Shortest Path Bridging for MAC

In This Chapter

This section describes advanced shortest path bridging for MAC configurations.

Topics in this section include:

Applicability

The example presented in this section is applicable to the 7950 XRS, 7750 SR-c4/c12, 7750 SR-7/12 and 7450 ESS-6/6v/7/12, and requires IOM3-XP/IMM or higher-based line cards. It is not supported in 7750 SR-1, 7450 ESS-1, 7710 SR or IOM-2 or lower-based line cards.

The configuration was tested on release 11.0R4. Shortest Path Bridging for MAC (SPBM) is supported from 10.0R4. SPB Static MAC, static I-service Instance Identifiers (ISIDs) and ISID-policies for SPB are supported from 11.0R4 onwards.

Overview

SPB enables a next generation control plane for Provider Backbone Bridges (PBB) and PBB-VPLS that adds the stability and efficiency of link state to unicast and multicast services (I-VPLS and Epipes). In addition, SPBM provides resiliency, load-balancing and multicast optimization without the need for any other control plane in the B-VPLS (for example, there is no need for Spanning Tree or G.8032 or MMRP).

SPBM exploits the complete knowledge of backbone addressing, which is a key consequence of the PBB hierarchy, by advertising and distributing the Backbone MACs (BMACs) through a link-state protocol, namely IS-IS. An immediate effect of this is that the old "flood-and-learn" can at last be turned off in the backbone and every B-VPLS node in the network will know what destination BMAC addresses are expected and valid. As a result of that, receiving an unknown unicast BMAC on a B-VPLS SAP/PW is indicative of an error, whereupon the frame is discarded (due to the Reverse Path Forwarding Check – RPFC – performed in SPBM) instead of flooded. Furthermore, SPBM allows condensing all the relevant information distribution (unicast and multicast) into a single control protocol: IS-IS.

SPBM can be easily enabled on the existing B-VPLS instances being used for multiplexing I-VPLS/Epipe services, providing the benefits summarized below:

- Per-service flood containment (for I-VPLS services) without the need for an additional protocol such as MMRP.
- Loop avoidance in the B-VPLS domain without the need for MSTP or other technologies.
- No unknown BMAC flooding in the B-VPLS domain.
- No need for MAC notification mechanisms or vMEPs in the B-VPLS to update the B-VPLS Forwarding Data Bases (FDBs) (vMEPs can still be configured though for OAM purposes).

Some other characteristics of the SPB implementation in the SR OS are listed below:

- The SR OS SPB implementation always uses Multi-Topology (MT) topology instance zero. However, up to four logical instances (that is, SPB instances in different B-VPLS services) are supported if different topologies are required for different services.
- Area addresses are not used and SPB is assumed to be a single area. SPB must be consistently configured on nodes in the system. SPB Regions information and IS-IS hello logic that detect mismatched configuration are not supported. IS-IS area is always zero.
- SPB uses All-Intermediate-Systems 09-00-2B-00-00-05 destination MAC to communicate.
- SPB Source ID is always zero.
- SPB uses a separate instance of IS-IS from the base IP IS-IS. IS-IS for SPB is configured in the SPB context under the B-VPLS component. Up to four ISIS-SPB instances are

supported, where the instance identifier can be any number between 1024 and 2047. The instance number is not in TLVs.

- Two ECT-ALGORITHMs (IEEE 802.1aq Equal Cost Tree Algorithms) per SPB instance are supported: low-path-id and high-path-id algorithms.
- SPB Link State Protocol Data Units (Link State Packets) contain BMACs, ISIDs (for multicast services) and link and metric information for an IS-IS database.
  - → Epipe ISIDs are not distributed in SR OS SPB allowing high scalability of PBB Epipes.
  - → I-VPLS ISIDs are distributed in SR OS SPB and the respective multicast group addresses (composed of PBB-OUI plus ISID) are automatically populated in a manner that provides automatic pruning of multicast to the subset of the multicast tree that supports an I-VPLS with a common ISID. This replaces the function of MMRP and is more efficient than MMRP.
- Multiple ISIS-SPB adjacencies between two nodes are not supported as per the IEEE 802.1aq standard specification. If multiple links between two nodes exist, LAG must be used.

Configuration

This section will describe the configuration of SPBM on the 7x50 as well as the available troubleshooting commands.

Basic SPBM Configuration

Figure 82 shows the topology used as an example of a basic SPBM configuration.

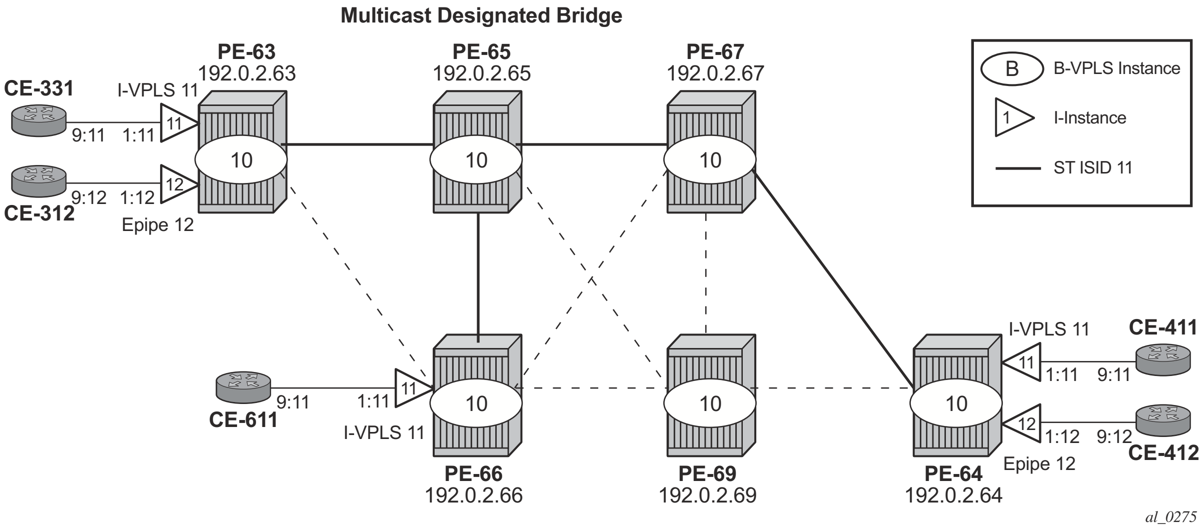

Figure 82: Basic SPBM Topology

Assume the following protocols and objects are configured beforehand:

- The six PEs shown in Figure 82 are running IS-IS for the global routing table with all the interfaces being Level-2.
- LDP is used as the MPLS protocol to signal transport tunnel labels.
- LDP SDPs are configured among the six PEs, in the way it is depicted in Figure 82 (dashed lines and bold lines among PEs).

Once the network infrastructure is properly running, the actual service configuration can be carried out. In the example, B-VPLS 10 will provide backbone connectivity for the services I-VPLS 11 and Epipe 12.

The SPBM configuration is only relevant to the B-VPLS instance and can be added to an existing B-VPLS, assuming that such a B-VPLS does not contain any non-SPB-compatible configuration parameters. The following parameters are not supported in SPB-enabled B-VPLS instances:

- mesh-sdps (only SAPs or spoke-sdps are supported in SPB-enabled B-VPLS)
- Spanning Tree Protocol (STP)
- Split-Horizon Groups
- Non-conditional Static-macs (configured under sap/spoke-sdps, see section about static BMAC configuration)
- G.8032
- Propagate-mac-flush and send-flush-on-failure
- Maximum number of macs (max-nbr-mac-addr)
- Bridge Protocol Data Unit (BPDU) translation
- Layer 2 Protocol Termination (L2PT)
- Mac-pinning
- Oper-groups
- Mac-move
- Any BGP, BGP-AD (BGP Autodiscovery) or BGP-VPLS (BGP Virtual Private LAN Services) parameters
- Endpoints
- Local/remote age
- Mac-notification
- Mac-protect
- Mesh-sdps
- Multiple MAC Registration Protocol (MMRP)
- Provider-tunnel
- Temporary flooding

Assuming all the parameters mentioned above are not configured in the B-VPLS (B-VPLS 10 in the example), SPBM can be enabled. The SPBM parameters are all configured under the **config>service>vpls(b-vpls)>spb** and **config>service>vpls(b-vpls)>spoke-sdp/sap>spb** contexts:

```
*A:PE-63# configure service vpls 10 spb ?
  - no spb
  - spb [<isis-instance>] [fid <fid>] [create]
 <isis-instance>      : [1024..2047]
 <fid>                : [1..4095]
      level           + Configure SPB level information
 [no] lsp-lifetime    - Configure LSP lifetime
 [no] lsp-wait        - Configure ISIS LSP wait times
 [no] overload        - Configure the local router so that it appears to be overloaded
 [no] overload-on-bo* - Configure the local router so that it appears to be overloaded at
boot up
 [no] shutdown        - Administratively enable or disable the operation of ISIS
 [no] spf-wait        - Configure ISIS SPF wait times

*A:PE-63# configure service vpls 10 spoke-sdp 35:10 spb ?
  - no spb
  - spb [create]
 <create>             : keyword
      level           + Configure SPB level information
 [no] lsp-pacing-int* - Configure the interval between LSP packets are sent from the inter-
face
 [no] retransmit-int* - Configure the minimum interval between LSP packets retransmission
for the given interface
 [no] shutdown        - Administratively Enable/disable the interface

*A:PE-63# configure service vpls 10 spoke-sdp 35:10 spb level 1 ?
  - level <[1..1]>
 [no] hello-interval  - Configure hello-interval for this interface
 [no] hello-multipli* - Configure hello-multiplier for this level
 [no] metric          - Configure IS-IS interface metric for IPv4 unicast
```

The parameters configured under the spb context refer to the SPB IS-IS and they should be configured following the same considerations as for the IS-IS base instance:

- spb [<isis-instance>] [fid <fid>] [create]

  → <isis-instance> identifies the SPB ISIS process. Up to four different ISIS SPB processes can be run in a system (range 1024 to 2047).

  → <fid> or *forwarding identifier* identifies the standard SPBM B-VID which is signaled in IS-IS with each advertised BMAC. Each B-VPLS has a single configurable FID.

- spb>lsp-lifetime <seconds>   : [350..65535]
- spb>lsp-wait <lsp-wait> [<lsp-initial-wait> [<lsp-second-wait>]]

  → <lsp-wait>   : [1..120] – seconds

  → <initial-wait>   : [0..100] – seconds

  → <second-wait>   : [1..100] – seconds

- spb>overload
- spb>overload-on-boot
- spb>spf-wait <spf-wait> [<spf-initial-wait> [<spf-second-wait>]]

  → <spf-wait>   : [1..120] – seconds

  → <initial-wait>   : [10..100000] – milliseconds

  → <second-wait>   : [1..100000] - milliseconds

- spoke-sdp/sap>spb>lsp-pacing-interval <milli-seconds> : [0..65535]
  - → spoke-sdp/sap>spb>retransmit-interval <seconds> : [1..65535]
  - → spoke-sdp/sap>spb>level 1>hello-interval <seconds> : [1..20000]
  - → spoke-sdp/sap>spb>level 1>hello-multiplier <multiplier> : [2..100]

In the same way lsp-wait (initial-wait) and spf-wait (initial wait) can be tuned in the base router IS-IS instance to minimize the convergence time (to 0 and 10 respectively), the equivalent SPB IS-IS parameters should also be adjusted so that failover time is minimized at the service level.

The following parameters are specific to SPBM (note that only IS-IS level 1 is supported for SPB):

- spb>level 1>bridge-priority <bridge-priority> : [0..15]
  - → This parameter will influence the election of the Multicast Designated bridge through which all the ST (Single Trees) for the multicast traffic will be established. The default value will be lowered on that node where the Multicast Designated bridge function is desired, normally because that node is the best connected node. Note that in the example of this document, PE-65 is the Multicast Designated bridge for B-VPLS 10 and therefore, PE-65 will be the root of the STs for the I-VPLS instances in that B-VPLS. Default value = 8.
- spb>level 1>ect-algorithm fid-range <fid-range> {low-path-id|high-path-id}
  - → This command defines the ect-algorithm used and the FIDs assigned. Two algorithms are supported: low-path-id and high-path-id. They can provide the required path diversity for an efficient load balancing in the B-VPLS. Default = fid-range 1-4095 low-path-id
- spb>level 1>forwarding-tree-topology unicast {spf|st}
  - → This command configures the type of tree that will be used for unicast traffic: shortest path tree or single tree. The multicast traffic (that encapsulated I-VPLS broadcast, unicast and multicast (BUM) traffic always uses the ST path (Single Tree path). Note that using SPF for unicast traffic can produce some packet re-ordering for unicast traffic compared to BUM traffic as different trees are used, therefore, when the B-VPLS transports I-VPLS traffic and the unicast and multicast trees do not follow the same path, it is recommended to use ST paths for unicast and multicast. Default value = spf.
- spoke-sdp/sap>spb>level 1>metric <ipv4-metric> : [1..16777215]
  - → This command configures the metric for each SPB interface (spoke-SDP or SAP). This value helps influence the SPF calculation in order to pick a certain path for the traffic to a remote system BMAC. Note that when the SPB link metric advertised by two peers is different, the maximum value is chosen according to the RFC 6329. Default metric = 10.

As an example, the following CLI output shows the relevant configuration of PE-63 and PE-65 (the Multicast Designated Bridge). SPB has to be created and enabled (no shutdown) at B-VPLS service level first and then created and enabled under each and every SAP/spoke-sdp in the B-VPLS. Non-SPB-enabled SAPs/spoke-sdps can exist in the SPB B-VPLS only if conditional static-macs are configured for them (refer to Static BMACs and Static ISIDs Configuration on page 579). Note that, as for regular B-VPLS services, the service-mtu has to be changed from the default value (1500) to a number 18-bytes greater than the I-VPLS service-mtu in order to allow for the PBB encapsulation.

```
*A:PE-63>config>service# info
----------------------------------------------
        pbb
            source-bmac 00:00:00:ca:fe:63
            mac-name "PE-63" 00:00:00:ca:fe:63
            mac-name "PE-64" 00:00:00:ca:fe:64
            mac-name "PE-65" 00:00:00:ca:fe:65
            mac-name "PE-66" 00:00:00:ca:fe:66
            mac-name "PE-67" 00:00:00:ca:fe:67
            mac-name "PE-69" 00:00:00:ca:fe:69
        exit
...
        vpls 10 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1024 fid 10 create
                overload-on-boot timeout 60
                spf-wait 2 50 100
                lsp-wait 8 0 1
                no shutdown
            exit
            spoke-sdp 35:10 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 36:10 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
        vpls 11 customer 1 i-vpls create
            pbb
                backbone-vpls 10
                exit
            exit
            stp
                shutdown
            exit
            sap 1/1/1:11 create
            exit
            no shutdown
        exit
        epipe 12 customer 1 create
            pbb
                tunnel 10 backbone-dest-mac "PE-64" isid 12
            exit
            sap 1/1/1:12 create
            exit
            no shutdown
        exit
```

As discussed, the *bridge-priority* will influence the election of the Multicast Designated Bridge. By making PE-65's bridge-priority zero, it ensures that PE-65 becomes the root of all the STs for B-VPLS 10 as long as the priority for the rest of the PEs is larger than zero. In case of a tie, the PE owning the lowest system BMAC will be elected as Multicast Designated Bridge. Figure 82 shows the ST for I-VPLS 11 (see a thicker continuous line representing the ST). Note that PE-65 is the root of the ST tree.

```
*A:PE-65>config>service# info
----------------------------------------------
        pbb
            source-bmac 00:00:00:ca:fe:65
            mac-name "PE-63" 00:00:00:ca:fe:63
            mac-name "PE-64" 00:00:00:ca:fe:64
            mac-name "PE-65" 00:00:00:ca:fe:65
            mac-name "PE-66" 00:00:00:ca:fe:66
            mac-name "PE-67" 00:00:00:ca:fe:67
            mac-name "PE-69" 00:00:00:ca:fe:69
        exit
...
        vpls 10 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1024 fid 10 create
                level 1
                    bridge-priority 0
                exit
                overload-on-boot timeout 60
                spf-wait 2 50 100
                lsp-wait 8 0 1
                no shutdown
            exit
            spoke-sdp 53:10 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 56:10 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 57:10 create
```

```
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 59:10 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
```

The rest of the nodes will be configured accordingly. Note that SPB instance 1024 will set up SPF (Shortest Path First) trees for unicast traffic and a ST (Single Tree) per ISID with PE-65 as the root-bridge (since it has the lowest bridge-priority) for BUM traffic. The ect-algorithm chosen for the B-VPLS FID (10) is the low-path-id (default one).

Once SPBM is configured as discussed above on all the six nodes, the six system BMACs and the ISID 11 will be advertised by SPB IS-IS.

The following show commands can help understand the IS-IS configuration for SPB 1024 and the BMACs populated by IS-IS:

- show service id spb base: provides the SPB configuration and parameters for a particular SPB B-VPLS.
- show service id 10 spb fdb: provides the B-VPLS FDB that has been populated by IS-IS, for the unicast and multicast entries.

```
*A:PE-63# show service id 10 spb base
===============================================================================
Service SPB Information
===============================================================================
Admin State        : Up                  Oper State         : Up
ISIS Instance      : 1024                FID                : 10
Bridge Priority    : 8                   Fwd Tree Top Ucast : spf
Fwd Tree Top Mcast : st
Bridge Id          : 80:00.00:00:00:ca:fe:63
Mcast Desig Bridge : 00:00.00:00:00:ca:fe:65
===============================================================================
ISIS Interfaces
===============================================================================
Interface                        Level CircID  Oper State   L1/L2 Metric
-------------------------------------------------------------------------------
sdp:35:10                        L1    65536   Up           10/-
sdp:36:10                        L1    65537   Up           10/-
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
FID ranges using ECT Algorithm
-------------------------------------------------------------------------------
1-4095    low-path-id
===============================================================================
*A:PE-63# show service id 10 spb fdb
==============================================================================
User service FDB information
==============================================================================
MAC Addr          UCast Source          State   MCast Source          State
------------------------------------------------------------------------------
00:00:00:ca:fe:64 35:10                 ok      35:10                 ok
00:00:00:ca:fe:65 35:10                 ok      35:10                 ok
00:00:00:ca:fe:66 36:10                 ok      35:10                 ok
00:00:00:ca:fe:67 35:10                 ok      35:10                 ok
00:00:00:ca:fe:69 35:10                 ok      35:10                 ok
------------------------------------------------------------------------------
Entries found: 5
==============================================================================
```

It can be seen above that the unicast (SPF) tree and the multicast (ST) tree differ with respect to PE-66.

The following commands help check the unicast and multicast topology for B-VPLS 10:

- **show service id 10 spb routes** provides a detailed view of the unicast and multicast routes computed by SPF. As shown below, the SPB unicast and multicast routes match on PE-65 since this node is the Multicast Designated Bridge. Unicast and multicast routes will differ on most other nodes.
- **show service id 10 spb mfib** and **show service id 10 mfib** shows information of the MFIB entries generated in the B-VPLS as well as the outgoing interface (OIF) associated with those MFIB entries.

```
*A:PE-65# show service id 10 spb routes
================================================================
MAC Route Table
================================================================
FID  MAC Addr                                     Ver.   Metric
      NextHop If                    SysID
----------------------------------------------------------------
Fwd Tree: unicast
----------------------------------------------------------------
10   00:00:00:ca:fe:63                            3      10
      sdp:53:10                     PE-63
10   00:00:00:ca:fe:64                            9      20
      sdp:57:10                     PE-67
10   00:00:00:ca:fe:66                            4      10
      sdp:56:10                     PE-66
10   00:00:00:ca:fe:67                            8      10
      sdp:57:10                     PE-67
10   00:00:00:ca:fe:69                            12     10
      sdp:59:10                     PE-69

Fwd Tree: multicast
----------------------------------------------------------------
10   00:00:00:ca:fe:63                            3      10
      sdp:53:10                     PE-63
10   00:00:00:ca:fe:64                            9      20
      sdp:57:10                     PE-67
```

```
10   00:00:00:ca:fe:66                            15     10
      sdp:56:10                     PE-66
10   00:00:00:ca:fe:67                            8      10
      sdp:57:10                     PE-67
10   00:00:00:ca:fe:69                            12     10
      sdp:59:10                     PE-69
----------------------------------------------------------------
No. of MAC Routes: 10
================================================================
 ================================================================
ISID Route Table
================================================================
FID  ISID                                         Ver.
      NextHop If                    SysID
----------------------------------------------------------------
10   11                                           16
      sdp:53:10                     PE-63
      sdp:56:10                     PE-66
      sdp:57:10                     PE-67
----------------------------------------------------------------
No. of ISID Routes: 1
================================================================


*A:PE-65# show service id 10 spb mfib
===============================================================================
User service MFIB information
===============================================================================
MAC Addr          ISID     Status
-------------------------------------------------------------------------------
01:1E:83:00:00:0B 11       Ok
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
A:PE-65# show service id 10 mfib
===============================================================================
Multicast FIB, Service 10
===============================================================================
Source Address  Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               01:1E:83:00:00:0B     b-sdp:53:10              Local    Fwd
                                      b-sdp:56:10              Local    Fwd
                                      b-sdp:59:10              Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

SPB Multicast Trees (STs) are pruned for each particular I-VPLS ISID, based on the advertisement of I-VPLS ISIDs in SPB IS-IS by each individual PE. Multicast B-VPLS traffic not belonging to any particular I-VPLS follows the *Default Tree*. The Default Tree is an ST for the B-VPLS which is not pruned and therefore reaches all the PE nodes in the B-VPLS. For instance, Ethernet-CFM CCM messages sent from vMEPs configured on the SPB B-VPLS will use the default tree. The default tree does not consume MFIB entries and can be checked in each node through the use of the following command:

```
*A:PE-69# tools dump service id 10 spb default-multicast-list
saps : { }
spoke-sdps : { 95:10 }
```

As it can be noted above, PE-69 is not part of the tree for I-VPLS 11. However, as with any SPB node part of B-VPLS 10, PE-69 is part of the default tree. Refer to Configuration of ISID-Policies in SPB B-VPLS on page 584 to see more use-cases for the Default Tree.

The following tools commands allow the operator to easily see the forwarding path (unicast and multicast) followed by the traffic to a remote node, with the aggregate metric from the source.

```
*A:PE-63# tools dump service id 10 spb fid 10 forwarding-path destination PE-64 forward-
ing-tree unicast

Hop  BridgeId                 Metric From Src
0    PE-63                     0
1    PE-65                     10
2    PE-67                     20
3    PE-64                     30

*A:PE-63# tools dump service id 10 spb fid 10 forwarding-path destination PE-64 forward-
ing-tree multicast

Hop  BridgeId                 Metric From Src
0    PE-63                     0
1    PE-65                     10
2    PE-67                     20
3    PE-64                     30
```

In large networks or networks where IP multicast, PBB and PBB-SPB services coexist, the data plane MFIB entries is a hardware resource that should be periodically checked. The tools dump service vpls-mfib-stats shows the total number of hardware MFIB entries (40k entries, in chassis-mode d) and the entries being used by IP multicast or PBB (MMRP or SPB). The tools dump service vpls-pbb-mfib-stats shows the breakdown between MFIB entries populated by MMRP or by SPB and the individual limits, system-wide and per service:

```
*A:PE-65# tools dump service vpls-mfib-stats
Service Manager VPLS MFIB info at 004 09:44:57.010:

Statistics last cleared at 000 00:00:00.000

            Statistic            |    Count
```

```
---------------------------------+-------------
             HW limit SG entries |      40959# total number of MFIB entries
              Current SG entries |          2
        Limit Non PBB SG entries |      16383# IP Multicast MFIB limit
      Current Non PBB SG entries |          0
...

*A:PE-65# tools dump service vpls-pbb-mfib-stats detail
Service Manager VPLS PBB MFIB statistics at 004 09:45:13.860:

Usage per Service
   ServiceId    MFIB User      Count
  ------------+--------------+-------
   10           spb                1
   30           spb                1
  ------------+--------------+-------
                       Total       2

MMRP
  Current Usage     :     0
  System Limit      :  8191 Full, 40959 ESOnly
  Per Service Limit :  2048 Full,  8192 ESOnly

SPB
  Current Usage     :     2
  System Limit      :  8191
  Per Service Limit :  8191
```

Finally the following debug commands can help monitor the SPB IS-IS process and the protocol PDU exchanges:

- debug service id <svcId> spb
- debug service id <svcId> spb adajacency
- debug service id <svcId> spb interface
- debug service id <svcId> spb l2db
- debug service id <svcId> spb lsdb
- debug service id <svcId> spb packet <detail>
- debug service id <svcId> spb spf

Control and User B-VPLS Configuration

The SR OS implementation of SPB allows a single SPB IS-IS instance to control the paths and FDBs of many B-VPLS instances. This is done by using the control B-VPLS, user B-VPLS and fate-sharing concepts.

The control B-VPLS will be SPB-enabled and configured with all the related SPB IS-IS parameters. Although the control B-VPLS might or might not have I-VPLS/Epipes directly attached, it must be configured on all the nodes where SPB forwarding is expected to be active. SPB uses the logical instance and a Forwarding ID (FID) to identify SPB locally on the node. That FID must be consistently configured on all the nodes where the B-VPLS exists. User B-VPLS are other instances of B-VPLS that are usually configured to separate the traffic for manageability reasons, QoS or ECT different treatment.

Figure 83 illustrates the control B-VPLS (B-VPLS 20) and user B-VPLS (B-VPLS 21) concept (in this case there is only one user B-VPLS but there might be many B-VPLS sharing fate with the same control B-VPLS). Note that both B-VPLS must share the same topology and both B-VPLS must share exactly the same interfaces. The user B-VPLS, which is linked to the control B-VPLS by its FID, follows (that is, inherits the state of) the control B-VPLS but may use a different ECT path in case of equal metric paths, like in this example: FID 20, that is, the control B-VPLS, follows the low-path-id ECT, whereas FID 21, for example, the user B-VPLS, follows the high-path-id ECT.

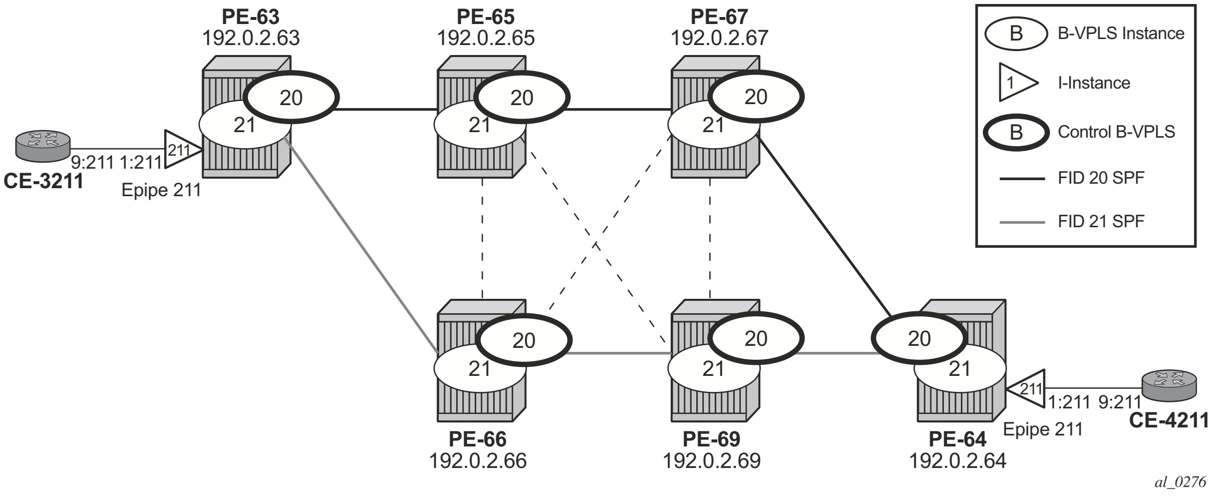

Figure 83: Control and User B-VPLS Test Topology

The configurations of B-VPLSs 20 and 21, on PE-63 and PE-65, are shown below. The **spbm-control-vpls 20 fid 21 in B-VPLS 21** command associates FID 21 to the user B-VPLS and links the B-VPLS to its control B-VPLS 20.

```
*A:PE-63>config>service# info
----------------------------------------------
...
        vpls 20 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1025 fid 20 create
                level 1
                    ect-algorithm fid-range 21-4095 high-path-id
                exit
                no shutdown
            exit
            spoke-sdp 35:20 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 36:20 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
        vpls 21 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spbm-control-vpls 20 fid 21
            spoke-sdp 35:21 create
                no shutdown
            exit
            spoke-sdp 36:21 create
                no shutdown
            exit
            no shutdown
        exit
        epipe 211 customer 1 create
            pbb
                tunnel 21 backbone-dest-mac "PE-64" isid 211
            exit
            sap 1/1/1:211 create
            exit
            no shutdown
        exit
...

*A:PE-65>config>service# info
----------------------------------------------
```

```
...
        vpls 20 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1025 fid 20 create
                level 1
                    ect-algorithm fid-range 21-4095 high-path-id
                exit
                no shutdown
            exit
            spoke-sdp 53:20 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 56:20 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 57:20 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 59:20 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
        vpls 21 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spbm-control-vpls 20 fid 21
            spoke-sdp 53:21 create
                no shutdown
            exit
            spoke-sdp 56:21 create
                no shutdown
            exit
            spoke-sdp 57:21 create
                no shutdown
            exit
            spoke-sdp 59:21 create
                no shutdown
            exit
            no shutdown
        exit
```

If there is a mismatch between the topology of a user B-VPLS and its control B-VPLS, only the user B-VPLS links and nodes that are in common with the control B-VPLS will function.

User B-VPLS instances supporting only unicast services (PBB-Epipes) may share the FID with the other B-VPLS (control or user). This is a configuration shortcut that reduces the LSP advertisement size for B-VPLS services but results in the same separation for forwarding between the B-VPLS services. In the case of PBB-Epipes only BMACs are advertised per FID but BMACs are populated per B-VPLS in the FIB. If I-VPLS services are to be supported on a B-VPLS that B-VPLS must have an independent FID.

Note that although user B-VPLS 21 does not have any SPB setting (other than the spbm-control-vpls) the spoke-sdps use the same SDPs as the parent control B-VPLS 20. The **show service id <user b-vpls> spb fate-sharing** command shows the control spoke-sdp/saps that control the user spoke-sdp/saps.

```
*A:PE-63# show service id 21 spb fate-sharing
===============================================================================
User service fate-shared sap/sdp-bind information
===============================================================================
Control   Control Sap/             FID       User      User Sap/
SvcId     SdpBind                            SvcId     SdpBind
-------------------------------------------------------------------------------
20        35:20                    21        21        35:21
20        36:20                    21        21        36:21
===============================================================================
```

SPBM Access Resiliency Configuration

The following example shows how to configure an I-VPLS/Epipe attached to an SPB-enabled B-VPLS when access resiliency is used.

In the tested SR OS release, multi-chassis LAG (MC-LAG) is the only resiliency mechanism supported for PBB-Epipes. The MC-LAG active node will advertise the MC-LAG BMAC (or sap-bmac) in SPB IS-IS. In case of failure, when the standby node takes over it will advertise the MC-LAG sap-bmac. Note that without SPB, the MC-LAG solution for PBB-Epipe required the use of mac-notification and periodic mac-notification. SPB provides a faster and more efficient solution without the need for any extra mac-notification mechanism. In the example described in this section, Epipe 31 uses MC-LAG access resiliency to get connected to the B-VPLS 30 on nodes PE-65 and PE-66.

As far as I-VPLS access resiliency is concerned, the same mechanisms supported for regular B-VPLS are supported for SPB-enabled B-VPLS, except for G.8032, which is not supported in the tested SR OS release. A very important aspect of the I-VPLS resiliency is a proper mac-flush propagation when there is a failure at the I-VPLS access links.

If the SPB-enabled B-VPLS uses B-SAPs for its connectivity to the backbone, there is no such a mac-flush propagation (since there is no TLDP). In this case, if MC-LAG is used and there is an MC-LAG switchover, the new active chassis will keep using the same source BMAC, such as the sap-bmac, and it will advertise it in the B-VPLS domain so that the remote FDBs can be properly updated. No mac-flush is required in this case.

When the B-VPLS uses spoke-sdps for its backbone connectivity, the traditional LDP MAC flush propagation mechanisms and commands can be used as follows:

- **send-flush-on-failure** works as expected when SPB is used at the B-VPLS. When configured, a flush-all-from-me event is triggered upon a SAP or spoke-sdp failure in the I-VPLS.
- **send-bvpls-flush** works as expected when SPB is used at the B-VPLS. Two variants are configurable: all-from-me/all-but-mine. Any I-VPLS SAP/spoke-sdp failure is propagated to the I-VPLS on the peers to flush their respective customer MACs (CMACs). It works only in conjunction with send-flush-on-failure configuration on I-VPLS. The associated ISID list is passed along with the LDP mac-flush message, which is flushed/retained according to the **all-from-me/all-but-me** flag.
- **send-flush-on-bvpls-failure** works as expected when SPB is used at the B-VPLS. A local B-VPLS failure is propagated to the I-VPLS, which then triggers a LDP mac-flush if it has any spoke-sdp on it.
- **propagate-mac-flush-from-bvpls** does no work when SPB is used at the B-VPLS (since failures within the B-VPLS are handled by SPB) and its configuration is blocked.

In the example described later in this section, I-VPLS 32 uses active/standby spoke-sdp resiliency to get connected to the B-VPLS 30 on nodes PE-67 and PE-69.

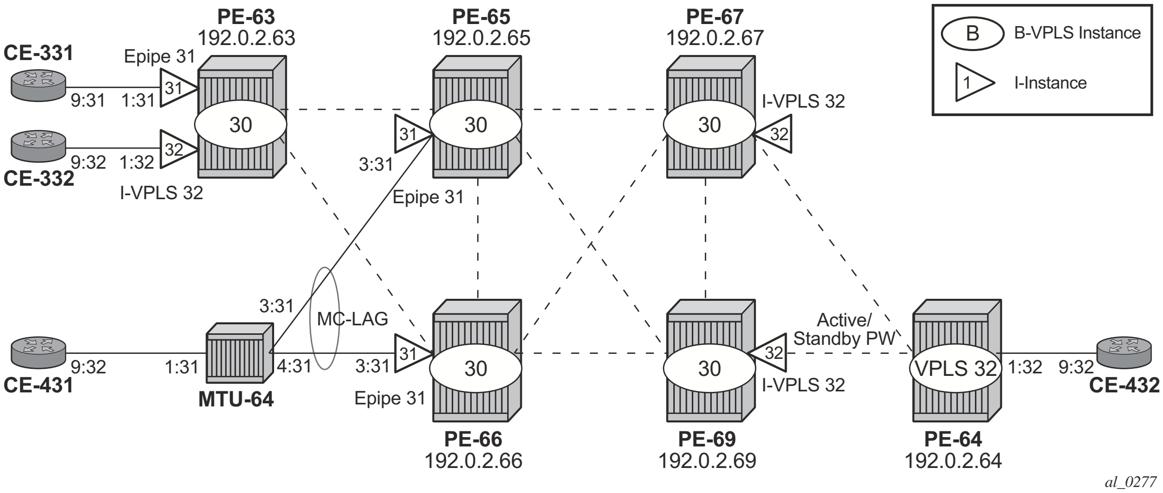

Figure 84: Access Resiliency Test Topology

As an example of MC-LAG connectivity, the Epipe 31 configuration is shown below. Just like for regular PBB-VPLS, a sap-bmac is used as source BMAC for the Epipe traffic from PE-65/66 to PE-63. A sap-bmac is a virtual BMAC formed from the configured source-bmac plus the MC-LAG LACP-KEY (if configured this way) and owned by the MC-LAG active chassis. The following CLI output shows the configuration of MC-LAG as well as the generation of the sap-bmac. Once it is properly configured and the MC-LAG and Epipe are up and running, SPB IS-IS will distribute the sap-bmac throughout the B-VPLS, as it does for the system BMACs and OAM vMEP MACs. In this example, PE-65 is the MC-LAG active node, hence the sap-bmac for Epipe 31 is generated from PE-65.

```
*A:PE-65>config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.0.2.66 create
                mc-lag
                    lag 1 lacp-key 1 system-id 00:00:00:00:65:66
                   system-priority 65535 source-bmac-lsb use-lacp-key
                    no shutdown
                exit
                no shutdown
            exit
        exit
----------------------------------------------

*A:PE-65>config>service# info
----------------------------------------------
        pbb
```

```
            source-bmac 00:00:00:ca:fe:65
            mac-name "PE-63" 00:00:00:ca:fe:63
            mac-name "PE-64" 00:00:00:ca:fe:64
            mac-name "PE-65" 00:00:00:ca:fe:65
            mac-name "PE-66" 00:00:00:ca:fe:66
            mac-name "PE-67" 00:00:00:ca:fe:67
            mac-name "PE-69" 00:00:00:ca:fe:69
        exit
...
        vpls 30 customer 1 b-vpls create
            service-mtu 2000
            pbb
                use-sap-bmac
            exit
            stp
                shutdown
            exit
            spb 1026 fid 30 create
                level 1
                    bridge-priority 0
                exit
                no shutdown
            exit
            spoke-sdp 53:30 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 56:30 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 57:30 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 59:30 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
        epipe 31 customer 1 create
            pbb
                tunnel 30 backbone-dest-mac "PE-63" isid 31
            exit
            sap lag-1:31 create
            exit
            no shutdown
        exit
...
```

```
*A:PE-66# show service id 30 spb fdb
===============================================================================
User service FDB information
===============================================================================
MAC Addr          UCast Source          State   MCast Source          State
-------------------------------------------------------------------------------
00:00:00:ca:00:01 65:30                 ok      65:30                 ok
00:00:00:ca:fe:63 63:30                 ok      65:30                 ok
00:00:00:ca:fe:65 65:30                 ok      65:30                 ok
00:00:00:ca:fe:67 67:30                 ok      65:30                 ok
00:00:00:ca:fe:69 69:30                 ok      65:30                 ok
-------------------------------------------------------------------------------
Entries found: 5
===============================================================================
```

The configuration for I-VPLS 32 on nodes PE-64 and PE-67 is shown below.

```
A:PE-64>config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            endpoint "CORE" create
                no suppress-standby-signaling
            exit
            sap 1/1/1:32 create
            exit
            spoke-sdp 47:32 endpoint "CORE" create
                stp
                    shutdown
                exit
                precedence primary
                no shutdown
            exit
            spoke-sdp 49:32 endpoint "CORE" create
                stp
                    shutdown
                exit
                no shutdown
            exit
            no shutdown
----------------------------------------------
*A:PE-67>config>service# info
----------------------------------------------
...
        vpls 30 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1026 fid 30 create
                no shutdown
            exit
            spoke-sdp 75:30 create
                spb create
                    no shutdown
                exit
                no shutdown
```

```
    exit
    spoke-sdp 76:30 create
        spb create
            no shutdown
        exit
        no shutdown
    exit
    spoke-sdp 79:30 create
        spb create
            no shutdown
        exit
        no shutdown
    exit
    no shutdown
exit
vpls 32 customer 1 i-vpls create
    send-flush-on-failure
    pbb
        backbone-vpls 30
        exit
        send-bvpls-flush all-from-me
 send-flush-on-bvpls-failure
    exit
    stp
        shutdown
    exit
    spoke-sdp 74:32 create
        no shutdown
    exit
    no shutdown
exit
```

As discussed, **send-flush-on-failure** and **send-bvpls-flush all-from-me** are configured in the I-VPLS. When the active spoke-sdp goes down on PE-67, a flush-all-from-me message will be propagated through the backbone and will flush the corresponding CMACs associated to the I-VPLS 32 in node PE-63. MAC flush-all-from-me messages are automatically propagated in the core up to the remote I-VPLS 32 on node PE-63 (there is no need for any propagate-mac-flush in the intermediate nodes). Note that the *send-flush-on-bvpls-failure* command works as expected. The command *propagate-mac-flush-from-bvpls* is never used when the B-VPLS is SPB-enabled (the command is not allowed).

Static BMACs and Static ISIDs Configuration

From 11.0R4 onwards, SR OS supports the interworking between SPB-enabled B-VPLS and non-SPB B-VPLS instances. With the addition of this feature, SPB networks can be connected to non-SPB capable nodes, for example third party vendor PBB switches or 7210 SAS nodes. This is possible through the use of conditional static BMACs and static ISIDs on the nodes doing the interworking function. Conditional static BMACs and static ISIDs can be associated to non-SPB B-VPLS SAPs or spoke-sdps.

The following example shows an SPB-enabled B-VPLS (40) on nodes PE-65, PE-66, PE-67 and PE-69. Node PE-63 supports PBB, but not SPB and it is connected by a MC-LAG to nodes PE-67 and PE-69. Services I-VPLS 41 and Epipe 42 have end-points on node PE-63. In this example, nodes PE-67 and PE-69 are acting as interworking nodes. They will be configured with the BMAC of PE-63 so that the MC-LAG active node advertises the non-SPB capable node BMAC into SPB IS-IS. The BMAC will be configured as a conditional static BMAC so that a given SPB node, such as PE-67 or 69, will only advertise PE-63’s BMAC if its connection to PE-63 is active. Besides the conditional static BMAC, nodes PE-67/69 should advertise the I-VPLS ISIDs defined in PE-63. Note that Epipe ISIDs are not advertised in SPB IS-IS, therefore it is not necessary to create a static ISID for Epipe 42.

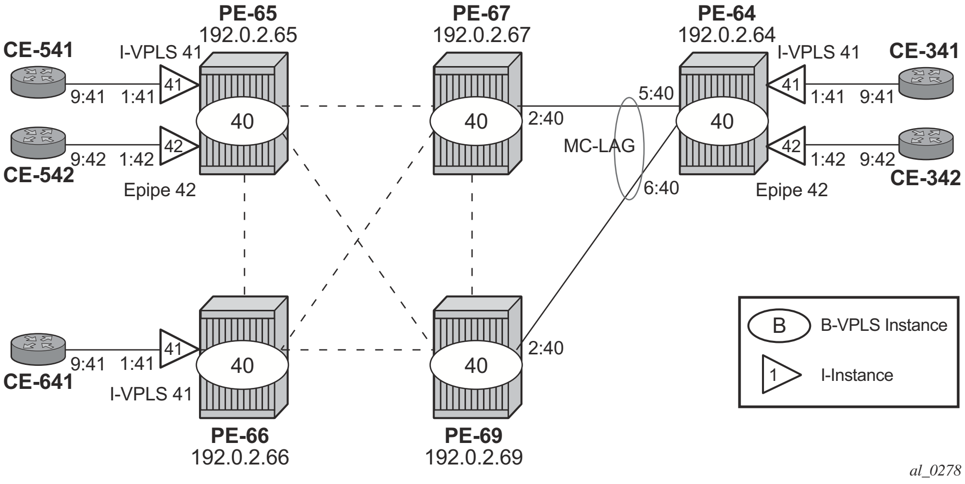

Figure 85: Access Resiliency Test Topology

The commands to configure conditional static BMACs and static ISIDs are shown below.

```
*A:PE-67# configure service vpls 40 static-mac mac  ?
  - mac <ieee-address> [create] sap <sap-id> [monitor {fwd-status}]
  - no mac <ieee-address>
  - mac <ieee-address> [create] spoke-sdp <sdp-id:vc-id> [monitor {fwd-status}]
```

```
*A:PE-67# configure service vpls 40 sap lag-1:40 static-isid range
  - no range <range-id>
  - range <range-id> isid <isid-value> [to <isid-value>] [create]

 <range-id>           : [1..8191]
 <isid-value>         : [1..16777215]
 <create>             : keyword
```

Note that the **monitor fwd-status** attribute identifies this to be a conditional MAC and is mandatory for static BMACs. This parameter instructs the 7x50 to advertise the BMAC only if the corresponding SAP/spoke-sdp is in forwarding state.

The configuration of the conditional static BMAC and static ISID is shown below.

```
*A:PE-67>config>service # info
----------------------------------------------
...
       vpls 40 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1027 fid 40 create
                spf-wait 10 10 1000
                no shutdown
            exit
            sap lag-1:40 create
                static-isid
                    range 1 create isid 41
                exit
            exit
            spoke-sdp 75:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 76:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 79:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            static-mac
                mac 00:00:00:ca:fe:63 create sap lag-1:40 monitor fwd-status
            exit
            no shutdown
        exit
----------------------------------------------
*A:PE-69>config>service# info
...
```

```
        vpls 40 customer 1 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1027 fid 40 create
                spf-wait 10 10 1000
                no shutdown
            exit
            sap lag-1:40 create
                static-isid
                    range 1 create isid 41
                exit
            exit
            spoke-sdp 95:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 96:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 97:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            static-mac
                mac 00:00:00:ca:fe:63 create sap lag-1:40 monitor fwd-status
            exit
            no shutdown
        exit
----------------------------------------------
```

Bear in mind that the configuration of the conditional static BMAC is different from the legacy static-mac command, configured within the sap/sdp-binding context. The latter static-mac is not conditional and it is always added to the FDB. The conditional static BMAC is added to the FDB based on the SAP/sdp-binding state (note that the conditional static BMAC is tagged in the FDB as **CStatic**, for Conditional Static).

```
*A:PE-67# show lag 1
===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Port-Threshold   Up-Link-Count   MC Act/Stdby
-------------------------------------------------------------------------------
1              up      up      0                1               active
===============================================================================


*A:PE-67# show service id 40 fdb pbb
```

```
=======================================================================
Forwarding Database, b-Vpls Service 40
=======================================================================
MAC               Source-Identifier     iVplsMACs  Epipes     Type/Age
-----------------------------------------------------------------------
00:00:00:ca:fe:63 sap:lag-1:40          0          0          CStatic
00:00:00:ca:fe:65 sdp:75:40             0          0          Spb
00:00:00:ca:fe:66 sdp:76:40             0          0          Spb
00:00:00:ca:fe:69 sdp:79:40             0          0          Spb
=======================================================================


*A:PE-69# show lag 1
===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Port-Threshold   Up-Link-Count   MC Act/Stdby
-------------------------------------------------------------------------------
1              up      down    0                0               standby
===============================================================================


*A:PE-69# show service id 40 fdb pbb
=======================================================================
Forwarding Database, b-Vpls Service 40
=======================================================================
MAC               Source-Identifier     iVplsMACs  Epipes     Type/Age
-----------------------------------------------------------------------
00:00:00:ca:fe:63 sdp:97:40             0          0          Spb
00:00:00:ca:fe:65 sdp:95:40             0          0          Spb
00:00:00:ca:fe:66 sdp:96:40             0          0          Spb
00:00:00:ca:fe:67 sdp:97:40             0          0          Spb
=======================================================================
```

The static-isid command identifies a set of ISIDs for I-VPLS services that are external to SPBM. These ISIDs are advertised as supported locally on this node unless altered by an isid-policy. Although the example above shows the use of the static-isid associated to a MC-LAG SAP, regular SAPs or spoke SDPs are also supported. ISIDs declared in this way become part of the ISID multicast and consume MFIBs. Multiple SPBM static-isid ranges are allowed under a SAP/ spoke SDP. ISIDs are advertised as if they were attached to the local BMAC. Only remote I-VPLS ISIDs need to be defined. In the MFIB, the backbone group MACs are then associated with the active SAP or spoke SDP.

Once the conditional static BMAC for PE-63 and the static-isid 41 (for I-VPLS 41) are configured as discussed, the advertised BMAC and ISID can be checked in the remote SPB nodes:

```
*A:PE-66# show service id 40 spb fdb
==============================================================================
User service FDB information
==============================================================================
MAC Addr          UCast Source          State   MCast Source          State
------------------------------------------------------------------------------
00:00:00:ca:fe:63 67:40                 ok      65:40                 ok
00:00:00:ca:fe:65 65:40                 ok      65:40                 ok
00:00:00:ca:fe:67 67:40                 ok      65:40                 ok
```

```
00:00:00:ca:fe:69 69:40                 ok      65:40                 ok
------------------------------------------------------------------------------
Entries found: 4
==============================================================================


*A:PE-66# show service id 40 spb mfib
===============================================================================
User service MFIB information
===============================================================================
MAC Addr          ISID     Status
-------------------------------------------------------------------------------
01:1E:83:00:00:29 41       Ok
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================


*A:PE-66# show service id 40 mfib
===============================================================================
Multicast FIB, Service 40
===============================================================================
Source Address  Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               01:1E:83:00:00:29     b-sdp:65:40              Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

Note that the group address terminates in hex 29, which corresponds to ISID 41.

The configured static-isids can be displayed with the following command (a range 41-100 has been added to the sap lag-1:40 to demonstrate this output):

```
*A:PE-69># show service id 40 sap lag-1:40 static-isids
===================================
Static Isid Entries
===================================
Entry          Range
-----------------------------------
1              41-100
===================================
```

Configuration of ISID-Policies in SPB B-VPLS

ISID policies are an optional aspect of SPBM which allow additional control of the advertisement of ISIDs and creation of MFIB entries for I-VPLS (Epipe services do not trigger ISID advertisements or the creation of MFIB entries). By default, if no ISID-policies are used, SPBM automatically advertises and populates MFIB entries for I-VPLS and static-isids. ISID-policies can be used on any SPB-enabled node with locally defined I-VPLS instances or static-isids. The isid-policy parameters are shown below:

```
*A:PE-67>config>service>vpls# isid-policy entry ?
  - entry <range-entry-id> [create]
  - no entry <range-entry-id>

 <range-entry-id>     : [1..8191]
 <create>             : keyword


 [no] advertise-local - Configure local advertisement of the range
 [no] range           - Configure ISID range for the entry
 [no] use-def-mcast   - Use default multicast tree to propogate ISID range
```

Where:

- **advertise-local** defines whether the local ISIDs (I-VPLS ISIDs linked to the B-VPLS) or static ISIDs contained in the configured range are advertised in SPBM.
- **use-def-mcast** controls whether the ISIDs contained in the range use MFIB entries (if **no use-def-mcast** is used) or just the default tree which does not use any MFIB entry.

The **isid-policy** becomes active as soon as it is defined, as opposed to other policies in SR OS, which require the policy itself to be applied within the configuration.

The typical use of ISID-policies is to reduce the number of ISIDs being advertised and/or to save MFIB space (in deployments where MFIB space is shared with MMRP and IP Multicast). The use of ISID-policies is recommended for I-VPLS where most of the traffic is unicast or for I-VPLS where the ISID end-points are present in all the Backbone Edge Bridges (BEBs) of the SPB network. In both cases, advertising ISIDs or consuming MFIB entries for those I-VPLSs has little value since no multicast (first case) or the default tree (second case) are as efficient as using MFIB entries.

The following configuration example will use the test topology in Figure 85. In this case, the objective of the **isid-policy** will be to use the default tree for all the I-VPLS services with ISIDs between 41 and 100, excluding the range 80-90. The following example shows the policy configuration in PE-67. Note that the same policy will be configured in the rest of the SPB nodes, that is, PE-65, PE-66 and PE-69.

```
*A:PE-67>config>service# vpls 40
*A:PE-67>config>service>vpls# info
----------------------------------------------
            service-mtu 2000
            stp
                shutdown
            exit
            spb 1027 fid 40 create
                spf-wait 10 10 1000
                no shutdown
            exit
            isid-policy
                entry 10 create
                    range 80 to 90
                exit
                entry 20 create
                    use-def-mcast
                    no advertise-local
                    range 41 to 79
                exit
                entry 30 create
                    use-def-mcast
                    no advertise-local
                    range 91 to 100
                exit
            exit
            sap lag-1:40 create
                static-isid
                    range 1 create isid 41 to 100
                exit
            exit
            spoke-sdp 75:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 76:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            spoke-sdp 79:40 create
                spb create
                    no shutdown
                exit
                no shutdown
            exit
            static-mac
                mac 00:00:00:ca:fe:63 create sap lag-1:40 monitor fwd-status
            exit
            no shutdown
----------------------------------------------
```

Note that the **no advertise-local** option can only be configured if the **use-def-mcast** option is also configured.

```
*A:PE-67>config>service>vpls>isid-policy>entry# no advertise-local
MINOR: SVCMGR #7855 Cannot set AdvLocal for entry - advertise-local or use-def-mcast 
option must be specified
```

Overlapping ISID values can be configured as long as the actions are consistent for the same ISID. Conflicting actions are shown in the CLI.

```
*A:PE-67>config>service>vpls>isid-policy>entry# use-def-mcast
MINOR: SVCMGR #7854 Cannot set UseDefMctree for entry - Conflicting Actions with Entry-20
```

The isid-policy configured for B-VPLS 40 in all the four nodes makes the SPB network to use the default tree for ISIDs 41-79 and 91-100 and not advertise those ISIDs in SPB IS-IS even if the ISID is locally defined (as in the case for ISIDs 41-100 in PE-67). As discussed in Basic SPBM Configuration on page 559, the default tree path can be checked from each node by using the **tools dump service id 40 spb default-multicast-list** command.

Due to entry 10 in the policy, ISIDs 80-90 will be advertised by PE-67 (active MC-LAG node). However, nodes PE-65 and PE-66 will not create any MFIB entry for those ISIDs until the corresponding I-VPLS ISIDs are locally created (or configured through static-isids). The following command executed on PE-65 proves that ISIDs 80-90 are indeed being advertised by PE-67:

```
*A:PE-65# show service id 40 spb database detail
===============================================================================
ISIS Database
===============================================================================
Displaying Level 1 database
-------------------------------------------------------------------------------
LSP ID    : PE-65.00-00                                  Level     : L1
...

-------------------------------------------------------------------------------
LSP ID    : PE-66.00-00                                  Level     : L1
...

-------------------------------------------------------------------------------
LSP ID    : PE-67.00-00                                  Level     : L1
…
  MT Capability :
    TLV Len        : 56
    MT ID          : 0
    SPBM Service ID:
    Sub TLV Len    : 52
      BMac Addr              : 00:00:00:ca:fe:67
      Base VID               : 40
      ISIDs                  :
        80       Flags:TR
        81       Flags:TR
        82       Flags:TR
        83       Flags:TR
        84       Flags:TR
        85       Flags:TR
        86       Flags:TR
        87       Flags:TR
        88       Flags:TR
        89       Flags:TR
        90       Flags:TR
...
-------------------------------------------------------------------------------
LSP ID    : PE-69.00-00                                  Level     : L1
...

Level (1) LSP Count : 4
===============================================================================
```

The **mfib** parameter in the **show service id 40 sap static-isids mfib** command can help understand the state of the MFIB entries added (or not) by the configured static-isid. The following possible states can be shown:

- If the static-isid is configured and programmed in the mfib the status is shown as:
  → ok
- If the static-isid is not configured and not programmed in the mfib, the reasons can be (order of priority):
  → useDefMCTree - isid policy is applied on the service for the isid.
  → sysMFibLimit - system mfib limit has been exceeded
  → addPending - adding pending due to processing delays
- If the static-isid is not configured but present in the mfib:
  → delPending - cleanup pending due to processing delays.

The following examples show some of these possible states:

```
*A:PE-69# show service id 40 sap lag-1:40 static-isids mfib
===================================
ISID Detail
===================================
ISID           Status
-----------------------------------
<<snip>>
80             ok
81             ok
<<snip>>
41             useDefMCTree
42             useDefMCTree
<<snip>>
8292           sysMFibLimit
8293           sysMFibLimit
<<snip>>
9253          addPending
9254          addPending
<<snip>>
```

```
=====================================
```

Conclusion

SR OS supports an efficient SPBM implementation in the context of a B-VPLS, where system BMACs, vMEP OAM BMACs and SAP-BMACs are advertised in SPB IS-IS. SPBM provides a simple solution where no other control plane protocol is required in the B-VPLS to take care of the resiliency, load-balancing and multicast optimization. The SPBM implementation in the SR OS provides scale optimization through the use of control and user B-VPLSs, allows the interworking between SPB networks and PBB networks, as well as the optimization of the MFIB resources and advertisement of ISIDs through the use of ISID-policies.