameters.

Variable	Value
VlanId	Identifies the customer VLAN.
DestMacAddress	Specifies the target MAC Address.
HostName	Specifies the target host name
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Ping transmission.
Messages	Specifies the number of L2Ping messages to be transmitted.
Status	Specifies the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Ping messages. • abort: the service aborted or is about to abort the L2Ping messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.
ResultOk	Indicates the result of the operation: • true: the L2Ping Messages will be (or have been) sent. • false: the L2Ping Messages will not be sent
Priority	Specifies a 3–bit value to be used in the VLAN header, if present in the transmitted frame.
TimeoutInt	Specifies the interval to wait for an L2Ping time-out. Range is 1–10 seconds. Default value is 3 seconds.
TestPattern	Specifies the test pattern to use in the L2Ping PDU: • allZero: Null signal without CRC-32 • allZeroCrc: Null signal with CRC-32 • pseudoRandomBitSequence: PRBS 1/2^31-1 without CRC-32 • pseudoRandomBitSequenceCrc: PBRS 1/2^31-1 with CRC-32 Default value is allZero.
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the Data Size is selected to be sent.

Variable	Value
FrameSize	Specifies the frame size. If the frame size is specified then the data size is internally calculated and the calculated data size is included in the data TLV. Range is 64–500.
SourceMode	Specifies the source modes of the transmit loopback service: • nodal • smltVirtual
SeqNumber	The transaction identifier/sequence number of the first loopback message (to be) sent
Result	Displays the L2Ping result.

Initiating an L2 traceroute

Use this procedure to trigger an L2 traceroute.

Prerequisites

On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **L2Ping/ L2Trace Route**.
2. Click the **L2 Traceroute** tab.
3. To start the traceroute, highlight an entry and click the **Start** button.
4. To stop the traceroute, click the **Stop** button.

Variable definitions

Use the data in the following table to configure the L2 traceroute parameters.

Variable	Value
VlanId	Specifies a value that uniquely identifies the Customer VLAN.
Priority	Specifies a 3 bit value to be used in the VLAN header, if present in the transmitted frame.
DestMacAddress	Specifies the target MAC address.

Variable	Value
HostName	Specifies the target host name.
DestIsHostName	Specifies whether the host name is (true) or is not (false) used for the L2Ping transmission.
Status	Indicates the status of the transmit loopback service: . • ready: the service is available. • transmit.tThe service is transmitting, or about to transmit, the L2Trace messages • abort: the service aborted or is about to abort the L2Trace messages. This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.
ResultOk	Indicates the result of the operation: • true: the L2Trace Messages will be (or have been) sent. • false: the L2Trace Messages will not be sent.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The value of the TTL Field in the L2Trace is defined by the originating MEP. The default value is 64.
SourceMode	Specifies the source mode of the transmit loopback service.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent.
Flag	L2Trace result flag indicating L2Trace status or error code: • none (1): No error • internalError (2): L2Trace Internal Error • invalidMac (3): Invalid Mac Address • mepDisabled (4): Mep must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received

Variable	Value
	• l2TraceToOwnMepMac (6): L2Trace to own Mep MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace • l2TraceLeafNode (9): On a leaf node in the ISID tree • l2TraceNotInTree (10): Not in the ISID tree • l2TraceSmltNotPrimary (11): Requested SMLT Source from Non-Primary Node

Viewing L2 traceroute results

Use this procedure to view L2 traceroute results.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **L2Ping/ L2Trace Route**.
2. Click the **L2 Traceroute** tab.
3. To view the traceroute results, highlight an entry and click the **Result** button.

The following table describes the fields in the L2 Traceroute Result tab.

Field	Definition
VlanId	A value that uniquely identifies the Customer Vlan.
SeqNumber	The transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which L2Trace's response is going to be returned.
Hop	The number of hops away from L2Trace initiator.
ReceiveOrder	An index to distinguish among multiple L2Trace responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Ttl field value for a returned L2Trace response.
SrcMac	Mac address of the MP that responds to the L2Trace request for this L2traceReply.

Field	Definition
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	The host name of the node that forwarded the L2Trace to the responding node.

Configuring L2 IP ping

Use this procedure to configure L2 IP ping

Prerequisites

On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **L2Ping/ L2Trace Route**.
2. Click the **L2 IP Ping** tab.
3. To add a new entry, click **Insert**, specify the destination IP address and optional parameters, and click **Insert**.
4. To start the L2 IP ping, highlight an entry and click the **Start** button.
5. To stop the L2 IP ping, click the **Stop** button.

Variable definitions

Use the data in the following table to configure the L2 IP Ping parameters.

Variable	Value
IpAddrType	Specifies the address type of destination IP Address (only IPv4 is supported).
IpAddr	Specifies the destination IP Address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the Name of the Virtual Router.
Messages	Specifies the number of L2IpPing messages to be transmitted per MAC/VLAN pair. Range is 1–200.

Variable	Value
Status	Specifies the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2IpPing messages. • abort: the service is aborted or about to abort the L2IpPing messages. This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time.
ResultOk	Indicates the result of the operation: • true: L2IpPing Messages will be or have been sent. • false: L2IpPing Messages will not be sent.
TimeoutInt	Specifies the interval to wait for an L2IpPing time-out with a range of 1–10 seconds with a default value of 3 seconds.
TestPattern	Specifies the test pattern to use in the L2IPPing PDU: • allZero: Null signal without CRC-32 • allZeroCrc: Null signal with CRC-32 • pseudoRandomBitSequence: PRBS 1/2^31-1 without CRC-32 • pseudoRandomBitSequenceCrc: PBRS 1/2^31-1 with CRC-32 Default value is allZero.
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the Data Size is selected to be sent. Range is 0–400.
PathsFound	Specifies the number of paths found to execute the l2ping command.

Viewing L2 IP Ping results

Use this procedure to view L2 IP ping results.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **L2Ping/ L2Trace Route**.
2. Click the **L2 IP Ping** tab.
3. To view the L2 IP ping results, highlight an entry and click the **Result** button.

Variable definitions

The following table describes the fields in the L2 IP Ping Result tab.

Field	Definition
IpAddrType	The address type of the destination IP Address.
IpAddr	Destination IP Address.
SendOrder	An index to distinguish among multiple L2Ping sessions. This value is assigned sequentially from 1. It denotes the order that sessions were sent. It correlates to the number of paths found.
VrfId	VRF ID.
VlanId	VLAN ID found from the L3 lookup and used for transmission.
DestMacAddress	An indication of the target MAC Address transmitted.
PortNum	Either the value '0', or the port number of the port used for the l2ping.
DestHostName	The host name of the responding node.
Size	The number of bytes of data sent.
PktsTx	Number of Packets transmitted for this vlan/mac.
PktsRx	Number of Packets received for this vlan/mac.
PercentLossWhole	Percentage of packet loss for this vlan/mac.
PercentLossFract	Percentage of packet loss for this vlan/mac.
MinRoundTrip	Minimum time for round-trip for this vlan/mac in us.
MaxRoundTrip	Maximum time for round-trip for this vlan/mac in us"
RttAvgWhole	Average time for round-trip for this vlan/mac in us.
RttAvgFract	Fractional portion of Average time for round-trip.
Flag	Result flag indicating status or error code: • 1 - No error • 2 - Internal Error • 3 - Invalid Ip

Field	Definition
	• 4 - L2Trace completed • 5 - Lookup failure for IP (no vlan/mac entries)

Configuring L2 IP traceroute

Use this procedure to configure L2 IP traceroute.

Prerequisites

On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **L2Ping/ L2Trace Route**.
2. Click the **L2 IP Traceroute** tab.
3. To add a new entry, click **Insert**, specify the destination IP address and, optionally, the TTL value, and click **Insert**.
4. To start the L2 IP traceroute, highlight an entry and click the **Start** button.
5. To stop the L2 IP traceroute, click the **Stop** button.

Variable definitions

Use the data in the following table to configure the L2 IP traceroute parameters.

Variable	Value
IPAdrrType	Specifies the address type of destination IP Address (only IPv4 is supported).
IPAddr	Specifies the destination IP Address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the Name of the Virtual Router.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The default value is 64

Variable	Value
Status	Indicates the status of the transmit loopback service: • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Trace messages. • abort: the service is aborted or about to abort the L2Trace messages. This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time.
ResultOk	Indicates the result of the operation: • true: the Trace Messages will be or have been sent. • false. the Trace Messages will not be sent
PathsFound	Specifies the number of paths found to execute the L2trace.

Viewing L2 IP traceroute results

Use this procedure to view L2 IP traceroute results.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **L2Ping/ L2Trace Route**.
2. Click the **L2 IP Traceroute** tab.
3. To view the L2 IP traceroute results, highlight an entry and click the **Result** button.

Variable definitions

The following table describes the fields in the L2 IP Traceroute Result tab.

Field	Definition
IpAddrType	The address type of destination IP Address.
IpAddr	Destination IP Address.
SendOrder	An index to distinguish among multiple L2Trace sessions. This value is assigned sequentially from 1. It

Field	Definition
	denotes the order that sessions were sent. It correlates to the number of paths found.
Hop	The number of L2 hops away from L2Trace initiator.
ReceiveOrder	An index to distinguish among multiple L2Trace responses with the same Send Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Ttl field value for a returned L2Trace response.
VrfId	VRF ID.
VlanId	Vlan found from L3 lookup and used for transmission.
DestMacAddress	An indication of the target MAC Address transmitted.
PortNum	Either the value '0', or the port number of the port used for the l2trace.
SeqNumber	The transaction identifier/sequence number used in linktrace message packet.
SrcMac	Mac address of the MP that responded to L2Trace request for this L2traceReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	The host name of the node that forwarded the L2Trace to the responding node.
Flag	L2Trace result flag indicating status or error code: • none (1): No error • internalError (2): L2Trace Internal Error • invalidMac (3): Invalid Mac Address • mepDisabled (4): Mep must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own Mep MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace

Triggering a loopback test

Use this procedure to trigger a loopback test.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **CFM**.
2. Click the **LBM** tab.
3. Configure the loopback test properties as required.
4. Click **Apply**.
5. To trigger the loopback test, in the Status field , select **transmit**.
6. Click **Apply**.

Variable definitions

Use the data in the following table to configure the loopback parameters.

Variable	Value
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the Maintenance EndPoint index value.
DomainName	Specifies the MD name in a string of 1–22 characters.
AssociationName	Specifies the MA name in a string of 1–22 characters.
DestMacAddress	Specifies the remote MAC address to reach the MEP/MIP. Range is 00:00:00:00:00:00 — FF:FF:FF:FF:FF:FF.
Messages	Specifies the number of loopback messages to be transmitted
VlanPriority	Specifies the priority. Range is 0 – 7.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent

Variable	Value
ResultOk	Indicates the result of the operation: • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent
Status	Indicates the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the Loopback messages. • abort: The service is aborted or about to abort the Loopback messages.
Result	Displays the LBM result.
TimeoutInt	Specifies the timeout interval in seconds. Range is 1 – 10. Default value is 3 seconds.
InterFrameInt	Specifies the interval between LBM frames with a range of (0..1000) msecs and a default value of 500 msecs. The value of 0 msecs indicates to send the frames as fast as possible.
TestPattern	Specifies the testfill pattern. Range is 1–4: • 1: allZero • 2: allZeroCrc • 3: pseudoRandomBitSequence • 4: pseduoRandomBitSequenceCrc
DataSize	Specifies the data TLV size. Range is 0 – 400.
FrameSize	Specifies the frame-size. Range is 64–500.
Sourcemode	Specifies the source mode. Range is: • nodal • smltVirtual

Triggering linktrace

Use the following procedure to trigger a linktrace.

Prerequisites

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **CFM**.
2. Click the **LTM** tab.
3. Configure the linktrace test properties as required.
4. Click **Apply**.
5. To trigger the linktrace test, in the Status field select **transmit**, and click **Apply**.

   OR

   Highlight an entry and click the **Start** button.
6. To stop the linktrace, click the **Stop** button.
7. To view the results of the linktrace, click the **Result** button.

Variable definitions

Use the data in the following table to configure the linktrace parameters.

Variable	Value
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the MEP index value.
DomainName	Specifies the MD name in a string of 1–22 characters.
AssociationName	Specifies the MA name in a string of 1–22 characters.
VlanPriority	Specifies the VLAN priority, a 3 bit value to be used in the VLAN tag, if present in the transmitted frame. Range is 0–7.
DestMacAddress	Specifies the remote MAC address to reach the MEP. Range is 00:00:00:00:00:00 — FF:FF:FF:FF:FF:FF.
Ttl	Indicates the number of hops remaining to this LTM. This value is decremented by 1 by each Bridge that handles the LTM. The decremented value is returned in the LTR. If the value is 0 on output, the LTM is not transmitted to the next hop. The value of the TTL Field in the LTM is specified at the originating MEP. The default value is 64. Range is 1 – 255.

Variable	Value
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent.
ResultOk	Indicates the result of the operation: • true. The Loopback Messages will be (or have been) sent. • false. The Loopback Messages will not be sent.
Status	Indicates the status of the transmit loopback service: • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the LTM messages. • abort: The service is aborted, or about to abort the LTM message
Result	Displays the LTM result.
Flag	Displays the LTM result flag indicating LTM status or error code. Each value represents a status or error case: • 1 - No error • 2 - Ltm Internal Error • 3 - Unknown Remote Maintenance End Point • 4 - Invalid Remote Maintenance End Point Mac Address • 5 - Unset Remote Maintenance End Point Mac Address • 6 - Mep must be enabled in order to perform LTM • 7 - No Ltr response received • 8 - Linktrace to own Mep MAC is not sent • 9 - Endpoint must be enabled in order to perform LTM • 10 - Pbt-trunk must be enabled in order to perform LTM • 11 - LTM completed • 12 - LTM leaf node
SourceMode	Specifies the source mode. Range is: • nodal • smltVirtual

Viewing linktrace results

Use this procedure to view linktrace results.

Procedure steps

1. From the navigation tree, select **Configuration** > **Edit** > **Diagnostics** > **CFM**.
2. Click the **LTM** tab.
3. Highlight an entry, and click the **Result** button.

Variable definitions

Use the data in the following table to use the Link Trace Replies tab.

Variable	Value
DomainIndex	Indicates the Maintenance Domain Index.
AssociationIndex	Indicates the Maintenance Association Index.
MepId	Indicates the Maintenance EndPoint ID.
SeqNumber	Indicates the transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which LTM response is going to be returned.
Hop	Indicates the number of hops away from the LTM initiator.
ReceiveOrder	Indicates the index value used to distinguish among multiple LTRs with the same LTR Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the LTRs.
Ttl	Indicates the Ttl field value for a returned LTR.
DomainName	Indicates the Maintenance Domain Name.
AssociationName	Indicates the Maintenance Association Name.
Forwarded	Indicates if a LTM was forwarded by the responding MP, as returned in the 'FwdYes' flag of the flags field.
TerminalMep	Displays a boolean value stating whether the forwarded LTM reached a MEP enclosing its MA, as returned in the Terminal MEP flag of the Flags field.
LastEgressIdentifier	Displays an octet field holding the Last Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Last Egress Identifier identifies the MEP

Variable	Value
	Linktrace Indicator that originated, or the Linktrace Responder that forwarded, the LTM to which this LTR is the response. This is the same value as the Egress Identifier TLV of that LTM.
NextEgressIdentifier	Displays an octet field holding the Next Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Next Egress Identifier Identifies the Linktrace Responder that transmitted this LTR, and can forward the LTM to the next hop. This is the same value as the Egress Identifier TLV of the forwarded LTM, if any. If the FwdYes bit of the Flags field is false, the contents of this field are undefined, and the field is ignored by the receiver.
RelayAction	Indicates the value returned in the Relay Action field.
SrcMac	Displays the MAC address of the MP that responded to the LTM request for this LTR.
IngressAction	Displays the value returned in the Ingress Action Field of the LTM. The value ingNoTlv indicates that no Reply Ingress TLV was returned in the LTM.
IngressMac	Displays the MAC address returned in the ingress MAC address field. If the rcCfmLtrReplyIngress object contains the value ingNoTlv(5), then the contents of this field are meaningless.
EgressAction	Displays the value returned in the Egress Action Field of the LTM. The value egrNoTlv(5) indicates that no Reply Egress TLV was returned in the LTM.
EgressMac	Displays the MAC address returned in the egress MAC address field. If the rcCfmLtrReplyEgress object contains the value egrNoTlv(5), then the contents of this field are meaningless.