NOKIA

7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM
VIRTUALIZED SERVICE ROUTER

LAYER 2 SERVICES AND EVPN GUIDE: VLL, VPLS, PBB, AND EVPN
RELEASE 15.0.R4

3HE 11970 AAAB TQZZA 01

Issue: 01

July 2017

Table of Contents

1 Getting Started

1.1 About This Guide

This guide describes Layer 2 service and EVPN functionality provided by the Nokia family of routers and presents examples to configure and implement various protocols and services.

This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- 7950 XRS
- VSR

Table 1 lists the available chassis types for each SR OS router.

Table 1 **Supported SR OS Router Chassis Types**

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-6/6v • 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e	• 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the *SR OS R15.0.Rx Software Release Notes*, part number 3HE 12060 000*x* TQZZA or the *VSR Release Notes*, part number 3HE 12092 000*x* TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.

Note: This guide generically covers Release 15.0.R*x* content and may contain some content that will be released in later maintenance loads. Refer to *SR OS R15.0.Rx Software Release Notes*, part number 3HE 12060 000*x* TQZZA or the *VSR Release Notes*, part number 3HE 12092 000*x* TQZZA, for information on features supported in each load of the Release 15.0.R*x* software.

1.2 Layer 2 Services and EVPN Configuration Process

Table 2 lists the tasks related to configuring and implementing Layer 2 Services and EVPN functionality.

This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 2 **Configuration Process**

Area	Task	Section
VLL Services	Configure services for MPLS-TP	Services Configuration for MPLS-TP
	Configure VCCV BFD	VCCV BFD support for VLL, Spoke-SDP Termination on IES and VPRN, and VPLS Services
	Configure pseudowire switching	Pseudowire Switching
	Configure a VLL service	Configuring a VLL Service with CLI
Virtual Private LAN Service (VPLS)	Configure a VPLS service	Configuring a VPLS Service with CLI
	VPLS service management	Service Management Tasks
Ethernet Virtual Private Networks (EVPNs)	Configure EVPN-VXLAN and EVPN-MPLS in the same VPLS service	Configuring EVPN-VXLAN and EVPN-MPLS in the Same VPLS Service
	Configure an EVPN service	Configuring an EVPN Service with CLI

2 VLL Services

2.1 ATM VLL (Apipe) Services

This section provides information about the ATM VLL (Apipe) services and implementation.

This feature is supported on the 7450 ESS platform in mixed-mode.

2.1.1 Apipe For End-to-End ATM Service

An Apipe provides a point-to-point ATM service between users connected to 7450 ESS or 7750 SR nodes on an IP/MPLS network. Users are either directly connected to a PE or through an ATM access network. In both cases, an ATM PVC (for example, a virtual channel (VC) or a virtual path (VP)) is configured on the PE. This feature supports local cross-connecting when users are attached to the same PE node. VPI/VCI translation is supported in the Apipe.

PE1, PE2, and PE3 receive standard UNI/NNI cells on the ATM Service Access Point (SAP) that are then encapsulated into a pseudowire packet using the N:1 cell mode encapsulation or AAL5 SDU mode encapsulation according to RFC 4717, *Encapsulation Methods for Transport of ATM Over MPLS Networks*. When using N:1 cell mode encapsulation, cell concatenation into a pseudowire packet is supported. In this application, both VC- and VP-level connections are supported.

The ATM pseudowire is initiated using Targeted LDP (T-LDP) signaling as specified in RFC 4447, *Pseudo-wire Setup and Maintenance using LDP*. The SDP can be an MPLS or a GRE type.

Figure 1 shows an example of Apipe for end-to-end ATM service.

Figure 1 **Apipe for End-to-End ATM Service**

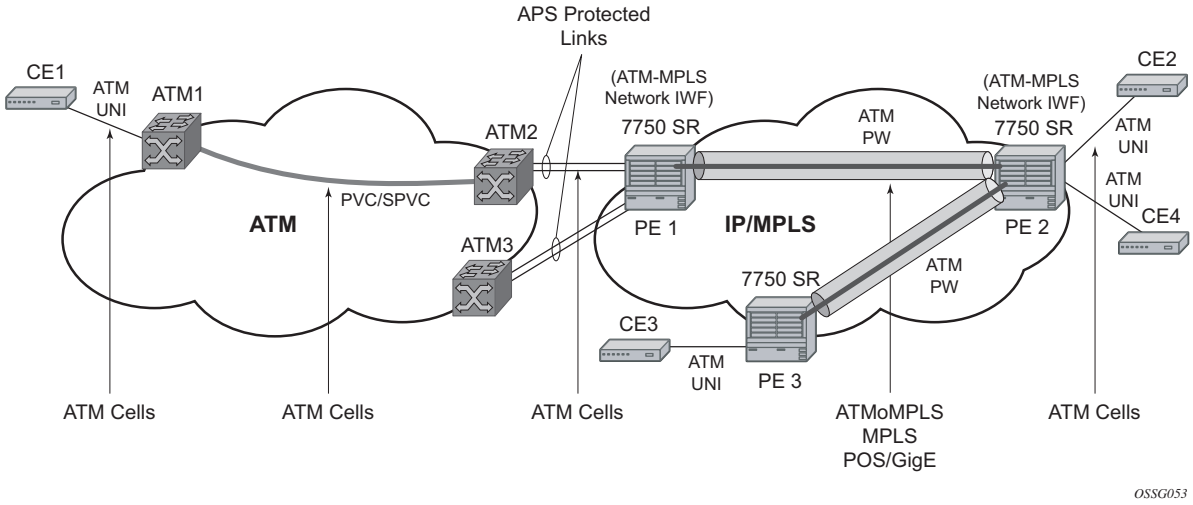

2.1.2 ATM Virtual Trunk Over IP/MPLS Packet Switched Network

For 7450 ESS or 7750 SR OS, ATM virtual trunk (VT) implements a transparent trunking of user and control traffic between two ATM switches over an ATM pseudowire. Figure 2 shows ATM 2 and ATM 3 switches that appear as if they are directly connected over an ATM link. Control traffic includes PNNI signaling and routing traffic.

Figure 2 **ATM VT Application**

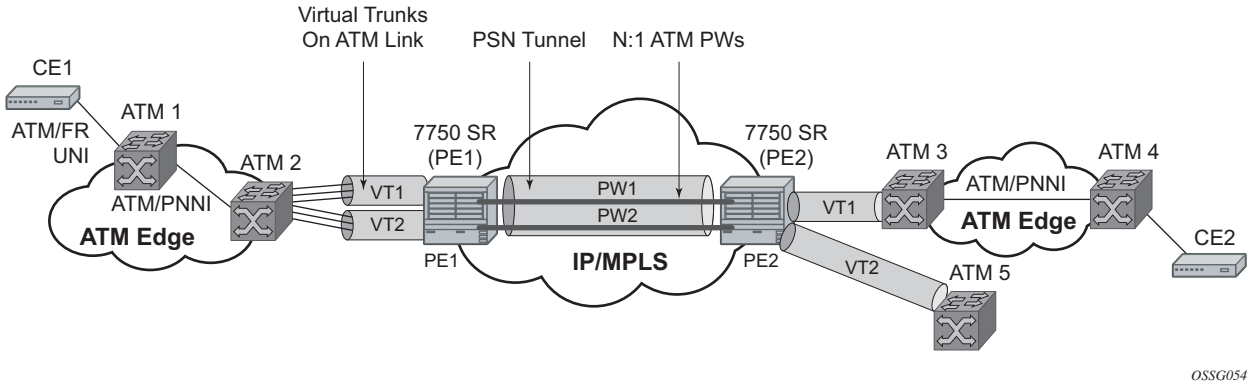

The VT SAP on a PE is identified by a tuple (port, VPI-range) meaning that all cells arriving on the specified port within the specified VPI range are fed into a single ATM pseudowire for transport across the IP/MPLS network. A user can configure the whole ATM port as a VT and does not need to specify a VPI range. No VPI/VCI translation is performed on ingress or egress. Cell order is maintained within a VT. As a special case, the two ATM ports could be on the same PE node.

By carrying all cells from all VPIs making up the VT in one pseudowire, a solution is provided that is robust; for example, black holes on some VPIs but not others. The solution is also operationally efficient since the entire VT can be managed as a single entity from the Network Manager (single point for configuration, status, alarms, statistics, and so on).

ATM virtual trunks use PWE3 N:1 ATM cell mode encapsulation to provide a cell-mode transport, supporting all AAL types, over the MPLS network. Cell concatenation on a pseudowire packet is supported. The SDP can be an MPLS or a GRE type.

The ATM pseudowire is initiated using Targeted LDP (T-LDP) signaling (defined in RFC 4447, *Pseudowire Setup and Maintenance using LDP*). In this application, there is no ATM signaling on the gateway nodes since both endpoints of the MPLS network are configured by the network operator. ATM signaling between the ATM nodes is passed transparently over the VT (along with user traffic) from one ATM port on a PE to another ATM port on a remote (or the same) SR PE.

2.1.3 Traffic Management Support

Traffic management support is supported only on the 7750 SR.

2.1.3.1 Ingress Network Classification

Classification is based on the EXP value of the pseudowire label and EXP-to-FC mapping is determined by the network ingress QoS policy.

2.1.3.2 Ingress Queuing and Shaping on the IOM

Each SAP of an ATM VLL has an associated single ingress service queue on the IOM. The default QoS policy configures this queue to have CIR=0 and PIR=line rate. Other QoS policies can be applied if they specify a single service queue. Applying a non-default QoS policy allows the CIR/PIR of the incoming traffic to be controlled, regardless of whether ATM policing is configured, and provides queuing and shaping to smooth traffic flows on the ingress of the network.

2.1.3.3 Egress Queuing and Shaping on the IOM

Each SAP of an ATM VLL has an associated single egress service queue on the IOM. The default QoS policy configures this queue to have CIR=0 and PIR=line rate. Other QoS policies can be applied if they specify a single service queue. Applying a non-default QoS policy allows the CIR/PIR of the outgoing traffic to be controlled, regardless of whether ATM shaping is configured.

2.1.3.4 Egress Shaping/Scheduling

Each SAP of an ATM VLL has an associated egress ATM traffic descriptor. The default traffic descriptor has service category UBR with zero MIR, resulting in endpoints associated with this descriptor being scheduled at the lowest priority on the ATM MDA. Egress traffic may be shaped or scheduled, depending on the configuration of the egress ATM traffic descriptor associated with the SAP. Table 3 describes how the different service categories, and shaping settings, and priorities affect egress transmission rates.

Shaping applies to CBR, rtVBR, and nrtVBR service categories and results in cells being transmitted in such a way as to satisfy a downstream ATM UPC function. For example, the transmission rate will be limited (in the case of CBR, there is a hard limit of PIR, while rtVBR/nrtVBR will transmit at SIR with short (constrained by MBS) bursts of up to PIR), and the inter-cell gap will also be controlled.

Service categories UBR and rtVBR are scheduled on the WRR scheduler with the configured rates (MIR for UBR+) determining the weight applied to the flow. Weights are between 1 and 255 and are determined by a formula applied to the configured rate. UBR flows (for example, those with no MIR) receive a weight of 1 and the maximum weight of 255 is reached by flows with configured rates of around 8 Mb/s. Scheduling does not apply a limit to the transmission rate; the available port bandwidth is shared out by the scheduler according to the weight, so if the other flows are quiescent, one flow may burst up to port bandwidth.

Shaping and scheduling of egress ATM VLL traffic is performed entirely at the ATM layer and is, therefore, not forwarding-class-aware. If the offered rate is greater than can be transmitted toward the customer (either because the shaping rate limits transmission or because the SAP does not receive sufficient servicing in the weighed round-robin used for scheduled SAPs), the per-VC queue will begin to discard traffic. These discards trigger the congestion control mechanisms in the MDA queues or in the IOM service egress queues associated with the SAP. For AAL5 SDU VLLs, these discards occur at the AAL5 SDU level. For N-to-1 VLLs, these discards occur at the level of the cell or a block of cells when cell concatenation is enabled.

Table 3 **Service Categories and Relative Priorities**

Flow Type	Transmission Rate	Priority
shaped CBR	Limited to configured PIR	Strict priority over all other traffic
shaped rtVBR	Limited to configured SIR, but with bursts up to PIR within MBS	Strict priority over all but shaped CBR
shaped nrtVBR	Limited to configured SIR, but with bursts up to PIR within MBS	Strict priority over all scheduled traffic
scheduled nrtVBR	Weighted share (according to SIR) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as UBR+ and UBR
scheduled UBR+	Weighted share (according to MIR) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as nrtVBR and UBR
scheduled UBR	Weighted share (with weight of 1) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as nrtVBR and UBR+

2.2 Circuit Emulation (Cpipe) Services

This section provides information about Circuit Emulation (Cpipe) services. Cpipe is supported for the 7450 ESS and 7750 SR only.

2.2.1 Mobile Infrastructure

Packet infrastructure is required within 2G, 2.5G, and 3G mobile networks to handle SMS messaging, web browsing, and emerging applications such as streaming video, gaming, and video on demand. Within existing 2.5G and 3G mobile networks, ATM is defined as the transport protocol. Within existing 2G networks, TDM is defined as the transport protocol. Due to the relatively low bit rate of existing handsets, most cell sites use 2 to 10 DS1s or E1s to transport traffic. When using ATM over multiple DS1/E1 links, Inverse Multiplexing over ATM (IMA) is very effective for aggregating the available bandwidth for maximum statistical gain and providing automatic resilience in the case of a link failure. Also, multiple DS1s or E1s are required to transport the 2G voice traffic.

Typically, low-cost devices are used at the many cell sites to transport multiple DS1 or E1 using ATM/IMA and TDM over an Ethernet/MPLS infrastructure. In Nokia applications, the circuit emulation would currently be performed using the 7705 SAR. This could be performed by DMXplore at the cell site. However, a large number of cell sites aggregate into a single switching center. Book-ending 7705 SAR nodes would require a very large number of systems at the switching center (see Figure 3). Therefore, a channelized OC3/STM1 solution is much more efficient at the switching center. Table 4 defines the cellsite backhaul types, CSR roles, and transport acronyms used in Figure 3.

With a channelized OC3/STM1 CES CMA/MDA in the 7750 SR, Nokia can provide a converged, flexible solution for IP/MPLS infrastructures for 2G/2.5G/3G mobile networks supporting both the CES (by CES CMA/MDA) and ATM/IMA transported traffic (by the ASAP MDA).

Figure 3 **Mobile Infrastructure**

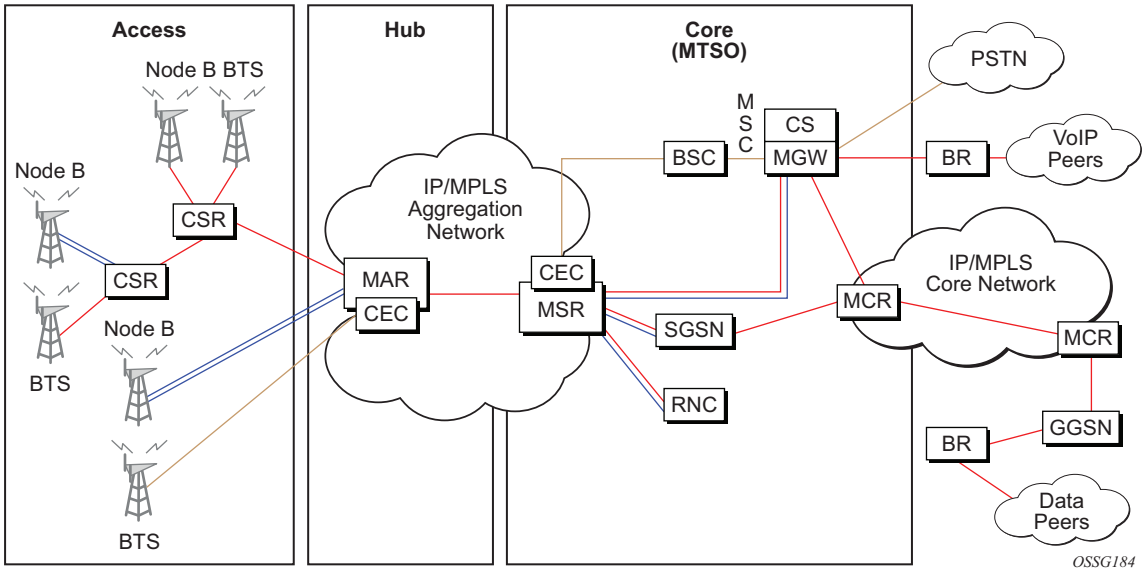

Table 4 **Mobile Infrastructure Definitions**

Cellsite Backhaul Type	CSR Role	Transport Acronyms
Microwave	Circuit emulation	CSR: Cellsite Service Router
xDSL	ATM IMA termination into pseudowire	MAR: Mobile Aggregation Router
Fiber, dark or light	Ethernet VLL switching	MSR: Mobile Service Router
ATM, ATM IMA	IP/MPLS aggregation	CEC: Circuit Emulation Concentrator
Leased line	—	MCR: Mobile Core Router
—	—	BR: Border Router

2.2.2 Circuit Emulation Modes

Two modes of circuit emulation are supported: unstructured and structured. Unstructured mode is supported for DS1 and E1 channels per RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*. Structured mode is supported for N*64 kb/s circuits as per RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*. Also, DS1, E1, and N*64 kb/s circuits are supported (per MEF8). TDM circuits are optionally encapsulated in MPLS or Ethernet as per the referenced standards in the following examples.

RFC 4553 (SAToP) MPLS PSN Encapsulation:

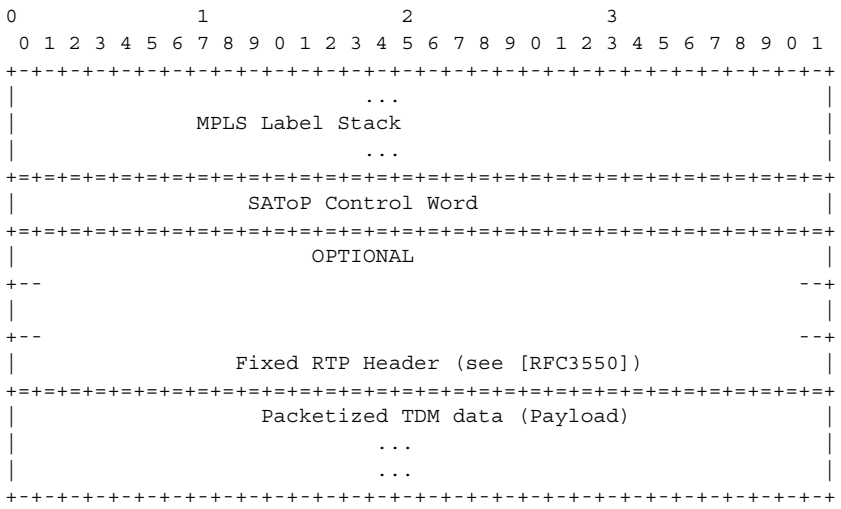

CESoPSN Packet Format for an MPLS PSN:

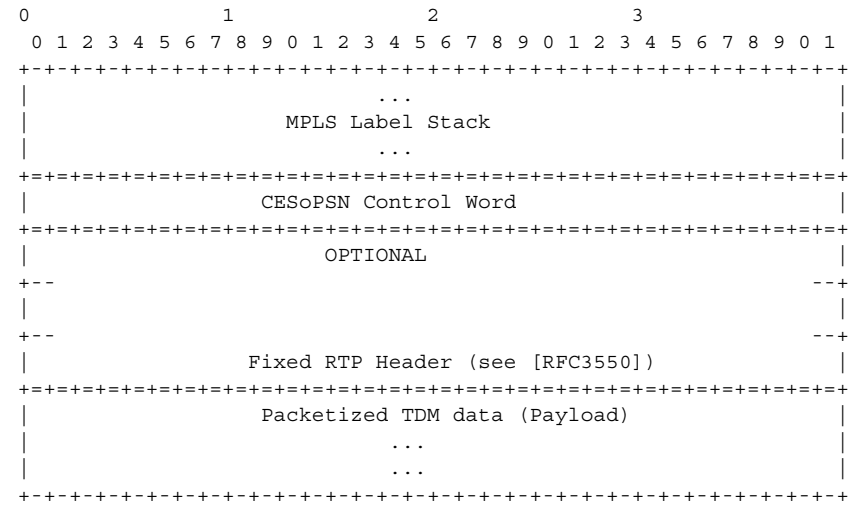

MEF8 PSN Encapsulation:

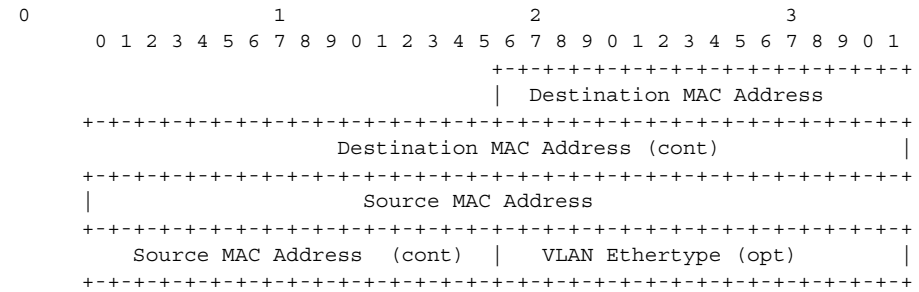

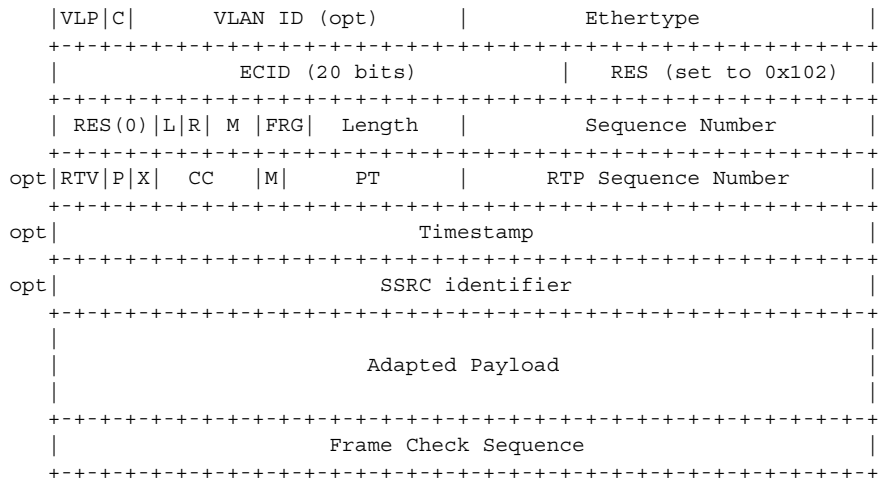

2.2.3 Circuit Emulation Parameters

2.2.3.1 Circuit Emulation Modes

All channels on the CES CMA/MDA are supported as circuits to be emulated across the packet network. Structure-aware mode is supported for N*64 kb/s channel groups in DS1 and E1 carriers. Fragmentation is not supported for circuit emulation packets (structured or unstructured).

For DS1 and E1 unstructured circuits, the framing can be set to unframed. When channel group 1 is created on an unframed DS1 or E1, it is automatically configured to contain all 24 or 32 channels, respectively.

N*64 kb/s circuit emulation supports basic and Channel Associated Signaling (CAS) options for timeslots 1 to 31 (channels 2 to 32) on E1 carriers and channels 1 to 24 on DS1 carriers. CAS in-band is supported; therefore, no separate pseudowire support for CAS is provided. CAS option can be enabled or disabled for all channel groups on a specific DS1 or E1. If CAS operation is enabled, timeslot 16 (channel 17) cannot be included in the channel group on E1 carriers. Control channel signaling (CCS) operation is not supported.

2.2.3.2 Absolute Mode Option

For all circuit emulation channels except those with differential clock sources, RTP headers in absolute mode can be optionally enabled (disabled by default). For circuit emulation channels that use differential clock sources, this configuration is blocked. All channel groups on a specific DS1 or E1 can be configured for the same mode of operation.

When enabled for absolute mode operation, an RTP header will be inserted. On transmit, the CES IWF will insert an incrementing (by 1 for each packet) timestamp into the packets. All other fields will be set to zero. The RTP header will be ignored on receipt. This mode is enabled for interoperability purposes only for devices that require an RTP header to be present.

2.2.3.3 Payload Size

For DS3, E3, DS1, and E1 circuit emulation, the payload size can be configurable in number of octets. The default values for this parameter are shown in Table 5. Unstructured payload sizes can be set to a multiple of 32 octets and minimally be 64 octets. TDM satellite supports only unstructured payloads.

Table 5 **Unstructured Payload Defaults**

TDM Circuit	Default Payload Size
DS1	192 octets
E1	256 octets

For N*64 kb/s circuits, the number of octets or DS1/E1 frames to be included in the TDM payload needs to be configurable in the range 4 to 128 DS1/E1 frames in increments of 1 or the payload size in octets. The default number of frames is shown in Table 6 with associated packet sizes. For the number of 64 kb/s channels included (N), the following number of default frames apply for no CAS: N=1, 64 frames; 2<=N<= 4, 32 frames; 5<=N<= 15, 16 frames; N>=16, 8 frames.

For CAS circuits, the number of frames can be 24 for DS1 and 16 for E1, which yields a payload size of N*24 octets for T1 and N*16 octets for E1. For CAS, the signaling portion is an additional ((N+1)/2) bytes, where N is the number of channels. The additional signaling bytes are not included in the TDM payload size, although they are included in the actual packet size shown in Table 6.

The full ABCD signaling value can be derived before the packet is sent. This occurs for every 24 frames for DS1 ESF and every 16 frames for E1. For DS1 SF, ABAB signaling is actually sent because SF framing only supports AB signaling every 12 frames.

Table 6 **Structured Number of Default Frames**

Num Timeslots	No CAS			DS1 CAS		E1 CAS	
	num-frames default	Default Payload	Minimum Payload	Payload (24 Frames)	Packet Size	Payload (16 Frames)	Packet Size
1	64	64	40	24	25	16	17
2	32	64	64	48	49	32	33
3	32	96	96	72	74	48	50
4	32	128	128	96	98	64	66
5	16	80	80	120	123	80	83
6	16	96	96	144	147	96	99
7	16	112	112	168	172	112	116
8	16	128	128	192	196	128	132
9	16	144	144	216	221	144	149
10	16	160	160	240	245	160	165
11	16	176	176	264	270	176	182
12	16	192	192	288	294	192	198
13	16	208	208	312	319	208	215
14	16	224	224	336	343	224	231
15	16	240	240	360	368	240	248
16	8	128	128	384	392	256	264
17	8	136	136	408	417	272	281
18	8	144	144	432	441	288	297
19	8	152	152	456	466	304	314
20	8	160	160	480	490	320	330
21	8	168	168	504	515	336	347
22	8	176	176	528	539	352	363

Table 6 **Structured Number of Default Frames (Continued)**

Num Timeslots	No CAS			DS1 CAS		E1 CAS	
	num-frames default	Default Payload	Minimum Payload	Payload (24 Frames)	Packet Size	Payload (16 Frames)	Packet Size
23	8	184	184	552	564	368	380
24	8	192	192	576	588	384	396
25	8	200	200	NA	NA	400	413
26	8	208	208	NA	NA	416	429
27	8	216	216	NA	NA	432	446
28	8	224	224	NA	NA	448	462
29	8	232	232	NA	NA	464	479
30	8	240	240	NA	NA	480	495
31	8	248	248	NA	NA	NA	NA

Note: num-frames DS1 CAS are multiples of 24; num-frames E1 is a multiple of 16.

2.2.3.4 Jitter Buffer

For each circuit, the maximum receive jitter buffer is configurable. Packet delay from this buffer starts when the buffer is 50% full, to give an operational packet delay variance (PDV) equal to 75% of the maximum buffer size. The default value for the jitter buffer is nominally 5 ms. However, for lower-speed N*64 kb/s circuits and CAS circuits, the following default values are used to align with the default number of frames (and resulting packetization delay) to allow at least two frames to be received before starting to playout the buffer. The jitter buffer is at least four times the packetization delay. The following default jitter buffer values for structured circuits apply:

Basic CES (DS1 and E1):

N=1, 32 ms

2<=N<=4, 16 ms

5<=N<=15, 8 ms

N>=16, 5 ms

2.2.3.5 CES Circuit Operation

The circuit status can be tracked to be either up, loss of packets, or administratively down. Statistics are available for the number of in-service seconds and the number of out-of-service seconds when the circuit is administratively up.

Jitter buffer overrun and underrun counters are available by statistics and optionally logged while the circuit is up. On overruns, excess packets are discarded and counted. On underruns, all ones are sent for unstructured circuits. For structured circuits, all ones or a user-defined data pattern is sent based on configuration. Also, if CAS is enabled, all ones or a user-defined signaling pattern is sent based on configuration.

For each CES circuit, alarms can be optionally disabled/enabled for stray packets, malformed packets, packet loss, receive buffer overrun, and remote packet loss. An alarm is raised if the defect persists for 3 seconds, and cleared when the defect no longer persists for 10 seconds. These alarms are logged and trapped when enabled.

2.2.4 Services for Transporting CES Circuits

Each circuit can be optionally encapsulated in MPLS, Ethernet packets. Circuits encapsulated in MPLS use circuit pipes (Cpipes) to connect to the far-end circuit. Cpipes support either SAP spoke-SDP or SAP-SAP connections. Cpipes are supported over MPLS and GRE tunnels. The Cpipe default service MTU is set to 1514 bytes.

Circuits encapsulated in Ethernet can be selected as a SAP in Epipes. Circuits encapsulated in Ethernet can be SAP spoke-SDP connections or Ethernet CEM SAP-to-Ethernet SAP for all valid Epipe SAPs. Circuits requiring CEM SAP-to-CEM SAP connections use Cpipes. A local and remote EC-ID and far-end destination MAC address can be configurable for each circuit. The CMA/MDA MAC address will be used as the source MAC address for these circuits.

For all service types, there are deterministic PIR=CIR values with class=EF parameters based on the circuit emulation parameters.

All circuit emulation services support the display of status of up, loss of packet (LOP), or admin down. Also, any jitter buffer overruns or underruns are logged.

Non-stop services are supported for Cpipes and CES over Epipes.

2.2.5 Network Synchronization Considerations

Each OC-3/STM-1 port can be independently configured to be loop-timed or node-timed. Each OC-3/STM-1 port can be configured to be a timing source for the node. TDM satellites only support node-timed mode.

Each DS-1 or E-1 channel without CAS signaling enabled can be independently configured to be loop-timed, node-timed, adaptive-timed, or differential-timed. Each DS-1 or E-1 channel with CAS signaling enabled can be independently configured to be loop-timed or node-timed. Adaptive timing and differential timing are not supported on DS-1 or E-1 channels with CAS signaling enabled. For the TDM satellite, each DS1/E1 channel can be loop-timed, node-timed, or differential-timed.

The adaptive recovered clock of a CES circuit can be used as a timing reference source for the node (ref1 or ref2). This is required to distribute network timing to network elements that only have packet connectivity to the network. One timing source on the CMA/MDA can be monitored for timing integrity. Both timing sources can be monitored if they are configured on separate CMA/MDAs while respecting the timing subsystem slot requirements.

If a CES circuit is being used for adaptive clock recovery at the remote end (such that the local end is now an adaptive clock master), Nokia recommends setting the DS-1/E-1 to be node-timed to prevent potential jitter issues in the recovered adaptive clock at the remote device. This is not applicable to TDM satellites.

For differential-timed circuits, the following timestamp frequencies are supported: 103.68 MHz (for recommended >100 MHz operation), 77.76 MHz (for interoperability with SONET/SDH-based systems such as TSS-5) and 19.44 MHz (for Y.1413 compliance). TDM satellite supports only 77.76 MHz.

Adaptive and differential timing recovery must comply with published jitter and wander specifications (G.823, G.824, and G.8261) for traffic interfaces under typical network conditions and for synchronous interfaces under specified packet network delay, loss, and delay variance (jitter) conditions. The packet network requirements to meet the synchronous interface requirements are to be determined during the testing phase.

On the 7450 ESS and 7750 SR CES CMA, a BITS port is also provided. The BITS port can be used as one of the two timing reference sources in the system timing subsystem. The operation of BITS ports configured as ref1 or ref2 is the same as existing ports configured as ref1 and ref2 with all options supported. The operation of the 7450 ESS or 7750 SR BITS source is unchanged and the BITS ports are not available on the CES MDAs (only SF/CPM BITS are available).

2.2.6 Cpipe Payload

Figure 4 shows the format of the CESoPSN TDM payload (with and without CAS) for packets carrying trunk-specific 64 kb/s service. In CESoPSN, the payload size is dependent on the number of timeslots used. This is not applicable to TDM satellite since only unstructured DS1/E1 is supported.

Figure 4 **CESoPSN MPLS Payload Format**

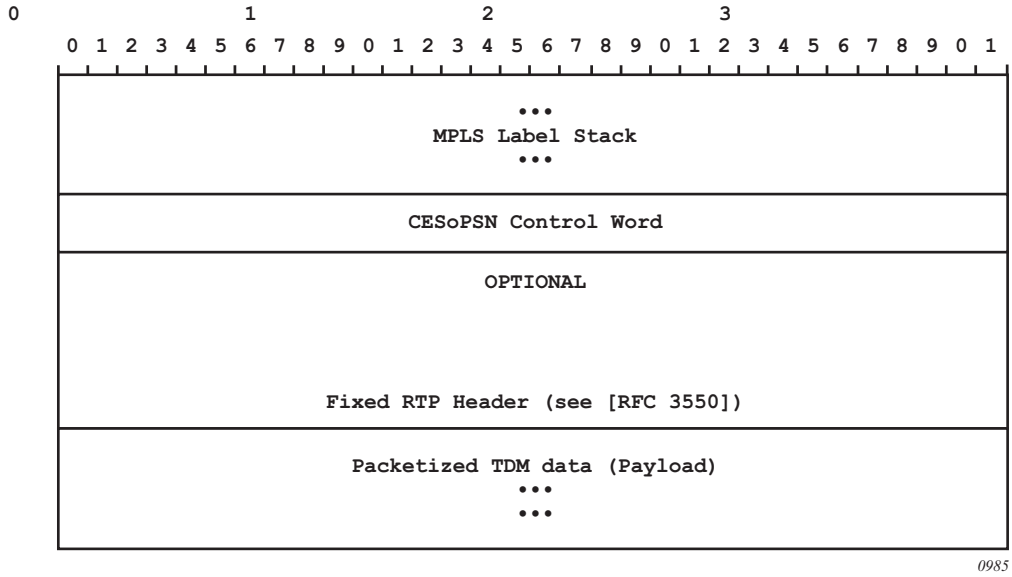

2.3 Ethernet Pipe (Epipe) Services

This section provides information about the Epipe service and implementation.

2.3.1 Epipe Service Overview

An Epipe service is the Nokia implementation of an Ethernet VLL based on the IETF “Martini Drafts” (draft-martini-l2circuit-trans-mpls-08.txt and draft-martini-l2circuit-encapmpls-04.txt) and the IETF Ethernet Pseudowire Draft (draft-so-pwe3-ethernet-00.txt).

An Epipe service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider IP, MPLS, or Provider Backbone Bridging (PBB) VPLS network. An Epipe service is completely transparent to the customer data and protocols. The Epipe service does not perform any MAC learning. A local Epipe service consists of two SAPs on the same node, whereas a distributed Epipe service consists of two SAPs on different nodes. SDPs are not used in local Epipe services.

Each SAP configuration includes a specific port or channel on which service traffic enters the router from the customer side (also called the access side). Each port is configured with an encapsulation type. If a port is configured with an IEEE 802.1Q (referred to as dot1q) encapsulation, a unique encapsulation value (ID) must be specified.

Figure 5 **Epipe/VLL Service**

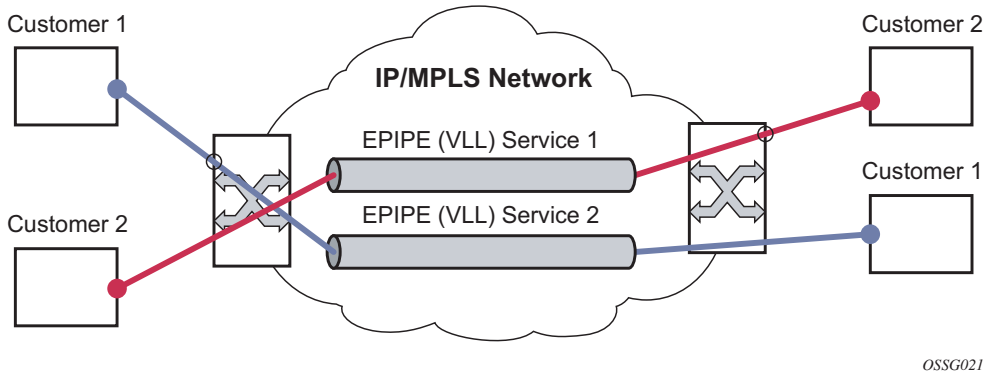

2.3.2 Epipe Service Pseudowire VLAN Tag Processing

Distributed Epipe services are connected using a pseudowire, which can be provisioned statically or dynamically and is represented in the system as a spoke-SDP. The spoke-SDP can be configured to process zero, one, or two VLAN tags as traffic is transmitted and received; see Table 7 and Table 8 for the ingress and egress tag processing. In the transmit direction, VLAN tags are added to the frame being sent. In the received direction, VLAN tags are removed from the frame being received. This is analogous to the SAP operations on a null, dot1q, and QinQ SAP.

The system expects a symmetrical configuration with its peer; specifically, it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. When removing VLAN tags from a spoke-SDP, the system attempts to remove the configured number of VLAN tags. If fewer tags are found, the system removes the VLAN tags found and forwards the resulting packet.

Because some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. With an asymmetrical behavior, a protocol extraction will not necessarily function as it would with a symmetrical configuration, resulting in an unexpected operation.

The VLAN tag processing is configured as follows on a spoke-SDP in an Epipe service:

- Zero VLAN tags processed — This requires the configuration of **vc-type ether** under the spoke-SDP, or in the related **pw-template**.
- One VLAN tag processed — This requires one of the following configurations:
  - **vc-type vlan** under the spoke-SDP or in the related **pw-template**
  - **vc-type ether** and **force-vlan-vc-forwarding** under the spoke-SDP or in the related **pw-template**
- Two VLAN tags processed — This requires the configuration of **force-qinq-vc-forwarding** under the spoke-SDP or in the related **pw-template**.

The **pw-template** configuration provides support for BGP VPWS services.

The following restrictions apply to VLAN tag processing:

- The configuration of **vc-type vlan** and **force-vlan-vc-forwarding** is mutually exclusive.
- **force-qinq-vc-forwarding** can be configured with the spoke-SDP signaled as either **vc-type ether** or **vc-type vlan**.
- The following are not supported with **force-qinq-vc-forwarding** configured under the spoke-SDP, or in the related **pw-template**:

– Multi-segment pseudowires.
– BGP VPWS routes are not accepted over an IBGP session.
– ETH-CFM MIPs and MEPs are not supported on dynamically signaled BGP QinQ PWs.

Table 7 and Table 8 describe the VLAN tag processing with respect to the zero, one, and two VLAN tag configuration described for the VLAN identifiers, Ethertype, ingress QoS classification (dot1p/DE), and QoS propagation to the egress (which can be used for egress classification and/or to set the QoS information in the innermost egress VLAN tag).

Table 7 **Epipe Spoke SDP VLAN Tag Processing: Ingress**

Ingress (Received on Spoke SDP)	**Zero VLAN Tags**	**One VLAN Tag**	**Two VLAN Tags**
VLAN identifiers	N/A	Ignored	Both inner and outer ignored
Ethertype (to determine the presence of a VLAN tag)	N/A	0x8100 or value configured under **sdp vlan-vc-etype**	Both inner and outer VLAN tags: 0x8100, or outer VLAN tag value configured under **sdp vlan-vc-etype** (inner VLAN tag value must be 0x8100)
Ingress QoS (dot1p/ DE) classification	N/A	Ignored	Both inner and outer ignored
QoeE (dot1p/DE) propagation to egress	Dot1p/DE=0	Dot1p/DE taken from received VLAN tag	Dot1p/DE taken from inner received VLAN tag

Table 8 **Epipe Spoke SDP VLAN Tag Processing: Egress**

Egress (Sent on Mesh or Spoke SDP)	Zero VLAN Tags	One VLAN Tag	Two VLAN Tags
VLAN identifiers (set in VLAN tags)	N/A	The tag is derived from one of the following: • the **vlan-vc-tag** value configured in **pw-template** or under the spoke-SDP • value from the inner tag received on a QinQ SAP or QinQ spoke-SDP • value from the VLAN tag received on a dot1q SAP or spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or spoke-SDP	Both the inner and outer VLAN tag are derived from one of the following: • The **vlan-vc-tag** value configured in **pw-template** or under the spoke-SDP • value from the inner tag received on a QinQ SAP or QinQ spoke-SDP • value from the VLAN tag received on a dot1q SAP or spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or spoke-SDP
Ethertype (set in VLAN tags)	N/A	0x8100 or value configured under **sdp vlan-vc-etype**	Both inner and outer VLAN tags: 0x8100, or outer VLAN tag value configured under **sdp vlan-vc-etype** (inner VLAN tag value will be 0x8100)

Table 8 **Epipe Spoke SDP VLAN Tag Processing: Egress (Continued)**

Egress (Sent on Mesh or Spoke SDP)	Zero VLAN Tags	One VLAN Tag	Two VLAN Tags
Egress QoS (dot1p/DE) (set in VLAN tags)	N/A	The tag taken from the innermost ingress service delimiting tag can be one of the following: • The inner tag received on a QinQ SAP or QinQ spoke-SDP • value from the VLAN tag received on a dot1q SAP or spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or spoke-SDP Note: Neither the inner nor outer dot1p/DE values can be explicitly set.	Both inner and outer dot1p/DE: Taken from the innermost ingress service delimiting tag can be one of the following: • The inner tag received on a QinQ SAP or QinQ spoke-SDP • value from the VLAN tag received on a dot1q SAP or spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or spoke-SDP Note: Neither the inner nor outer dot1p/DE values can be explicitly set.

Any non-service delimiting VLAN tags are forwarded transparently through the Epipe service. SAP egress classification is possible on the outermost customer VLAN tag received on a spoke-SDP using the **ethernet-ctag** parameter in the associated SAP egress QoS policy.

2.3.3 Epipe Up Operational State Configuration Option

By default, the operational state of the Epipe is tied to the state of the two connections that comprise the Epipe. If either of the connections in the Epipe are operationally down, the Epipe service that contains that connection will also be operationally down. The operator can configure a single SAP within an Epipe that does not affect the operational state of that Epipe, using the optional **ignore-oper-state** command. Within an Epipe, if a SAP that includes this optional command

becomes operationally down, the operational state of the Epipe will not transition to down. The operational state of the Epipe will remain up. This does not change that the SAP is down and no traffic can transit an operationally down SAP. Removing and adding this command on the fly will evaluate the operational state of the service, based on the SAPs and the addition or deletion of this command.

Service OAM (SOAM) designers may consider using this command if an operationally up MEP configured on the operationally down SAP within an Epipe is required to receive and process SOAM PDUs. When a service is operationally down, this is not possible. For SOAM PDUs to continue to arrive on an operationally up, MEP configured on the failed SAP, the service must be operationally up. Consider the case where an operationally up MEP is placed on a UNI-N or E-NNI and the UNI-C on E-NNI peer is shutdown in such a way that it causes the SAP to become operationally down.

Two connections must be configured within the Epipe; otherwise, the service will be operationally down regardless of this command. The **ignore-oper-state** functionality will only operate as intended when the Epipe has one ingress and one egress. This command is not to be used for Epipe services with redundant connections that provide alternate forwarding in case of failure, even though the CLI does not prevent this configuration.

Support is available on Ethernet SAPs configured on ports or Ethernet SAPs configured on LAG. However, it is not allowed on SAPs using LAG profiles or if the SAP is configured on a LAG that has no ports.

2.3.4 Epipe with PBB

A PBB tunnel may be linked to an Epipe to a B-VPLS. MAC switching and learning is not required for the point-to-point service. All packets ingressing the SAP are PBB encapsulated and forwarded to the PBB tunnel to the backbone destination MAC address. Likewise, all the packets ingressing the B-VPLS destined for the ISID are PBB de-encapsulated and forwarded to the Epipe SAP. A fully specified backbone destination address must be provisioned for each PBB Epipe instance to be used for each incoming frame on the related I-SAP. If the backbone destination address is not found in the B-VPLS FDB, packets may be flooded through the B-VPLSs.

All B-VPLS constructs may be used including B-VPLS resiliency and OAM. Not all generic Epipe commands are applicable when using a PBB tunnel.

2.3.5 Epipe over L2TPv3

The L2TPv3 feature provides a framework to transport Ethernet pseudowire services over an IPv6-only network without MPLS. This architecture relies on the abundance of address space in the IPv6 protocol to provide unique far-end and local-end addressing that uniquely identify each tunnel and service binding.

L2TPv3 provides the capability of transporting multiple Epipes (up to 16K per system), by binding multiple IPv6 addresses to each node and configuring one SDP per Epipe.

Because the IPv6 addressing uniqueness identifies the customer and service binding, the L2TPv3 control plane is disabled in this mode.

L2TPv3 is supported on non-12e 7750 SR and 7450 ESS (mixed mode) and 7950 XRS platforms.

ETH-CFM is supported for OAM services.

Figure 6 **L2TPv3 SDP**

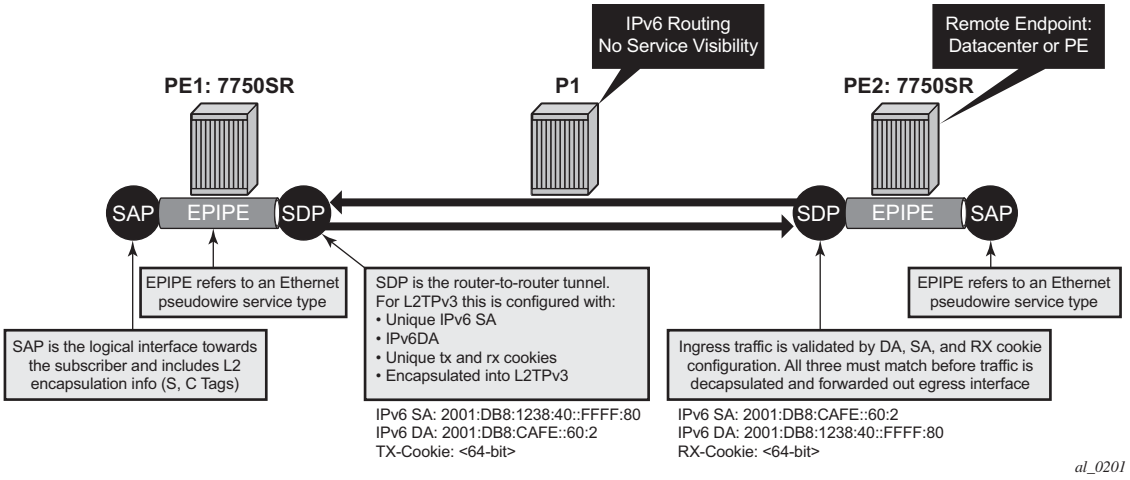

2.3.6 Ethernet Interworking VLL

Figure 7 provides an example of an Ethernet interworking VLL. The Ethernet interworking VLL provides a point-to-point Ethernet VLL service between Frame Relay (FR) attached users, ATM-attached users, and Ethernet-attached users across an IP/MPLS packet switched network. It effectively provides ATM and FR bridged encapsulation termination on the existing Epipe service of the 7750 SR.

Figure 7 **Application of Ethernet Interworking VLL**

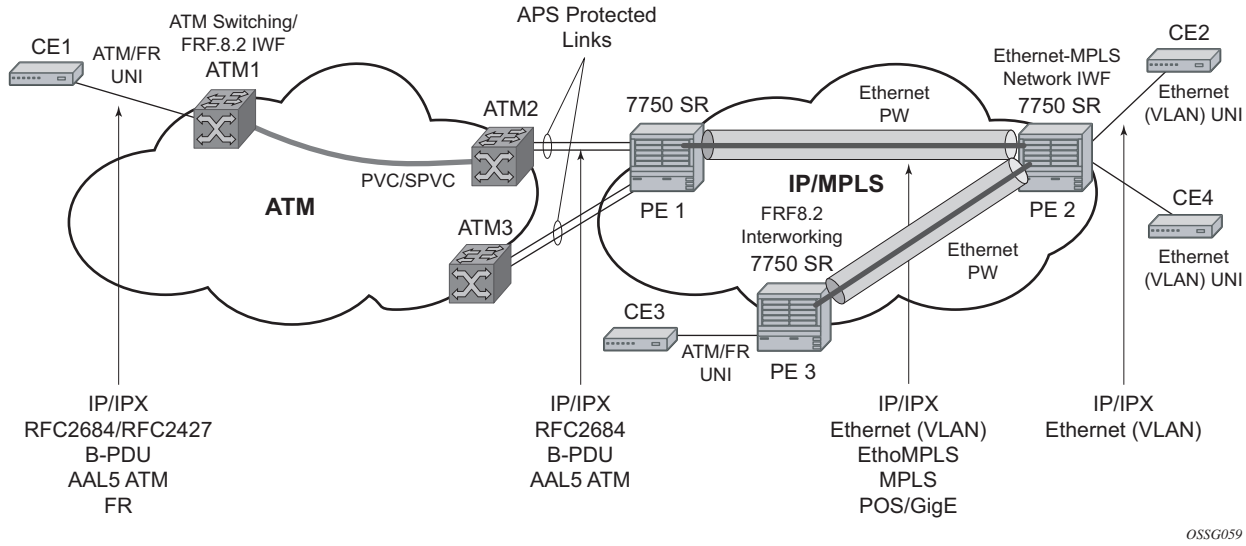

The following connectivity scenarios are supported:

- a Frame Relay or ATM user connected to a ATM network communicating with a Ethernet user connected to a 7750 SR PE node on a IP/MPLS network
- a Frame Relay or ATM user connected to 7750 SR PE node communicating with an Ethernet user connected to a 7750 SR PE node on a IP/MPLS network. This feature supports local cross-connecting when these users are attached to the same 7750 SR PE node.

Users attach over an ATM UNI with RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, tagged/untagged bridged Ethernet PDUs, a FR UNI using RFC 2427, *Multiprotocol Interconnect over Frame Relay*, tagged/untagged bridged Ethernet PDUs, or an Ethernet tagged/untagged UNI interface. However, the VCI/VPI and the data-link connection identifier (DLCI) are the identifiers of the SAP in the case of ATM and FR, respectively, and the received tags are transparent to the service, so are preserved.

The Ethernet pseudowire is established using T-LDP signaling and can use the **ether** or **vlan** VC types on the SDP. The SDP can be either an MPLS or GRE type.

2.3.7 VLL CAC

The VLL Connection Admission Control (CAC) is supported for the 7750 SR only and provides a method to administratively account for the bandwidth used by VLL services inside an SDP that consists of RSVP LSPs.

The service manager keeps track of the available bandwidth for each SDP. The SDP available bandwidth is applied through a configured booking factor. An administrative bandwidth value is assigned to the spoke-SDP. When a VLL service is bound to an SDP, the amount of bandwidth is subtracted from the adjusted available SDP bandwidth. When the VLL service binding is deleted from the SDP, the amount of bandwidth is added back into the adjusted SDP available bandwidth. If the total adjusted SDP available bandwidth is overbooked when adding a VLL service, a warning is issued and the binding is rejected.

This feature does not guarantee bandwidth to a VLL service because there is no change to the datapath to enforce the bandwidth of an SDP by means such as shaping or policing of constituent RSVP LSPs.

2.3.8 MC-Ring and VLL

To support redundant VLL access in ring configurations, the multi-chassis ring (MC-Ring) feature is applicable to VLL SAPs. A conceptual drawing of the operation is shown in Figure 8. The specific CPE that is connected behind the ring node has access to both BSAs through the same VLAN provisioned in all ring nodes. There are two SAPs (with the same VLAN) provisioned on both nodes.

If a closed ring status occurs, one of the BSAs becomes the master and will signal an active status bit on the corresponding VLL pseudowire. Similarly, the standby BSA will signal a standby status. With this information, the remote node can choose the correct path to reach the CPE. In case of a broken ring, the node that can reach the ring node, to which the CPE is connected by RNCV check, will become master and will signal corresponding status on its pseudowire.

The mapping of individual SAPs to the ring nodes is done statically through CLI provisioning. To keep the convergence time to a minimum, MAC learning must be disabled on the ring node so all CPE originated traffic is sent in both directions. If the status is operationally down on the SAP on the standby BSA, that part of the traffic will be blocked and not forwarded to the remote site.

Figure 8 **MC-Ring in a Combination with VLL Service**

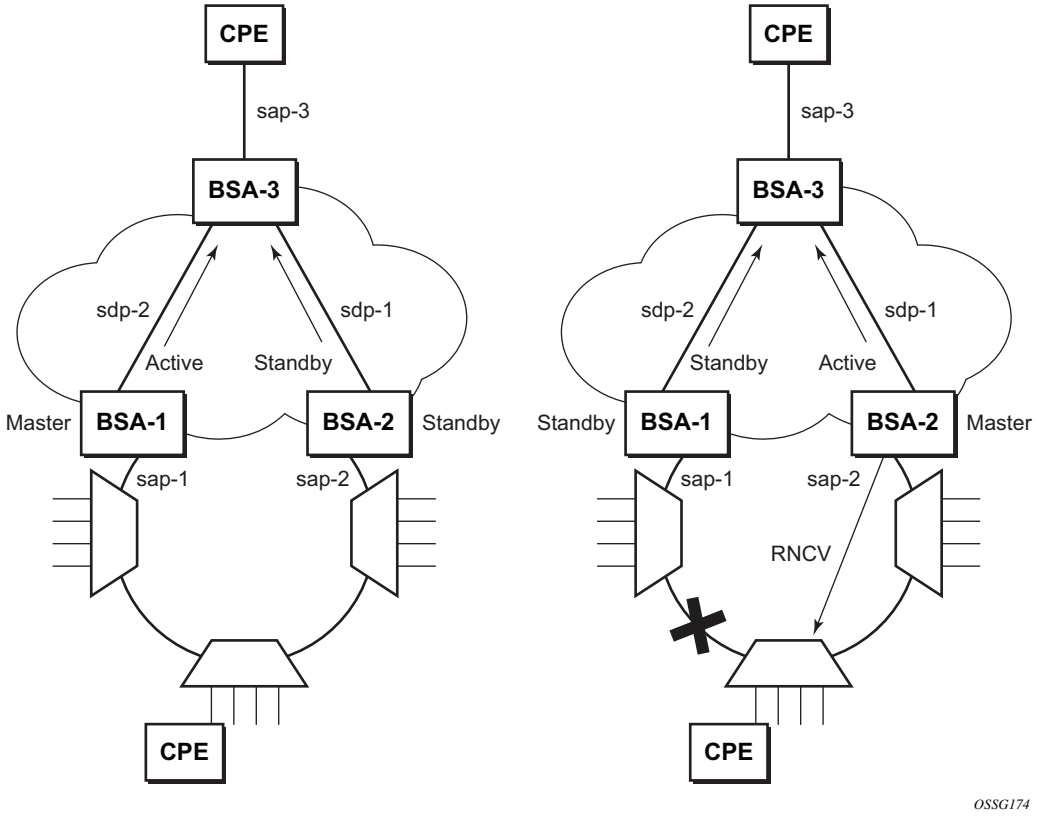

For further information about Multi-Chassis Ring Layer 2 (with ESM), refer to the *Advanced Configuration Guide*.

2.4 Frame Relay VLL (Fpipe) Services

This section provides information about the Frame Relay VLL (Fpipe) service and implementation. Fpipe is supported for the 7750 SR only.

2.4.1 Frame Relay VLL

Figure 9 shows an application of a Frame Relay VLL. The Fpipe provides a point-to-point Frame Relay service between users connected to 7750 SR nodes on the IP/MPLS network. Users are connected to the 7750 SR.

Figure 9 **Application of a Fpipe**

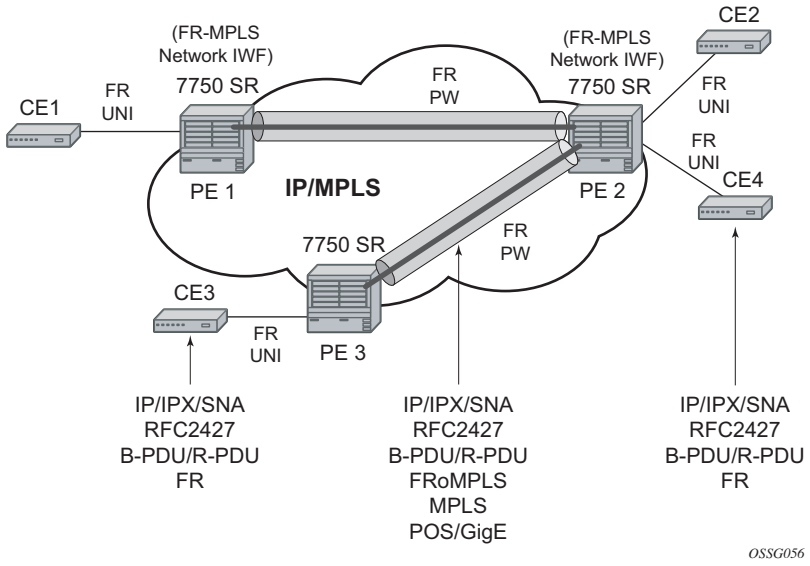

PE nodes using Frame Relay PVCs PE1, PE2, and PE3 receive a standard Q.922 Core frame on the Frame Relay SAP and encapsulate it into a pseudowire packet according to the 1-to-1 Frame Relay encapsulation mode in RFC 4619, *Encapsulation Methods for Transport of Frame Relay Over MPLS Networks*. The 7750 SR Fpipe feature supports local cross-connecting when the users are attached to the same 7750 SR PE node.

The FR pseudowire is initiated using T-LDP signaling as specified in RFC 4447, *Pseudo-wire Setup and Maintenance using LDP*. The SDP can be an MPLS or a GRE type.

2.4.2 Frame Relay-to-ATM Interworking (FRF.5) VLL

Figure 10 provides an example of a point-to-point Frame Relay service between users where one user is connected to an existing ATM network, the other to a 7750 SR PE node on an IP/MPLS network.

Figure 10 **Frame Relay-to-ATM Network Interworking (FRF.5) VLL**

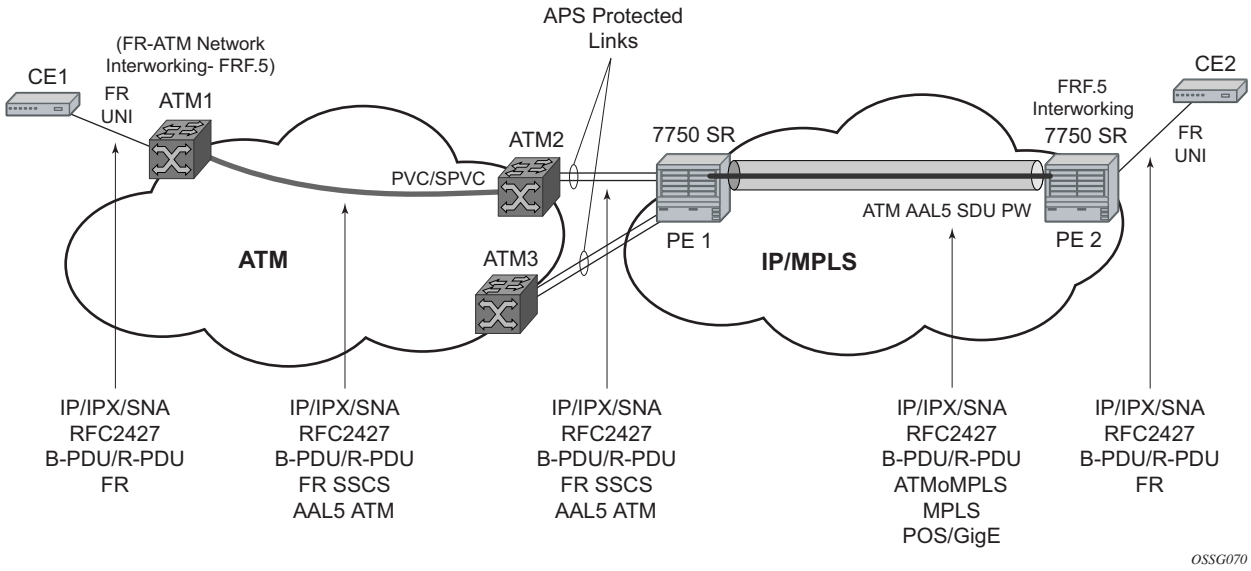

This VLL uses an ATM AAL5 SDU pseudowire between the 7750 SR SR PE nodes. It is configured by adding a FR SAP to an Apipe service using vc-type atm-sdu. The 7750 SR SR PE2 node performs an FRF.5 interworking function to interwork the ingress and egress data paths as well as the operations required in an FR and an ATM VLL.

The pseudowire is initiated using Targeted LDP signaling as specified in the IETF Draft *draft-ietf-pwe3-control-protocol-xx.txt*. The SDP can be an MPLS or a GRE type.

2.4.3 Traffic Management Support

2.4.3.1 Frame Relay Traffic Management

Traffic management of Fpipes is supported for the 7750 SR only and is achieved through the application of ingress and egress QoS policies to SAPs like other Frame Relay SAPs. No queuing occurs on the MDA; all queuing, policing, and shaping occurs on the IOM and, therefore, traffic management is forwarding-class-aware. Forwarding classes may be determined by inspecting the DSCP marking of contained IP packets (for example) and this will determine both the queuing and the EXP bit setting of packets on an Fpipe.

2.4.3.2 Ingress SAP Classification and Marking

Ingress SAP classification and marking is supported for the 7450 ESS and 7750 SR only. DE=0 frames are subject to the CIR marking algorithm in the queue. Drop preference for these packets will follow the state of the CIR bucket associated with the ingress queue. The value is marked in the drop preference bit of the internal header and in the DE bit in the Q.922 frame header. DE=1 frames are classified in “out-of-profile” state and are not overwritten by the CIR marking in the ingress queue. The drop preference is set to high.

2.4.3.3 Egress Network EXP Marking

FC-to-EXP mapping is supported for the 7450 ESS and 7750 SR only and is as per the Network Egress QoS policy. Marking of the EXP field in both label stacks is performed.

2.4.3.4 Ingress Network Classification

Classification is supported for the 7450 ESS and 7750 SR only and is based on the EXP value of the pseudowire label and EXP-to-FC mapping is as per Network Ingress QoS policy.

2.5 IP Interworking VLL (Ipipe) Services

This section provides information about IP Interworking VLL (Ipipe) services.

2.5.1 Ipipe VLL

Figure 11 provides an example of IP connectivity between a host attached to a point-to-point access circuit (FR, ATM, PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts appear to be on the same LAN segment. This feature is supported for the 7450 ESS and 7750 SR and enables service interworking between different link layer technologies. A typical use of this application is in a Layer 2 VPN when upgrading a hub site to Ethernet while keeping the spoke sites with their existing Frame Relay or ATM IPv4 (7750 SR only) routed encapsulation.

Figure 11 **IP Interworking VLL (Ipipe)**

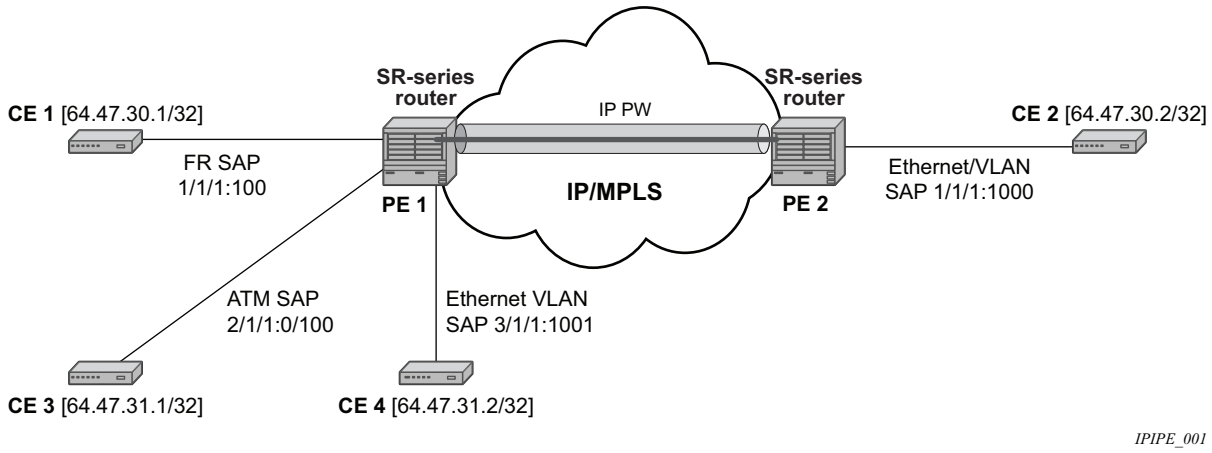

The ATM SAP is supported by the 7750 SR only. It carries the IPv4 packet using RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, *VC-Mux or LLC/SNAP routed PDU encapsulation.*

The Frame Relay SAP uses RFC 2427*, Multiprotocol Interconnect over Frame Relay,* routed PDU encapsulation of an IPv4 packet. A PPP interface uses RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*, PPP IPCP encapsulation of an IPv4 packet. A Cisco-HDLC SAP uses the routed IPv4 encapsulation. The pseudowire uses the IP Layer 2 transport pseudowire encapsulation type.

Note: The Ipipe is a point-to-point Layer 2 service. All packets received on one SAP of the Ipipe will be forwarded to the other SAP. No IP routing of customer packets occurs.

2.5.2 IP Interworking VLL Datapath

In Figure 11, PE 2 is manually configured with both CE 1 and CE 2 IP addresses. These are host addresses and are entered in /32 format. PE 2 maintains an ARP cache context for each IP interworking VLL. PE 2 responds to ARP request messages received on the Ethernet SAP. PE 2 responds with the Ethernet SAP configured MAC address as a proxy for any ARP request for CE 1 IP address. PE 2 silently discards any ARP request message received on the Ethernet SAP for an address other than that of CE 1. Likewise, PE 2 silently discards any ARP request message with the source IP address other than that of CE 2. In all cases, PE 2 keeps track of the association of IP to MAC addresses for ARP requests it receives over the Ethernet SAP.

To forward unicast frames destined for CE 2, PE 2 needs to know the CE 2 MAC address. When the Ipipe SAP is first configured and administratively enabled, PE2 sends an ARP request message for CE 2 MAC address over the Ethernet SAP. Until an ARP reply is received from CE2, providing the CE2 MAC address, unicast IP packets destined for CE2 will be discarded at PE2. IP broadcast and IP multicast packets are sent on the Ethernet SAP using the broadcast or direct-mapped multicast MAC address.

To forward unicast frames destined for CE 1, PE 2 validates the MAC destination address of the received Ethernet frame. The MAC address should match that of the Ethernet SAP. PE 2 then removes the Ethernet header and encapsulates the IP packet directly into a pseudowire without a control word. PE 1 removes the pseudowire encapsulation and forwards the IP packet over the Frame Relay SAP using RFC 2427, *Multiprotocol Interconnect over Frame Relay,* routed PDU encapsulation.

To forward unicast packets destined for CE1, PE2 validates the MAC destination address of the received Ethernet frame. If the IP packet is unicast, the MAC destination must match that of the Ethernet SAP. If the IP packet is multicast or broadcast, the MAC destination address must be an appropriate multicast or broadcast MAC address.

The other procedures are similar to the case of communication between CE 1 and CE 2, except that the ATM SAP and the Ethernet SAP are cross-connected locally and IP packets do not get sent over an SDP.

A PE does not flush the ARP cache unless the SAP goes administratively or operationally down. The PE with the Ethernet SAP sends unsolicited ARP requests to refresh the ARP cache every “T” seconds. ARP requests are staggered at an increasing rate if no reply is received to the first unsolicited ARP request. The value of T is configurable by the user through the mac-refresh CLI command.

2.5.3 Extension to IP VLL for Discovery of Ethernet CE IP Address

VLL services provide IP connectivity between a host attached to a point-to-point access circuit (FR, ATM, PPP) with routed PDU encapsulation and a host attached to an Ethernet interface. Both hosts appear to be on the same IP interface. This feature is supported only for IPv4 payload.

In deployments where it is not practical for operators to obtain and configure their customer CE address, the following behaviors apply:

- A service comes up without prior configuration of the CE address parameter under both the SAP and the spoke-SDP.
- Operators rely solely on received ARP messages from the Ethernet SAP-attached CE device to update the ARP cache with no further check of the validity of the source IP address of the ARP request message and the target IP address being resolved.
- The LDP address list TLV signaling the learned CE IP address to the remote PE is supported. This is to allow the PE with the FR SAP to respond to an invFR ARP request message received from the FR-attached CE device. Only Ethernet SAP and FR SAP can learn the CE address through ARP and invFR ARP, respectively. The 7450 ESS and 7750 SR OS do not support invATM ARP on an ATM interface.

2.5.3.1 VLL Ethernet SAP Processes

The operator can enable the following CE address discovery processes by configuring the **ce-address-discovery** in the **config>service>ipipe** context.

- The service is brought up without the CE address parameter configured at either the SAP or the spoke-SDP.

- The operator cannot configure the **ce-address** parameter under the **config>service>ipipe>sap** or **config>service>ipipe>spoke-sdp** context when the **ce-address-discovery** in the **config>service>ipipe** context is enabled. Conversely, the operator is not allowed to enable the **ce-address-discovery** option under the Ipipe service if it has a SAP and/or spoke-SDP with a user-entered **ce-address** parameter.
- While an ARP cache is empty, the PE does not forward unicast IP packets over the Ethernet SAP but forwards multicast/broadcast packets. target IP address being resolved.
- The PE waits for an ARP request from the CE to learn both IP and MAC addresses of the CE. Both entries are added into the ARP cache. The PE accepts any ARP request message received over Ethernet SAP and updates the ARP cache IP and MAC entries with no further check of the source IP address of the ARP request message or of the target IP address being resolved.
- The 7450 ESS, 7750 SR, and 7950 XRS routers will always reply to a received ARP request message from the Ethernet SAP with the SAP MAC address and a source IP address of the target IP address being resolved without any further check of the latter.
- If the router received an address list TLV from the remote PE node with a valid IP address of the CE attached to the remote PE, the router will not check the CE IP address against the target IP address being resolved when replying to an ARP request over the Ethernet SAP.
- The ARP cache is flushed when the SAP bounces or when the operator manually clears the ARP cache. This results in the clearing of the CE address discovered on this SAP. However, when the SAP comes up initially or comes back up from a failure, an unsolicited ARP request is not sent over the Ethernet SAP.
- If the Ipipe service uses a spoke-SDP, the router includes the address list TLV in the interface parameters field of the pseudowire Forwarding Equivalent Class (FEC) TLV in the label mapping message. The address list TLV contains the current value of the CE address in the ARP cache. If no address was learned, an address value of 0.0.0.0 must be used.
- If the remote PE included the address list TLV in the received label mapping message, the local router updates the remote PE node with the most current IP address of the Ethernet CE using a T-LDP notification message with the TLV status code set to 0x0000002C and containing an LDP address list. The notification message is sent each time an IP address different from the current value in the ARP cache is learned. This includes when the ARP is flushed and the CE address is reset to the value of 0.0.0.0.
- If the remote PE did not include the address list TLV in the received label mapping message, the local router will not send any notification messages containing the address list TLV during the lifetime of the IP pseudowire.

- If the operator disables the **ce-address-discovery** option under the VLL service, service manager instructs LDP to withdraw the service label and the service is shutdown. The pseudowire labels will only be signaled and the service will come up if the operator re-enters the option again or manually enters the **ce-address** parameter under SAP and spoke-SDP.

2.5.3.1.1 VLL FR SAP Procedures

The operator enables the following CE address dynamic learning procedures by enabling the **ce-address-discovery** option under the VLL service on the 7450 ESS or 7750 SR.

- Allow the service to come up without the CE address parameter configured at both the SAP and spoke-SDP. If one or both parameters are configured, they are ignored.
- The operator cannot configure the **ce-address** parameter under SAP or spoke-SDP when the **ce-address-discovery** option under the VLL service is enabled. Conversely, the operator is not allowed to enable the **ce-address-discovery** option under the Ipipe service if it has a SAP and/or spoke-SDP with a user-entered **ce-address** parameter.
- If the router receives an invFR ARP request message over the FR SAP, it updates the ARP cache with the FR CE address. It also replies with the IP address of the CE attached to the remote PE if a valid address was advertised in the address list TLV by this remote PE. Otherwise, the router updates the ARP cache but does not reply to the invFR ARP.
- If the Ipipe service uses a spoke-SDP, the router includes the address list TLV in the interface parameters field of the pseudowire FEC TLV in the label mapping message. The address list TLV contains the current value of the CE address in the ARP cache. If no address was learned, then an address value of 0.0.0.0 is used.
- If the remote PE included the address list TLV in the received label mapping message, the local router updates the remote PE node with the most current IP address of the FR CE using a T-LDP status notification message containing an LDP address list. The notification message is sent each time an IP address different from the current value in the ARP cache is learned. This includes when the ARP is flushed and the CE address is reset to the value of 0.0.0.00.
- If the remote PE did not include the address list TLV in the received label mapping message, the local router does not send any notification messages containing the address list TLV during the lifetime of the IP pseudowire.

2.5.3.1.2 VLL ATM SAP Procedures

The operator enables the following CE address dynamic learning procedures by enabling the **ce-address-discovery** option under the VLL service on the 7750 SR.

- Allow the service to come up without the **ce-address** parameter configured at both the SAP and spoke-SDP. If one or both parameters are configured, they are ignored.
- The operator is not allowed to configure the **ce-address** parameter under the SAP or spoke-SDP when the **ce-address-discovery** option under the VLL service is enabled. Conversely, the operator is not allowed to enable the **ce-address-discovery** option under the Ipipe service if it has a SAP and/or spoke-SDP with a user-entered **ce-address** parameter.
- If the router receives an invATM ARP request message over the ATM SAP, the router silently discards it. The router does not support receiving or sending of an invATM ARP message.
- If the Ipipe service uses a spoke-SDP, the router includes the address list TLV in the interface parameters field of the pseudowire FEC TLV in the label mapping message. The address list TLV contains an address value of 0.0.0.0.
- If the remote PE included the address list TLV in the received label mapping message, the local router will not make further updates to the address list TLV to the remote PE node using a T-LDP status notification message since the learned IP address of the ATM-attached CE will never change away from the value of 0.0.0.0.
- If the remote PE did not include the address list TLV in the received label mapping message, the local router will not send any notification messages containing the address list TLV during the lifetime of the IP pseudowire.

2.5.3.1.3 VLL PPP/IPCP and Cisco-HDLC SAP Procedures

The procedures are similar to the case of an ATM SAP. The remote CE address can only be learned in the case of a PPP SAP but is not sent in the address list TLV to the remote PE in both PPP and Cisco-HDLC SAP cases.

2.5.4 IPv6 Support on IP Interworking VLL

The 7450 ESS, 7750 SR, and 7950 XRS nodes support both the transport of IPv6 packets and the interworking of IPv6 Neighbor discovery/solicitation messages on an IP Interworking VLL. IPv6 capability is enabled on an Ipipe using the **ce-address-discovery ipv6** command in the CLI.

2.5.4.1 IPv6 Datapath Operation

The IPv6 Datapath operation uses ICMPv6 extensions to automatically resolve IP address and link address associations. These are IP packets, as compared to ARP and invARP in IPv4, which are separate protocols and not based on IP packets. Manual configuration of IPv6 addresses is not supported on the IP Interworking VLL.

Each PE device intercepts ICMPv6 Neighbor Discovery (RFC 2461) packets, whether received over the SAP or over the pseudowire. The device inspects the packets to learn IPv6 interface addresses and CE link-layer addresses, modifies these packets as required according to the SAP type, then forwards them toward the original destination. The PE is also capable of generating packets to interwork between CEs (by using IPv6 Neighbor Discovery) and CEs that use other neighbor discovery protocols to bring up the link; for example, IPv6CP for PPP.

The PE device learns the IPv6 interface addresses for its directly-attached CE and other IPv6 interface addresses for the far-end CE. The PE device also learns the link-layer address of the local CE and uses it when forwarding traffic between the local and far-end CEs.

As with IPv4, the SAP accepts both unicast and multicast packets. For unicast packets, the PE checks that the MAC address/IP addresses are consistent with that in the ARP cache before forwarding; otherwise, the packet is silently discarded. Multicast packets are validated and forwarded. If more than one IP address is received per MAC address in a neighbor discovery packet, or if multiple neighbor discovery packets are received for a specific MAC address, the currently cached address is overwritten with the most recent value.

Figure 12 illustrates the data path operation for IPv6 on an IP Interworking VLL between the Ethernet and PPP (IPv6CP) SAPs.

Figure 12 **Data Path for Ethernet CE to PPP Attached CE**

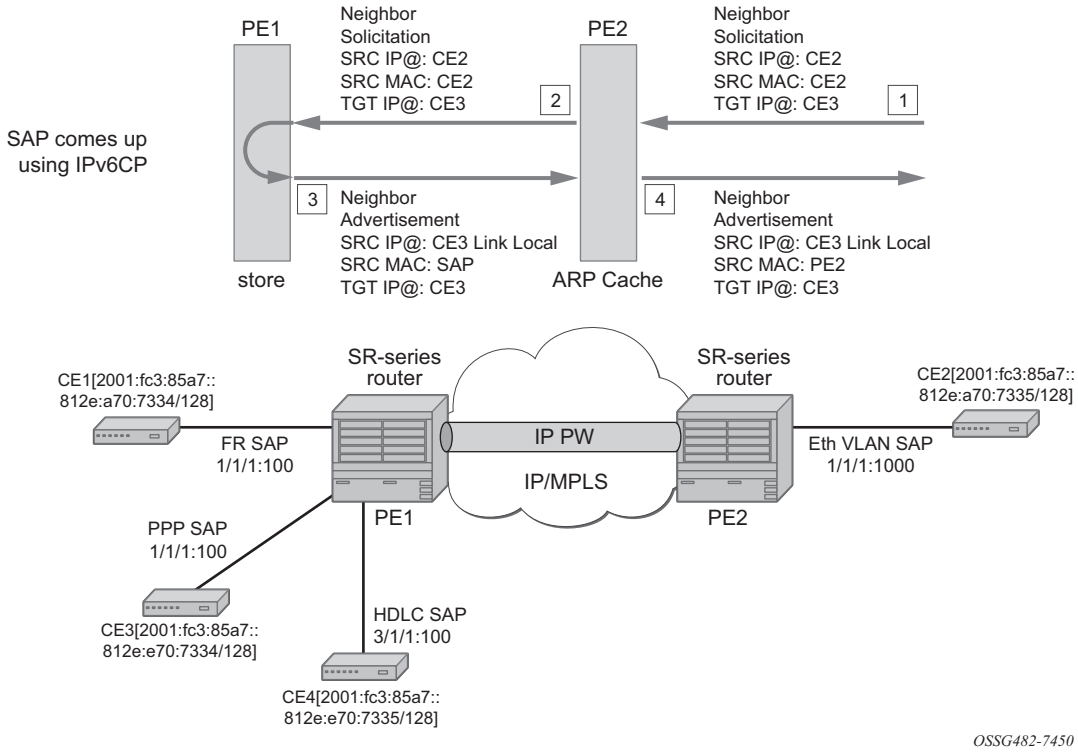

With reference to neighbor discovery between Ethernet and PPP CEs in Figure 12, the steps are as follows:

1. Ethernet-attached CE2 sends a Neighbor Solicitation message toward PE2 in order to begin the neighbor discovery process.
2. PE2 snoops this message, and the MAC address and IP address of CE2 is stored in the ARP cache of PE2 before forwarding the Neighbor Solicitation on the IP pseudowire to PE1.
3. PE1 snoops this message that arrives on the IP pseudowire and stores the IP address of the remote CE2. Since CE3 is attached to a PPP SAP, which uses IPv6CP to bring up the link, PE1 generates a neighbor advertisement message and sends it on the Ipipe toward PE2.
4. PE2 receives the neighbor advertisement on the Ipipe from PE1. It must replace the Layer 2 address in the neighbor advertisement message with the MAC address of the SAP before forwarding to CE2.

2.5.4.2 IPv6 Stack Capability Signaling

The 7750 SR, 7450 ESS, and 7950 XRS support IPv6 capability negotiation between PEs at the ends of an IP interworking VLL. Stack capability negotiation is performed if stack-capability-signaling is enabled in the CLI. Stack capability negotiation is disabled by default. Therefore, it must be assumed that the remote PE supports both IPv4 and IPv6 transport over an Ipipe.

A stack-capability sub-TLV is signaled by the two PEs using T-LDP so that they can agree on which stacks they should be using. By default, the IP pseudowire will always be capable of carrying IPv4 packets. Therefore, this capability sub-TLV is used to indicate if other stacks need to be supported concurrently with IPv4.

The stack-capability sub-TLV is a part of the interface parameters of the pseudowire FEC. This means that any change to the stack support requires that the pseudowire be torn down and re-signaled.

A PE that supports IPv6 on an IP pseudowire must signal the stack-capability sub-TLV in the initial label mapping message for the pseudowire. For the 7750 SR, 7450 ESS, and 7950 XRS, this means that the stack-capability sub-TLV must be included if both the **stack-capability-signaling** and **ce-address-discovery ipv6** options are enabled under the VLL service.

In Release 14.0, if one PE of an IP interworking VLL supports IPv6, while the far-end PE does not support IPv6 (or ce-address-discovery ipv6 is disabled), the pseudowire does not come up.

If a PE that supports IPv6 (that is, stack-capability-signaling ipv6 is enabled) has already sent an initial label mapping message for the pseudowire, but does not receive a stack-capability sub-TLV from the far-end PE in the initial label mapping message, or one is received but it is set to a reserved value, then the PE assumes that a configuration error has occurred. That is, if the remote PE did not include the stack-capability sub-TLV in the received label mapping message, or it does include the sub-TLV but with the IPv6 bit cleared, and if stack-capability-signaling is enabled, the local node with ce-address-discovery ipv6 enabled withdraws its pseudowire label with the LDP status code “IP Address type mismatch”.

If a 7750 SR, 7450 ESS, and 7950 XRS PE that supports IPv6 (that is, stack-capability-signaling ipv6 is enabled) has not yet sent a label mapping message for the pseudowire and does not receive a stack-capability sub-TLV from the far-end PE in the initial label mapping message, or one is received but it is set to a reserved value, the PE assumes that a configuration error has occurred and does not send a label mapping message of its own.

If the IPv6 stack is not supported by both PEs, or at least one of the PEs does support IPv6 but does not have the **ce-address-discovery ipv6** option selected in the CLI, IPv6 packets received from the AC are discarded by the PE. IPv4 packets are always supported.

If IPv6 stack support is implemented by both PEs, but the **ce-address-discovery ipv6** command was not enabled on both so that the IP pseudowire came up with only IPv4 support, and one PE is later toggled to **ce-address-discovery ipv6**, then that PE sends a label withdraw with the LDP status code meaning “Wrong IP Address Type” (Status Code 0x0000004B9).

If the IPv6 stack is supported by both PEs and, therefore, the pseudowire is established with IPv6 capability at both PEs, but the **ce-address-discovery ipv6** command on one PE is later toggled to **no ce-address-discovery ipv6** so that a PE ceases to support the IPv6 stack, then that PE sends a label withdraw with the LDP status code meaning “Wrong IP Address Type”.

2.6 Services Configuration for MPLS-TP

MPLS-TP PWs are supported in Epipe, Apipe, and Cpipe VLLs and Epipe spoke termination on IES/VPRN and VPLS, I-VPLS, and B-VPLS on the 7450 ESS and 7750 SR only.

This section describes how SDPs and spoke-SDPs are used with MPLS-TP LSPs and static pseudowires with MPLS-TP OAM. It also describes how to conduct test service throughput for PWs, using lock instruct messages and loopback configuration.

2.6.1 MPLS-TP SDPs

Only MPLS SDPs are supported.

An SDP used for MPLS-TP supports the configuration of an MPLS-TP identifier as the far-end address as an alternative to an IP address. IP addresses are used if IP/MPLS LSPs are used by the SDP, or if MPLS-TP tunnels are identified by IPv4 source/destination addresses. MPLS-TP node identifiers are used if MPLS-TP tunnels are used.

Only static SDPs with signaling off support MPLS-TP spoke-SDPs.

The following CLI shows the MPLS-TP options:

```
config
   service
      sdp 10 [mpls | GRE | [ldp-enabled] [create]
         signaling <off | on>
         [no] lsp <xyz>
         [no] accounting-policy <policy-id>
         [no] adv-mtu-override
         [no] booking-factor <percentage>
         [no] class-forwarding
         [no] collect-stats
         [no] description <description-string>
         [no] far-end <ip-address> | [node-id
              {<ip-address> | <0…4,294,967,295>} [global-id <global-id>]]
         [no] tunnel-far-end <ip-address>
         [no] keep-alive
         [no] mixed-lsp-mode
         [no] metric <metric>
         [no] network-domain <network-domain-name>
         [no] path-mtu <mtu>
         [no] pbb-etype <ethertype>
         [no] vlan-vc-etype <ethertype>
         [no] shutdown
```

The **far-end node-id** *ip-address* **global-id** *global-id* command is used to associate an SDP far end with an MPLS-TP tunnel whose far-end address is an MPLS-TP node ID. If the SDP is associated with an RSVP-TE LSP, the far end must be a routable IPv4 address.

The system accepts the node-id being entered in either 4-octet IP address format (a.b.c.d) or unsigned integer format.

The SDP far end refers to an MPLS-TP node-id/global-id only if:

- delivery type is MPLS
- signaling is off
- keep-alive is disabled
- mixed-lsp-mode is disabled
- adv-mtu-override is disabled

An LSP can only be allowed to be configured if the far-end information matches the lsp far end information (whether MPLS-TP or RSVP).

- Only one LSP is allowed if the far end is an MPLS-TP node-id/global-id.
- MPLS-TP or RSVP-TE LSPs are supported. However, LDP and BG LSPs are not blocked in CLI.

Signaling LDP or BGP is blocked if:

- far-end node-id/global-id is configured
- control-channel-status is enabled on any spoke (or mate vc-switched spoke)
- pw-path-id is configured on any spoke (or mate vc-switched spoke)
- IES/VPRN interface spoke control-word is enabled

The following commands are blocked if a far-end node-id/global-id is configured:

- class-forwarding
- tunnel-far-end
- mixed-lsp-mode
- keep-alive
- ldp or bgp-tunnel
- adv-mtu-override

2.6.2 VLL Spoke SDP Configuration

The system can be a T-PE or an S-PE for a pseudowire (a spoke-SDP) supporting MPLS-TP OAM. MPLS-TP related commands are applicable to spoke-SDPs configured under all services supported by MPLS-TP pseudowires. All commands and functions that are applicable to spoke-SDPs are supported, except for those that explicitly depend on T-LDP signaling of the pseudo-wire, or as stated following. Likewise, all existing functions on a specified service SAP are supported if the spoke-SDP that it is mated to is MPLS-TP.

vc-switching is supported.

The following describes how to configure MPLS-TP on an Epipe VLL. However, a similar configuration applies to other VLL types.

A spoke-SDP bound to an SDP with the **mpls-tp** keyword cannot be **no shutdown** unless the ingress label, the egress label, the control word, and the pw-path-id are configured, as follows:

```
config
   service
      epipe
         [no] spoke-sdp sdp-id[:vc-id]
              [no] hash-label
              [no] standby-signaling-slave

         [no] spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
             [create] [vc-switching] [no-endpoint | {endpoint [icb]}]
            egress
               vc-label <out-label>
            ingress
               vc-label <in-label>
            control-word
            bandwidth <bandwidth>
            [no] pw-path-id
               agi <agi>
               saii-type2 <global-id:node-id:ac-id>
               taii-type2 <global-id:node-id:ac-id>
               exit
            [no] control-channel-status
[no] refresh-timer <value>
request-timer <request-timer-secs> retry-timer <retry-timer-secs> timeout-
multiplier <multiplier>
no request-timer
               [no] acknowledgment
               [no] shutdown
               exit
```

The **pw-path-id** context is used to configure the end-to-end identifiers for an MS-PW. These may not coincide with those for the local node if the configuration is at an S-PE. The SAII and TAII are consistent with the source and destination of a label mapping message for a signaled PW.

The **control-channel-status** command enables static pseudowire status signaling. This is valid for any spoke-SDP where **signaling none** is configured on the SDP (for example, where T-LDP signaling is not in use). The refresh timer is specified in seconds, from 10-65535, with a default of 0 (off). This value can only be changed if **control-channel-status** is **shutdown**.

Commands that rely on PW status signaling are allowed if control-channel-status is configured for a spoke-SDP bound to an SDP with signaling off, but the system will use control channel status signaling rather than T-LDP status signaling. The ability to configure control channel status signaling on a specified spoke-SDP is determined by the credit-based algorithm described earlier. Control channel status for a pseudowire only counts against the credit-based algorithm if the pseudowire is in a **no shutdown** state and has a non-zero refresh timer and a non-zero request timer.

A shutdown of a service will result in the static PW status bits for the corresponding PW being set.

The spoke-SDP is held down unless the **pw-path-id** is complete.

The system will accept the node-id of the pw-path-id saii or taii being entered in either 4-octet IP address format (a.b.c.d) or unsigned integer format.

The control-word must be enabled to use MPLS-TP on a spoke-SDP.

The optional acknowledgment to a static PW status message is enabled using the **acknowledgment** command. The default is **no acknowledgment**.

The **pw-path-id** is only configurable if all of the following are true:

- in network mode D
- sdp signaling is off
- control-word is enabled (control-word is disabled by default)
- on service type Epipe, VPLS, Cpipe, or IES/VPRN interface
- An MPLS-TP node-id/global-id is configured under the **config>router>mpls>mpls-tp** context. This is required for OAM to provide a reply address.

In the vc-switching case, if configured to make a static MPLS-TP spoke SDP to another static spoke SDP, the TAII of the spoke-SDP must match the SAII of its mate, and the SAII of the spoke-SDP must match the TAII of its mate.

A control-channel-status no shutdown is allowed only if all of the following are true:

- in network-mode D
- sdp signaling is off

- control-word is enabled (control-word by default is disabled)
- the service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN interface
- pw-status-signaling is enabled (as follows)
- pw-path-id is configured for this spoke

The **hash-label** option is only configurable if SDP far end is not node-id/global-id.

The control channel status request mechanism is enabled when the **request-timer** *timer* parameter is non-zero. When enabled, this overrides the normal RFC-compliant refresh timer behavior. The refresh timer value in the status packet defined in RFC 6478 is always set to zero. The refresh-timer in the sending node is taken from the request-timer <timer1> timer. The two mechanisms are not compatible with each other. One node sends a request timer while the other is configured for refresh timer. In a specified node, the request timer can only be configured with both acknowledgment and refresh timers disabled.

When configured, the procedures following are used instead of the RFC 6478 procedures when a PW status changes.

The CLI commands to configure control channel status requests are as follows:

```
[no] control-channel-status
     [no] refresh-timer <value> //0,10-65535, default:0
     [no] request-timer <timer1> retry-timer <timer2>
                  [timeout-multiplier <value>]
     [no] shutdown
     exit
```

request-timer <timer1>: 0, 10-65535, defaults: 0.

- This parameter determines the interval at which PW status messages are sent, including a reliable delivery TLV, with the “request” bit set (as follows). This cannot be enabled if refresh-timer is not equal to zero (0).

retry-timer <timer2>: 3-60s

- This parameter determines the timeout interval if no response to a PW status is received. This defaults to zero (0) when **no retry-timer**.

timeout-multiplier <value> - 3 to 15

- If a requesting node does not get a response after retry-timer × multiplier, the node must assume that the peer is down. This defaults to zero (0) when **no retry-timer**.

2.6.2.1 Epipe VLL Spoke SDP Termination on IES, VPRN, and VPLS

All existing commands (except for those explicitly specified following) are supported for spoke-SDP termination on IES, VPRN, and VPLS (VPLS, I-VPLS and B-VPLS and routed VPLS) services. Also, the MPLS-TP commands listed preceding are supported. The syntax, default values, and functional behavior of these commands is the same as for Epipe VLLs, as specified preceding.

Also, the PW Control Word is supported on spoke-SDP termination on IES/VPRN interfaces for pseudowires of type “Ether” with statically assigned labels (signaling off) for spoke-SDPs configured with MPLS-TP Identifiers.

The following CLI commands under spoke-SDP are blocked for spoke-SDPs with statically assigned labels (and the SDP has signaling off) and MPLS-TP identifiers:

- **no status-signaling** — This command causes the spoke-SDP to fall back to using PW label withdrawal as a status signaling method. However, T-LDP is not supported on MPLS-TP SDPs. Control channel status signaling should always be used for signaling PW status. Since active/standby dual-homing into a routed VPLS requires the use of T-LDP label withdrawal as the method for status signaling, active/standby dual-homing into routed VPLS is not supported if the spoke-SDPs are MPLS-TP.
- **propagate-mac-flush** — This command requires the ability to receive MAC Flush messages using T-LDP signaling and is blocked.

2.6.3 Configuring MPLS-TP Lock Instruct and Loopback

MPLS-TP supports lock instruct and loopback for PWs.

2.6.3.1 MPLS-TP PW Lock Instruct and Loopback Overview

The lock instruct and loopback capability for MPLS-TP PWs includes the ability to:

- administratively lock a spoke-SDP with MPLS-TP identifiers
- divert traffic to and from an external device connected to a SAP
- create a data path loopback on the corresponding PW at a downstream S-PE or T-PE that was not originally bound to the spoke-SDP being tested
- forward test traffic from an external test generator into an administratively locked PW, while simultaneously blocking the forwarding of user service traffic

MPLS-TP provides the ability to conduct test service throughput for PWs, using lock instruct messages and loopback configuration. To conduct a service throughput test, you can apply an administrative lock at each end of the PW. This creates a test service that contains the SAP connected to the external device. Lock request messaging is not supported. You can also configure a MEP to send a lock instruct message to the far-end MEP. The lock instruct message is carried in a G-ACh on Channel 0x0026. A lock can be applied using the CLI or NMS. The forwarding state of the PW can be either active or standby.

After locking a PW, you can put it into loopback mode (for two-way tests) so the ingress data path in the forward direction is cross-connected to the egress data path in the reverse direction of the PW. This is accomplished by configuring the source MEP to send a loopback request to an intermediate MIP or MEP. A PW loopback is created at the PW level, so everything under the PW label is looped back. This distinguishes a PW loopback from a service loopback, where only the native service packets are looped back. The loopback is also configured through CLI or NMS.

The following MPLS-TP lock instruct and loopback functionality is supported:

- An MPLS-TP loopback can be created for an Epipe, Cpipe or Apipe VLL.
- Test traffic can be inserted at an Epipe, Cpipe or Apipe VLL endpoint or at an Epipe spoke-sdp termination on a VPLS interface.

2.6.3.2 Lock PW Endpoint Model

You can administratively lock a spoke-SDP by locking the host service using the **admin-lock** parameter of the **tools** command. The following conditions and constraints apply:

- Both ends of a PW or MS-PW represented by a spoke-SDP must be administratively locked.
- Test traffic can be injected into the spoke-SDP using a SAP defined within a test service. The test service must be identified in the **tools** command at one end of the locked PW.
- All traffic is forwarded to and from the test SAP defined in the test service, which must be of a type that is compatible with the spoke-SDP.
- Traffic to and from a non-test SAP is dropped. If no test SAP is defined, all traffic received on the spoke-SDP is dropped, and all traffic received on the paired SAP is also dropped.
- If a spoke-SDP is administratively locked, it is treated as operationally down. If a VLL SAP is paired with a spoke-SDP that is administratively locked, the SAP OAM treats this as if the spoke-SDP is operationally down.

- If a VPLS interface is paired to a spoke-SDP that is administratively locked, the L2 interface is taken down locally.
- Control-channel-status must be shutdown prior to administratively locking a spoke-SDP.

2.6.3.3 PW Redundancy and Lock Instruct and Loopback

It is possible to apply an administrative lock and loopback to one or more spoke-SDPs within a redundant set. That is, it is possible to move a spoke-SDP from an existing endpoint to a test service. When an administrative lock is applied to a spoke-SDP, it becomes operationally down and cannot send or receive traffic from the normal service SAP or spoke interface. If the lock is applied to all the spoke-SDPs in a service, all the spoke-SDPs will become operationally down.

2.6.3.4 Configuring a Test SAP for an MPLS-TP PW

A test SAP is configured under a unique test service type. This looks similar to a normal service context, but will normally only contain a SAP configuration:

```
config
   service
      epipe <service-id> [test][create]
         [no] sap <sap-id>
         [no] shutdown
       [no] shutdown
config
   service
      apipe <service-id> [vc-type {atm-vcc | atm-sdu | atm-vpc | atm-cell}
        [test][create]
         [no] sap <sap-id>
         [no] shutdown
       [no] shutdown
config
   service
      cpipe <service-id> [vc-type {satop-e1 | satop-t1 | cesopsn | cesopsncas}
        [test][create]
         [no] sap <sap-id>
         [no] shutdown
       [no] shutdown
```

You can define test SAPs appropriate to any service or PW type supported by MPLS-TP, including an Apipe, Cpipe or Epipe. The following test SAP types are supported:

- Ethernet NULL, 1q, Q-in-Q
- ATM VC, VP, VT, and so on

- TDM E1, E3, DS0, DS3, and so on

The following constraints and conditions apply:

- Up to a maximum of 16 test services can be configured per system.
- It is possible to configure access ingress and access egress QoS policies on a test SAP, as well as any other applicable SAP-specific commands and overrides.
- Vc-switching and spoke-SDP are blocked for services configured under the test context.
- The **test** keyword is mutually exclusive with vc-switching and customer.
- Valid commands under a compatible test service context do not need to be blocked just because the service is a test service.

2.6.3.5 Configuring an Administrative Lock

An administrative lock is configured on a spoke-SDP using the **admin-lock** option of the **tools perform** command, as follows:

```
tools
    perform
      service-id <svc-id>
         admin-lock
            pw
               sdp <sdp-id> admin-lock [test-svc-id <id>]
```

The following conditions and constraints apply for configuring an administrative lock:

- The lock can be configured either on a spoke-SDP that is bound to a SAP, another spoke-SDP or a VPLS interface.
- The lock is only allowed if a PW path ID is defined (for example, for static PWs with MPLS-TP identifiers).
- The lock cannot be configured on spoke-SDPs that are an Inter-Chassis Backup (ICB) or if the vc-switching keyword is present.
- The control-channel-status must be shutdown. The operator should also shutdown control-channel-status on spoke-SDPs belonging to an MS-PW at an S-PE whose far ends are administratively locked at its T-PEs. This should be enforced throughout the network management if using the 5620 SAM.
- When enabled, all traffic on the spoke-SDP is sent to and from a paired SAP that has the **test** keyword present, if such a SAP exists in the X endpoint (see Pseudowire Redundancy Service Models). Otherwise, all traffic to and from the paired SAP is dropped.

- The lock can be configured at a spoke-SDP that is bound to a VLL SAP or a VPLS interface.
- The **test-svc-id** parameter refers to the test service that should be used to inject test traffic into the service. The test service must be of a compatible type to the existing spoke-SDP under test (see Table 9).
- If the **test-svc-id** parameter is not configured on an admin-locked spoke-SDP, user traffic is blocked on the spoke-SDP.

The service manager should treat an administrative lock as a fault from the perspective of a paired SAP that is not a test SAP. This will cause the appropriate SAP OAM fault indication.

Table 9 maps supported real services to their corresponding test services.

Table 9 **Mapping of Real Services to Test Service Types**

Service	Test Service
CPIPE	CPIPE
EPIPE	EPIPE
APIPE	APIPE
VPLS	EPIPE
PBB VPLS	EPIPE

2.6.3.6 Configuring a Loopback

If a loopback is configured on a spoke-SDP, all traffic on the ingress direction of the spoke-sdp and associated with the ingress vc-label is forwarded to the egress direction of the spoke-SDP. A loopback may be configured at either a T-PE or an S-PE. It is recommended that an administrative lock is configured before configuring the loopback on a spoke-SDP. This is enforced by the NMS.

A data path loopback is configured using a **tools perform** command, as follows:

```
tools
   perform
      service-id <svc-id>
         loopback
            pw
               sdp <sdp-id>:<vc-id> {start | stop}
```

The following constraints and conditions apply for PW loopback configuration:

- The spoke-SDP cannot be an ICB or be bound to a VPLS interface.
- A PW path ID must be configured, that is, the spoke-SDP must be static and use MPLS-TP identifiers.
- The spoke-SDP must be bound to a VLL mate SAP or another spoke-SDP that is not an ICB.
- The control-channel-status must be shutdown.
- The following are disabled on a spoke-SDP for which a loopback is configured:
  - Filters
  - PW shaping
- Only network port QoS is supported.

2.6.4 Switching Static MPLS-TP to Dynamic T-LDP Signaled PWs

Some use cases for MPLS-TP require an MPLS-TP based aggregation network and an IP-based core network to interoperate, so providing the seamless transport of packet services across static MPLS-TP and dynamically signaled domains using an MS-PW. In this environment, end-to-end VCCV Ping and VCCV Trace may be used on the MS-PW, as illustrated in Figure 13.

Figure 13 **Static - Dynamic PW Switching with MPLS-TP**

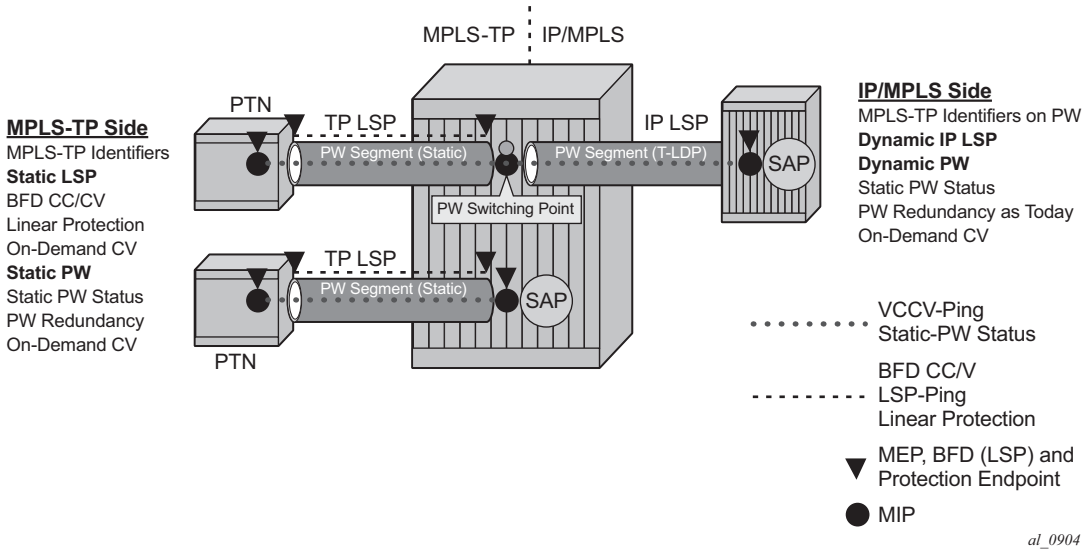

Services are backhauled from the static MPLS-TP network on the left to the dynamic IP/MPLS network on the right. The router acts as an S-PE interconnecting the static and dynamic domains.

The router implementation supports such use cases through the ability to mate a static MPLS-TP spoke SDP, with a defined *pw-path-id*, to a FEC128 spoke SDP. The dynamically signaled spoke SDP must be MPLS; GRE PWs are not supported, but the T-LDP signaled PW can use any supported MPLS tunnel type (for example, LDP, RSVP-TE, static, BGP). The control-word must be enabled on both mate spoke SDPs.

Mapping of control channel status signaling to and from T-LDP status signaling at the router S-PE is also supported.

The use of VCCV Ping and VCCV Trace on an MS-PW composed of a mix of static MPLS-TP and dynamic FEC128 segments is described in more detail in the 7450 ESS, 7750 SR, and 7950 XRS OAM and Diagnostics Guide.

2.7 VCCV BFD support for VLL, Spoke-SDP Termination on IES and VPRN, and VPLS Services

This section provides information about VCCV BFD support for VLL, spoke-SDP Termination on IES and VPRN, and VPLS Services. VCCV BFD is supported on the 7450 ESS and 7750 SR only.

2.7.1 VCCV BFD Support

The SR OS supports RFC 5885, which specifies a method for carrying BFD in a pseudowire-associated channel. This enables BFD to monitor the pseudowire between its terminating PEs, regardless of how many P routers or switching PEs the pseudowire may traverse. This makes it possible for faults that are local to individual pseudowires to be detected, whether or not they also affect forwarding for other pseudowires, LSPs, or IP packets. VCCV BFD is ideal for monitoring specific high-value services, where detecting forwarding failures (and potentially restoring from them) in the minimal amount of time is critical.

VCCV BFD is supported on VLL services using T-LDP spoke-SDPs or BGP VPWS. It is supported for Apipe, Cpipe, Epipe, Fpipe, and Ipipe VLL services.

VCCV BFD is supported on IES/VPRN services with T-LDP spoke -SDP termination (for Epipes and Ipipes).

VCCV BFD is supported on LDP- and BGP-signaled pseudowires, and on pseudowires with statically configured labels, whether signaling is off or on for the SDP. VCCV BFD is not supported on MPLS-TP pseudowires.

VCCV BFD is supported on VPLS services (both spoke-SDPs and mesh SDPs). VCCV BFD is configured by:

- configuring generic BFD session parameters in a BFD template
- applying the BFD template to a spoke-SDP or pseudowire-template binding, using the **bfd-template** *template_name* command
- enabling the template on that spoke-SDP, mesh SDP, or pseudowire-template binding using the **bfd-enable** command

2.7.2 VCCV BFD Encapsulation on a Pseudowire

The SR OS supports IP/UDP encapsulation for BFD. With this encapsulation type, the UDP headers are included on the BFD packet. IP/UDP encapsulation is supported for pseudowires that use router alert (VCCV Type 2), and for pseudowires with a control word (VCCV Type 1). In the control word case, the IPv4 channel (channel type 0x0021) is used. On the node, the destination IPv4 address is fixed at 127.0.0.1 and the source address is 127.0.0.2.

VCCV BFD sessions run end-to-end on a switched pseudowire. They do not terminate on an intermediate S-PE; therefore, the TTL of the pseudowire label on VCCV BFD packets is always set to 255 to ensure that the packets reach the far-end T-PE of an MS-PW.

2.7.3 BFD Session Operation

BFD packets flow along the full length of a PW, from T-PE to T-PE. Since they are not intercepted at an S-PE, single-hop initialization procedures are used.

A single BFD session exists per pseudowire.

BFD runs in asynchronous mode.

BFD operates as a simple connectivity check on a pseudowire. The BFD session state is reflected in the MIBs and in the **show>service id>sdp>vccv-bfd session** command. Therefore, BFD operates in a similar manner to other proactive OAM tools, such as SAA with VCCV Ping. BFD is not used to change the operation state of the pseudowire or to modify pseudowire redundancy. Mapping the BFD state to SAP OAM is not supported.

VCCV BFD runs in software with a minimum supported timer interval of 1 s.

BFD is only used for fault detection. While RFC 5885 provides a mode in which VCCV BFD can be used to signal pseudowire status, this mode is only applicable for pseudowires that have no other status signaling mechanism in use. LDP status and static pseudowire status signaling always take precedence over BFD-signaled PW status, and BFD-signaled pseudowire status is not used on pseudowires that use LDP status or static pseudowire status signaling mechanisms.

2.7.4 Configuring VCCV BFD

Generic BFD session parameters are configured for VCCV using the **bfd-template** command, in the **config**>**router**>**bfd** context. However, there are some restrictions.

For VCCV, the BFD session cannot terminate on the CPM network processor. Therefore, an error is generated if the user tries to bind a BFD template using the **type cpm-np** command within the **config**>**router**>**bfd**>**bfd-template** context.

As well, the minimum supported value for the **transmit-interval** and **receive-interval** commands when BFD is used for VCCV-BFD is 1s. Attempting to bind a BFD template with any unsupported transmit or receive interval will generate an error.

Finally, attempting to commit changes to a BFD template that is already bound to a pseudowire where the new values are invalid for VCCV BFD will result in an error.

If the preceding BFD timer values are changed in a specified template, any BFD sessions on pseudowires to which that template is bound will try to renegotiate their timers to the new values.

Commands within the BFD-template use a **begin-commit** model. To edit any value within the BFD template, a **begin** command needs to be executed after the template context has been entered. However, a value will still be stored temporarily in the template-module until the **commit** command is issued. When the **commit** is issued, values will be used by other modules such as the MPLS-TP module and BFD module.

For pseudowires where the pseudowire template does not apply, a named BFD template is configured on the spoke-SDP using the **config service** [**epipe** | **cpipe** | **apipe** | **fpipe** | **ipipe**] **spoke-sdp bfd-template** *name* command, then enabled using the **config service** [**epipe** | **cpipe** | **apipe** | **fpipe** | **ipipe**] **spoke-sdp bfd-enable** command. For example, LDP-signaled spoke-SDPs for a VLL service that uses the pseudowire ID FEC (FEC128) or spoke-SDPs with static pseudowire labels with or without MPLS-TP identifiers.

Configuring and enabling a BFD template on a static pseudowire already configured with MPLS-TP identifiers (that is, with a pw-path-id) or on a spoke-SDP with a configured pw-path-id is not supported. Likewise, if a BFD template is configured and enabled on a spoke-SDP, a pw-path-id cannot be configured on the spoke-SDP.

The **bfd-enable** command is blocked on a spoke-SDP configured with VC-switching. This is because VCCV BFD always operates end-to-end on an MS-pseudowire. It is not possible to extract VCCV BFD packets at the S-PE.

For IES and VPRN spoke-SDP termination where the pseudowire template does not apply (that is, where the spoke-SDP is signaled with LDP and uses the pseudowire ID FEC (FEC128)), the BFD template is configured using the **config service ies** | **vprn if spoke-sdp bfd-template** *name* command, then enabled using the **config service ies** | **vprn if spoke-sdp bfd-enable** command.

For H-VPLS, where the pseudo-wire template does not apply (that is, LDP-VPLS spoke and mesh SDPs that use the pseudo-wire ID FEC(FEC128)) the BFD template is configured using the **config service vpls spoke-sdp bfd-name** *name* command or the **config service vpls mesh-sdp bfd-name** *name* command.VCCV BFD is then enabled with the bfd-enable command under the VPLS spoke-SDP or mesh-SDP context.

Pseudo-wires where the pseudo-wire template does apply and that support VCCV BFD are as follows:

- BGP-AD, which is signaled using the Generalized pseudowire ID FEC (FEC129) with Attachment Individual Identifier (AII) type I
- BGP VPLS
- BGP VPWS

For these pseudowire types, a named BFD template is configured and enabled from the pseudowire template binding context.

For BGP VPWS, the BFD template is configured using the **config service epipe bgp pw-template-binding bfd-template** *name* command, then enabled using the **config service epipe bgp pw-template-binding bfd-enable** command.

2.8 Pseudowire Switching

The pseudowire switching feature provides the user with the ability to create a VLL service by cross-connecting two spoke-SDPs. This feature allows the scaling of VLL and VPLS services in a large network in which the otherwise full mesh of PE devices would require thousands of Targeted LDP (T-LDP) sessions per PE node.

Services with one SAP and one spoke-SDP are created normally on the PE; however, the target destination of the SDP is the pseudowire switching node instead of what is normally the remote PE. Also, the user configures a VLL service on the pseudowire switching node using the two SDPs.

The pseudowire switching node acts in a passive role with respect to signaling of the pseudowires. It waits until one or both of the PEs sends the label mapping message before relaying it to the other PE. This is because it needs to pass the interface parameters of each PE to the other.

A pseudowire switching point TLV is inserted by the switching pseudowire to record its system address when relaying the label mapping message. This TLV is useful in a few situations:

- It allows for troubleshooting of the path of the pseudowire especially if multiple pseudowire switching points exist between the two PEs.
- It helps in loop detection of the T-LDP signaling messages where a switching point would receive back a label mapping message it had already relayed.
- The switching point TLV is inserted in pseudowire status notification messages when they are sent end-to-end or from a pseudowire switching node toward a destination PE.

Pseudowire OAM is supported for the manual switching pseudowires and allows the pseudowire switching node to relay end-to-end pseudowire status notification messages between the two PEs. The pseudowire switching node can generate a pseudowire status and send it to one or both of the PEs by including its system address in the pseudowire switching point TLV. This allows a PE to identify the origin of the pseudowire status notification message.

In the following example, the user configures a regular Epipe VLL service PE1 and PE2. These services each consist of a SAP and a spoke-SDP. However, the target destination of the SDP is not the remote PE, but the pseudowire switching node. Also, the user configures an Epipe VLL service on the pseudowire switching node using the two SDPs.

```
|7450 ESS, 7750 SR, and 7950 XRS PE1 (Epipe)|---sdp 2:10---|7450 ESS, 7750 SR, and
7950 XRS PW SW (Epipe)|---sdp 7:15---|7450 ESS, 7750 SR, and 7950 XRS PE2 (Epipe)|
```

Configuration examples are in Configuring Two VLL Paths Terminating on T-PE2.

2.8.1 Pseudowire Switching with Protection

Pseudowire switching scales VLL and VPLS services over a multi-area network by removing the need for a full mesh of targeted LDP sessions between PE nodes. Figure 14 illustrates the use of pseudowire redundancy to provide a scalable and resilient VLL service across multiple IGP areas in a provider network.

In the network in Figure 14, PE nodes act as masters and pseudowire switching nodes act as slaves for the purpose of pseudowire signaling. A switching node will need to pass the SAP interface parameters of each PE to the other PEs. T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node; for example, S-PE1. The label mapping message will include the SAP interface parameters, such as MTU, in the label mapping message. S-PE1 checks the FEC against the local information and, if a match exists, appends the optional pseudowire switching point TLV to the FEC TLV in which it records its system address. T-PE1 then relays the label mapping message to S-PE2. S-PE2 performs similar operations and forwards a label mapping message to T-PE2.

The same procedures are followed for the label mapping message in the reverse direction; for example, from T-PE2 to T-PE1. S-PE1 and S-PE2 will make the spoke-SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

Figure 14 **VLL Resilience with Pseudowire Redundancy and Switching**

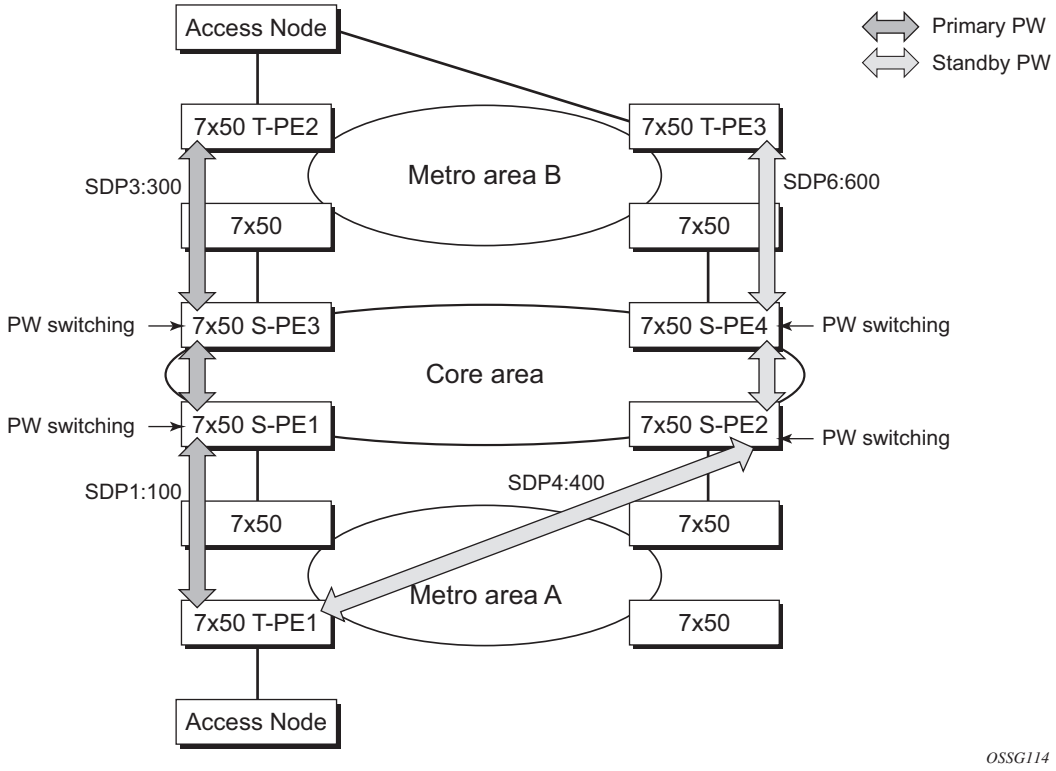

The pseudowire switching TLV is useful in a few situations. First, it allows for troubleshooting of the path of the pseudowire, especially if multiple pseudowire switching points exist between the two T-PE nodes. Second, it helps in loop detection of the T-LDP signaling messages where a switching point receives back a label mapping message that the point already relayed. Finally, it can be inserted in pseudowire status messages when they are sent from a pseudowire switching node toward a destination PE.

Pseudowire status messages can be generated by the T-PE nodes and/or the S-PE nodes. Pseudowire status messages received by a switching node are processed, and passed on to the next hop. An S-PE node appends the optional pseudowire switching TLV, with the S-PEs system address added to it, to the FEC in the pseudowire status notification message, only if that S-PE originated the message or the message was received with the TLV in it. Otherwise, the message was originated by a T-PE node and the S-PE should process and pass the message without changes, except for the VC-ID value in the FEC TLV.

2.8.2 Pseudowire Switching Behavior

In the network in Figure 14, PE nodes act as masters and pseudowire switching nodes act as slaves for the purpose of pseudowire signaling. This is because a switching node will need to pass the SAP interface parameters of each PE to the other. T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node; for example, S-PE1. It will include the SAP interface parameters, such as MTU, in the label mapping message. S-PE1 checks the FEC against the local information and, if a match exists, appends the optional pseudowire switching point TLV to the FEC TLV in which it records its system address. T-PE1 then relays the label mapping message to S-PE2. S-PE2 performs similar operation and forwards a label mapping message to T-PE2.

The same procedures are followed for the label mapping message in the reverse direction; for example, from T-PE2 to T-PE1. S-PE1 and S-PE2 will effect the spoke-SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

Pseudowire status messages can be generated by the T-PE nodes and/or the S-PE nodes. Pseudowire status messages received by a switching node are processed, then passed on to the next hop. An S-PE node appends the optional pseudowire switching TLV, with its system address added to it, to the FEC in the pseudowire status notification message, only if it originated the message or the message was received with the TLV in it. Otherwise, the message was originated by a T-PE node and the S-PE should process and pass the message without changes, except for the VC-ID value in the FEC TLV.

The merging of the received T-LDP status notification message and the local status for the spoke-SDPs from the service manager at a PE complies with the following rules:

- When the local status for both spoke-SDPs is up, the S-PE passes any received SAP or SDP binding generated status notification message unchanged; for example, the status notification TLV is unchanged but the VC-ID in the FEC TLV is set to value of the pseudowire segment to the next hop.
- When the local operational status for any of the spokes is down, the S-PE always sends an SDP-binding down status bits regardless of whether the received status bits from the remote node indicated SAP up or down or SDP-binding up or down.

2.8.2.1 Pseudowire Switching TLV

The format of the pseudowire switching TLV is as follows:

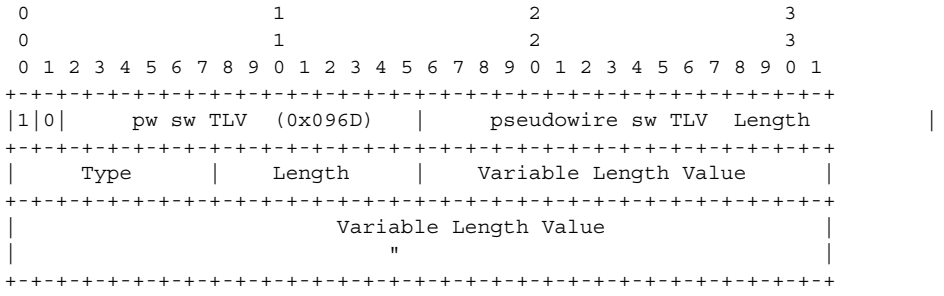

PW sw TLV Length — Specifies the total length of all the following pseudowire switching point TLV fields in octets

Type — Encodes how the Value field is to be interpreted

Length — Specifies the length of the Value field in octets

Value — Octet string of Length octets that encodes information to be interpreted as specified by the Type field

2.8.2.2 Pseudowire Switching Point Sub-TLVs

Following is information specific to pseudowire switching point sub-TLVs:

- Pseudowire ID of last pseudowire segment traversed — Sub-TLV type that contains a pseudowire ID in the format of the pseudowire ID
- Pseudowire switching point description string — An optional description text string of up to 80 characters
- IP address of pseudowire switching point — An options sub-TLV; IP V4 or V6 address of the pseudowire switching point.
- MH VCCV capability indication

2.8.3 Static-to-Dynamic Pseudowire Switching

When one segment of the pseudowire cross-connect at the S-PE is static while the other is signaled using T-LDP, the S-PE operates much like a T-PE from a signaling perspective and as an S-PE from a data plane perspective.

The S-PE signals a label mapping message as soon as the local configuration is complete. The control word C-bit field in the pseudowire FEC is set to the value configured on the static spoke-SDP.

When the label mapping for the egress direction is also received from the T-LDP peer, and the information in the FEC matches that of the local configuration, the static-to-dynamic cross-connect is created.

It is possible that end nodes of a static pseudowire segment can be misconfigured. In this case, an S-PE or T-PE node may be receiving packets with the wrong encapsulation, so that an invalid payload could be forwarded over the pseudowire or the SAP, respectively. Also, if the S-PE or T-PE node is expecting the control word in the packet encapsulation and the received packet comes with no control word, but the first nibble below the label stack is 0x0001, the packet may be mistaken for a VCCV OAM packet and may be forwarded to the CPM. In that case, the CPM will perform a check of the IP header fields such as version, IP header length, and checksum. If any of these fail the VCCV packet will be discarded.

2.8.4 Ingress VLAN Swapping

This feature is supported on VPLS and VLL services where the end-to-end solution is built using two node solutions (requiring SDP connections between the nodes).

In VLAN swapping, only the vlan-id value is copied to the inner VLAN position. The Ethertype of the inner tag will be preserved and all consecutive nodes will work with that value. Similarly, the dot1p bits value of outer-tag will not be preserved.

Figure 15 describes a network where, at user-access side (DSLAM-facing SAPs), every subscriber is represented by several QinQ SAPs with inner-tag encoding service and outer-tag encoding subscriber (DSL line). At the aggregation side (BRAS- or PE-facing SAPs) every subscriber is represented by DSL line number (inner VLAN tag) and DSLAM (outer VLAN tag). The effective operation on the VLAN tag is to drop the inner tag at the access side and push another tag at the aggregation side.

Figure 15 **Ingress VLAN Swapping**

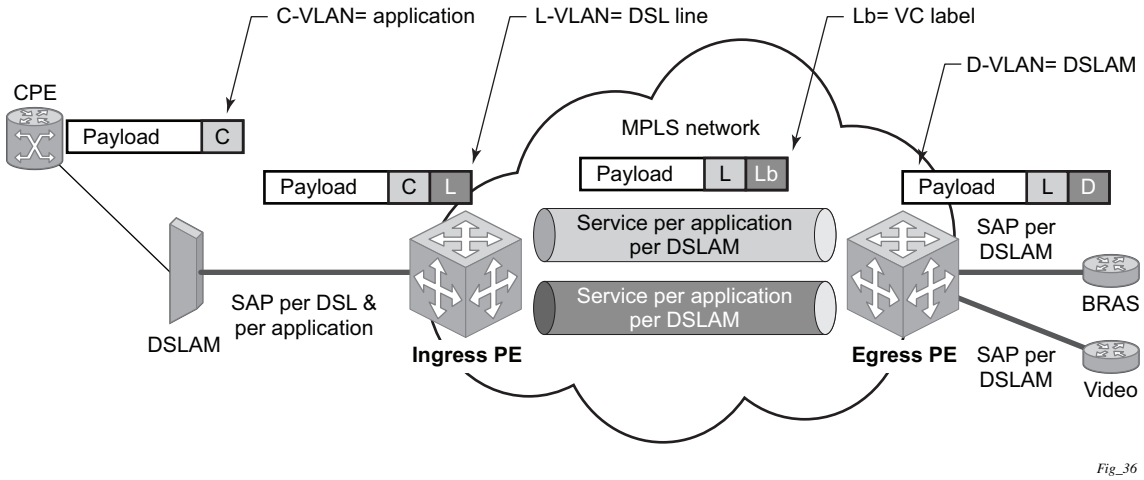

2.8.4.1 Ingress VLAN Translation

Figure 16 indicates an application where different circuits are aggregated in the VPLS-based network.The access side is represented by an explicit do1q encapsulated SAP. Because the VLAN ID is port specific, those connected to different ports might have the same VLAN. The aggregation side is aggregated on the same port; therefore, a unique VLAN ID is required.

Figure 16 **Ingress VLAN Translation**

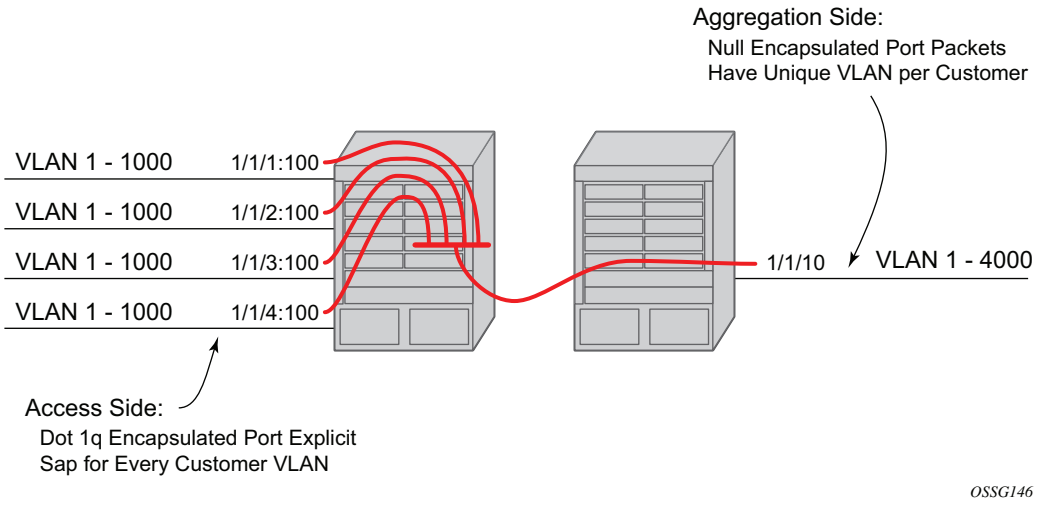

2.8.5 Pseudowire Redundancy

Pseudowire redundancy provides the ability to protect a pseudowire with a pre-provisioned secondary standby pseudowire and to switch traffic over to that secondary standby pseudowire in case of a SAP and/or network failure condition. Normally, pseudowires are redundant by the virtue of the SDP redundancy mechanism. For instance, if the SDP is an RSVP LSP and is protected by a secondary standby path and/or by Fast-Reroute paths (FRR), the pseudowire is also protected. However, there are two applications in which SDP redundancy does not protect the end-to-end pseudowire path:

- There are two different destination PE nodes for the same VLL service. The main use case is the provision of dual-homing of a CPE or access node to two PE nodes located in different POPs. The other use case is the provision of a pair of active and standby BRAS nodes, or active and standby links to the same BRAS node, to provide service resiliency to broadband service subscribers.
- The pseudowire path is switched in the middle of the network and the pseudowire switching node fails.

Pseudowire and VPLS link redundancy extends link-level resiliency for pseudowires and VPLS to protect critical network paths against physical link or node failures. These innovations enable the virtualization of redundant paths across the metro or core IP network to provide seamless and transparent fail-over for point-to-point and multi-point connections and services. When deployed with multi-chassis LAG, the path for return traffic is maintained through the pseudowire or VPLS switchover, which enables carriers to deliver “always on” services across their IP/MPLS networks.

2.8.6 Dynamic Multi-Segment Pseudowire Routing

2.8.6.1 Overview

Dynamic Multi-Segment Pseudowire Routing (Dynamic MS-PWs) enable a complete multi-segment pseudowire to be established, while only requiring per-pseudowire configuration on the T-PEs. No per-pseudowire configuration is required on the S-PEs. End-to-end signaling of the MS-PW is achieved using T-LDP, while multi-protocol BGP is used to advertise the T-PEs, allowing dynamic routing of the MS-PW through the intervening network of S-PEs. Dynamic multi-segment pseudowires are described in the IETF Draft *draft-ietf-pwe3-dynamic-ms-pw-13.txt*.

Figure 17 illustrates the operation of dynamic MS-PWs.

Figure 17 **Dynamic MS-PW Overview**

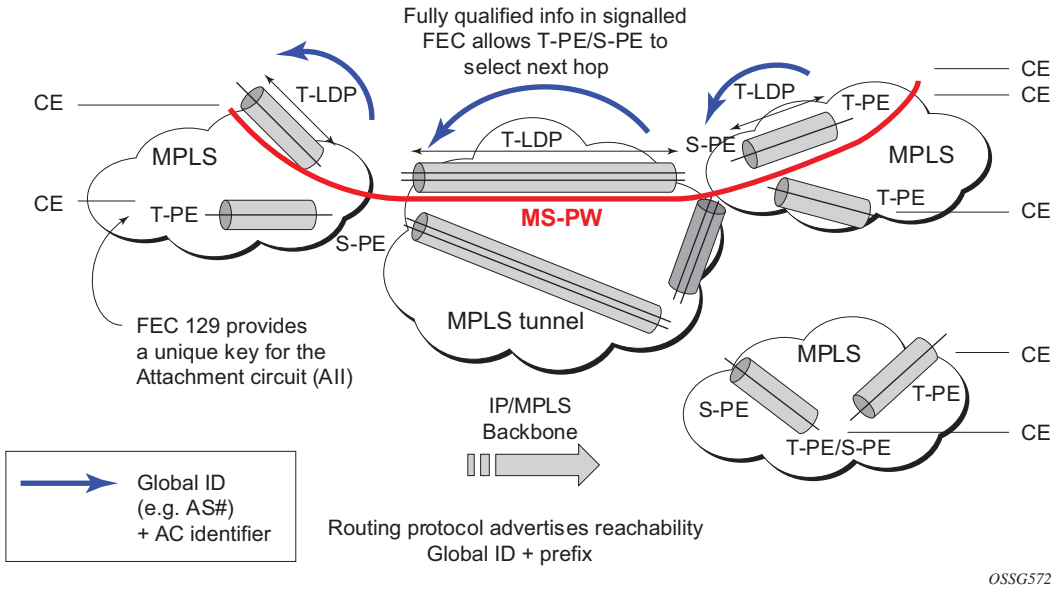

The FEC 129 AII Type 2 structure depicted in Figure 18 is used to identify each individual pseudowire endpoint:

Figure 18 **MS-PW Addressing using FEC129 AII Type 2**

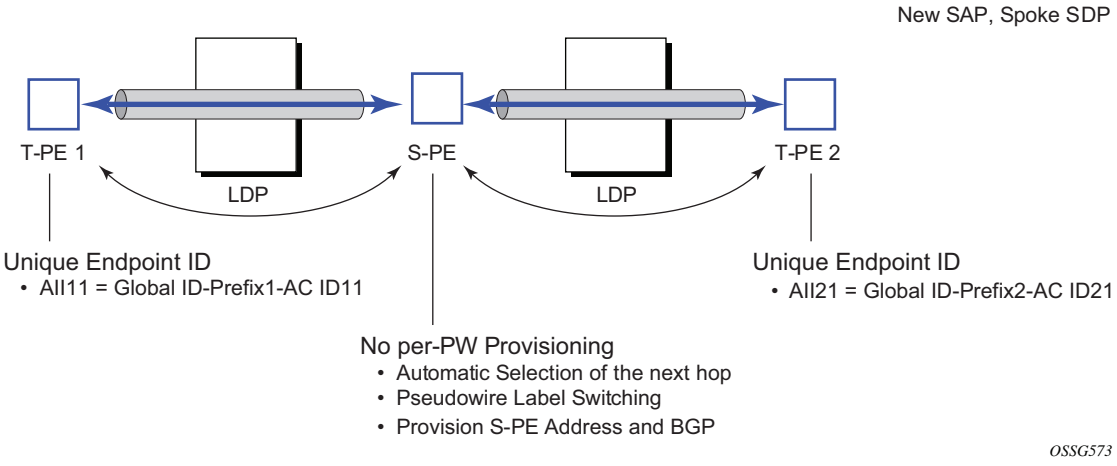

A 4-byte global-id followed by a 4-byte prefix and a 4-byte attachment circuit ID are used to provide for hierarchical, independent allocation of addresses on a per-service provider network basis. The first 8 bytes (global-id + prefix) may be used to identify each individual T-PE or S-PE as a loopback Layer 2 address.

The AII type is mapped into the MS-PW BGP NLRI (a BGP AFI of L2VPN, and SAFI for network layer reachability information for dynamic MS-PWs). As soon as a new T- PE is configured with a local prefix address of global id: prefix, pseudowire routing will proceed to advertise this new address to all the other T- PEs and S-PEs in the network, as depicted in Figure 19.

Figure 19 **Advertisement of PE Addresses by PW Routing**

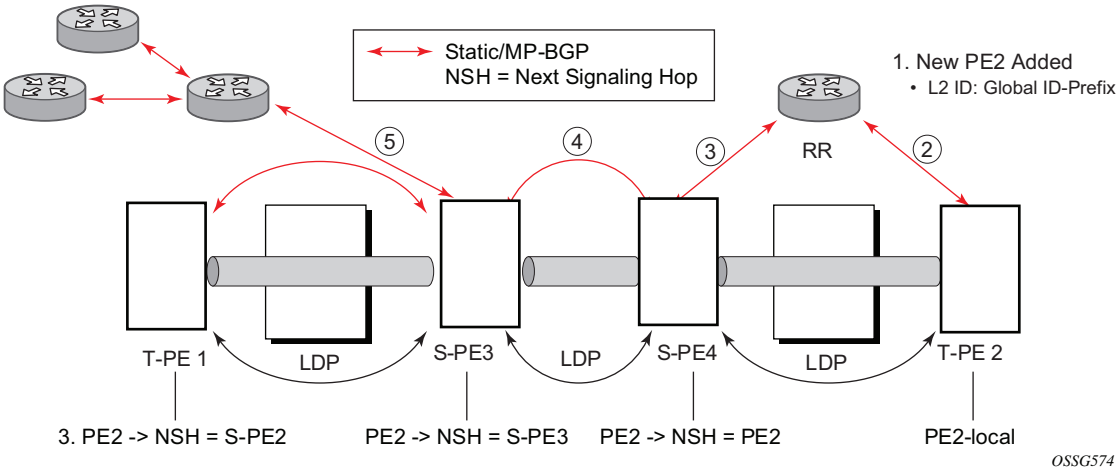

In step 1 of Figure 19, a new T-PE (T-PE2) is configured with a local prefix.

Next, in steps 2 to 5, MP-BGP will use the NLRI for the MS-PW routing SAFI to advertise the location of the new T-PE to all the other PEs in the network. Alternatively, static routes may be configured on a per T-PE/S-PE basis to accommodate non-BGP PEs in the solution.

As a result, pseudowire routing tables for all the S-PEs and remote T-PEs are populated with the next hop to be used to reach T-PE2.

VLL services can then be established, as illustrated in Figure 20.

Figure 20 **Signaling of Dynamic MS-PWs using T-LDP**

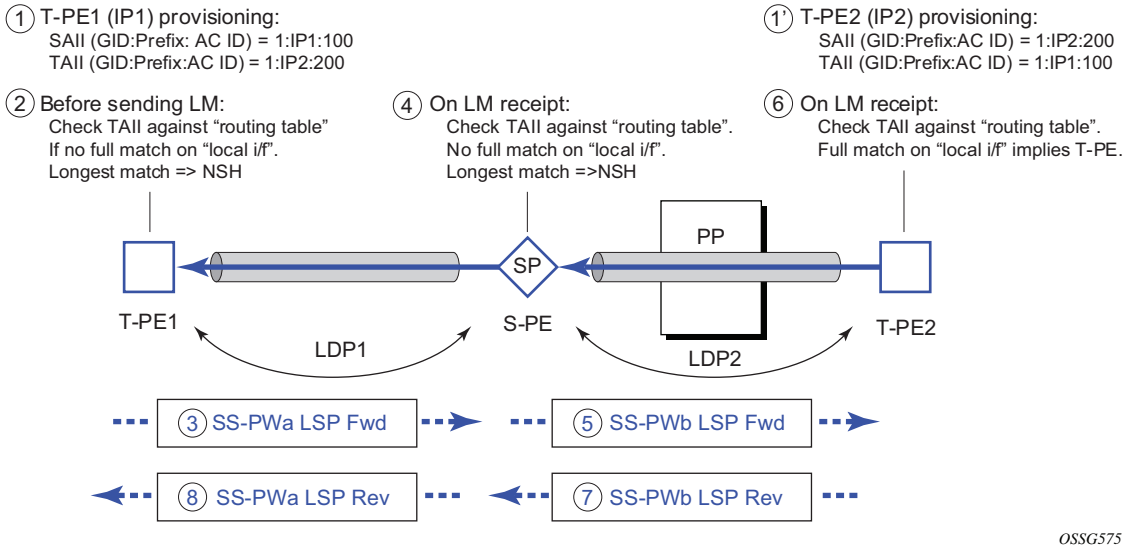

In step 1 and 1' of Figure 20 the T-PEs are configured with the local and remote endpoint information: Source AII (SAII) and Target AII (TAII). On the router, the AIIs are locally configured for each spoke-SDP, according to the model shown in Figure 21. Therefore the router provides for a flexible mapping of the AII to SAP. That is, the values used for the AII are through local configuration, and it is the context of the spoke-SDP that binds it to a specific SAP.

Figure 21 **Mapping of AII to SAP**

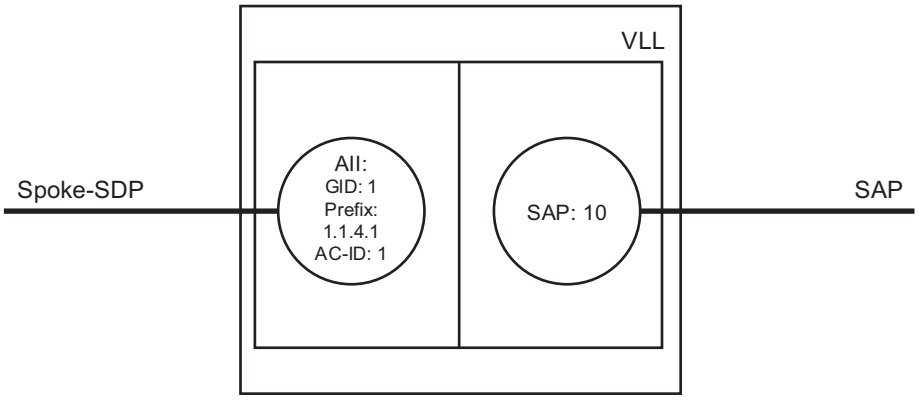

Before T-LDP signaling starts, the two T-PEs decide on an active and passive relationship using the highest AII (comparing the configured SAII and TAII) or the configured precedence. Next, the active T-PE (in the IETF draft, this is referred to as the source T-PE or ST-PE) checks the PW routing table to determine the next signaling hop for the configured TAII using the longest match between the TAII and the entries in the PW routing table.

This signaling hop is then used to choose the T-LDP session to the chosen next-hop S-PE. Signaling proceeds through each subsequent S-PE using similar matching procedures to determine the next signaling hop. Otherwise, if a subsequent S-PE does not support dynamic MS-PW routing, so uses a statically configured PW segment, the signaling of individual segments follows the procedures already implemented in the PW Switching feature.

BGP can install a PW AII route in the PW routing table with ECMP next-hops. However, when LDP needs to signal a PW with matching TAII, it will choose only one next-hop from the available ECMP next-hops. PW routing supports up to 4 ECMP paths for each destination.

The signaling of the forward path ends when the PE matches the TAII in the label mapping message with the SAII of a spoke-SDP bound to a local SAP. The signaling in the reverse direction can now be initiated, which follows the entries installed in the forward path. The PW routing tables are not consulted for the reverse path. This ensures that the reverse direction of the PW follows exactly the same set of S-PEs as the forward direction.

This solution can be used in either a MAN-WAN environment or in an Inter-AS/Inter-Provider environment as depicted in Figure 22.

Figure 22 **VLL Using Dynamic MS-PWs, Inter-AS Scenario**

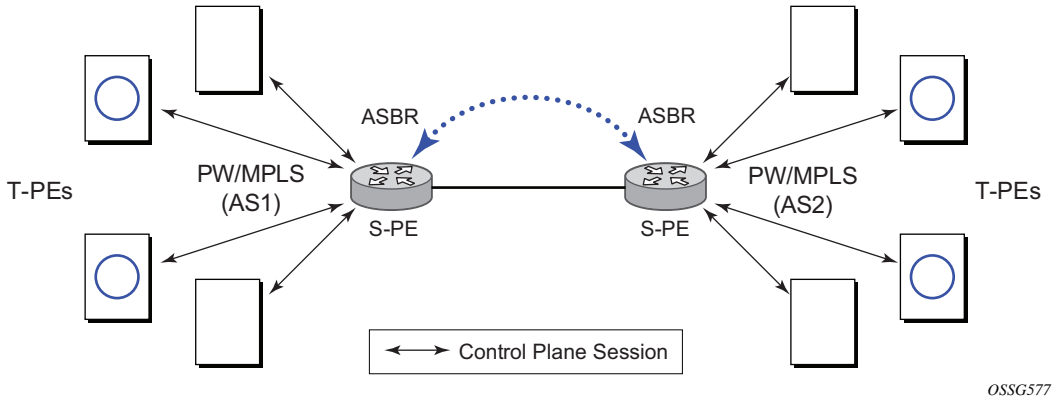

Data plane forwarding at the S-PEs uses pseudowire service label switching, as per the pseudowire switching feature.

2.8.6.2 Pseudowire Routing

Each S-PE and T-PE has a pseudowire routing table that contains a reference to the T-LDP session to use to signal to a set of next hop S-PEs to reach a specific T-PE (or the T-PE if that is the next hop). For VLLs, this table contains aggregated AII Type 2 FECs and may be populated with routes that are learned through MP-BGP or that are statically configured.

MP-BGP is used to automatically distribute T-PE prefixes using the new MS-PW NLRI, or static routes can be used. The MS-PW NLRI is composed of a Length, an 8-byte route distinguisher (RD), a 4-byte global-id, a 4-byte local prefix, and (optionally) a 4-byte AC-ID. Support for the MS-PW address family is configured in CLI under the **config**>**router**>**bgp**>**family ms-pw** context.

MS-PW routing parameters are configured in the **config**>**service**>**pw-routing** context.

To enable support for dynamic MS-PWs on a 7750 SR, 7450 ESS, or 7950 XRS node to be used as a T-PE or S-PE, a single, globally unique, S-PE ID, known as the S-PE address, is first configured under **config>service>pw-routing** on each node to be used as a T-PE or S-PE. The S-PE address has the format global-id:prefix. It is not possible to configure any local prefixes used for pseudowire routing or to configure spoke-SDPs using dynamic MS-PWs at a T-PE unless an S-PE address has already been configured. The S-PE address is used as the address of a node used to populate the switching point TLV in the LDP label mapping message and the pseudowire status notification sent for faults at an S-PE.

Each T-PE is also configured with the following parameters:

- Global-id — This is a 4-byte identifier that uniquely identifies an operator or the local network.
- Local prefix — One or more local (Layer 2) prefixes (up to a maximum of 16), which are formatted in the style of a 4-octet IPv4 address. A local prefix identifies a T-PE or S-PE in the PW routing domain.
- For each local prefix, at least one 8-byte RD can be configured. It is also possible to configure an optional BGP community attribute.

For each local prefix, BGP then advertises each global-id/prefix tuple and unique RD and community pseudowire using the MS-PW NLRI, based on the aggregated FEC129 AII Type 2 and the Layer 2 VPN/PW routing AFI/SAFI 25/6, to each T-PE/S-PE that is a T-LDP neighbor, subject to local BGP policies.

The dynamic advertisement of each of these pseudowire routes is enabled for each prefix and RD using the **advertise-bgp** command.

An export policy is also required in order to export MS-PW routes in MP-BGP. This can be done using a default policy, such as the following:

```
*A:lin-123>config>router>policy-options# info
----------------------------------------------
            policy-statement "ms-pw"
                default-action accept
                exit
            exit
----------------------------------------------
```

However, this would export all routes. A recommended choice is to enable filtering per-family, as follows:

```
*A:lin-123>config>router>policy-options# info
----------------------------------------------
            policy-statement "to-mspw"
                entry 1
                    from
                        family ms-pw
                    exit
                    action accept
                    exit
                exit
            exit
----------------------------------------------
```

The following command is then added in the **config>router>bgp** context:

```
export "to-mspw"
```

Local-preference for IBGP and BGP communities can be configured under such a policy.

2.8.6.2.1 Static Routing

As well as support for BGP routing, static MS-PW routes may also be configured using the **config services pw-routing static-route** command. Each static route comprises the target T-PE global-id and prefix, and the IP address of the T-LDP session to the next hop S-PE or T-PE that should be used.

If a static route is set to 0, this represents the default route. If a static route exists to a specified T-PE, this default route is used in preference to any BGP route that may exist.

2.8.6.2.2 Explicit Paths

A set of default explicit routes to a remote T-PE or S-PE prefix may be configured on a T-PE under **config>services>pw-routing** using the **path name** command. Explicit paths are used to populate the explicit route TLV used by MS-PW T-LDP signaling. Only strict (fully qualified) explicit paths are supported.

It is possible to configure explicit paths independently of the configuration of BGP or static routing.

2.8.6.3 Configuring VLLs using Dynamic MS-PWs

One or more spoke-SDPs may be configured for distributed Epipe VLL services. Dynamic MS-PWs use FEC129 (also known as the Generalized ID FEC) with Attachment Individual Identifier (AII) Type 2 to identify the pseudowire, as opposed to FEC128 (also known as the PW ID FEC) used for traditional single segment pseudowires and for pseudowire switching. FEC129 spoke-SDPs are configured under the **spoke-sdp-fec** command in the CLI.

FEC129 AII Type 2 uses a Source Attachment Individual Identifier (SAII) and a Target Attachment Individual Identifier (TAII) to identify the end of a pseudowire at the T-PE. The SAII identifies the local end, while the TAII identifies the remote end. The SAII and TAII are each structured as follows:

- Global-id — This is a 4-byte identifier that uniquely identifies an operator or the local network.
- Prefix — A 4-byte prefix, which should correspond to one of the local prefixes assigned under pw-routing.
- AC-ID — A 4-byte identifier for the local end of the pseudowire. This should be locally unique within the scope of the global-id:prefix.

2.8.6.3.1 Active/Passive T-PE Selection

Dynamic MS-PWs use single-sided signaling procedures with double-sided configuration; a fully qualified FEC must be configured at both endpoints. That is, one T-PE (the source T-PE, ST-PE) of the MS-PW initiates signaling for the MS-PW, while the other end (the terminating T-PE, TT-PE) passively waits for the label mapping message from the far end. This termination end only responds with a label mapping message to set up the opposite direction of the MS-PW when it receives the label mapping from the ST-PE. By default, the router will determine which T-PE is the ST-PE (the active T-PE) and which is the TT-PE (the passive T-PE) automatically, based on comparing the SAII with the TAII as unsigned integers. The T-PE with

SAII>TAII assumes the active role. However, it is possible to override this behavior using the signaling {**master** | **auto**} command under **spoke-sdp-fec**. If master is selected at a specified T-PE, that T-PE will assume the active role. If a T-PE is at the endpoint of a spoke-SDP that is bound to an VLL SAP and single-sided auto-configuration is used (see 2.8.6.3.2), then that endpoint is always passive. Therefore, signaling master should only be used when it is known that the far end will assume a passive behavior.

2.8.6.3.2 Automatic Endpoint Configuration

Automatic endpoint configuration allows the configuration of an endpoint without specifying the TAII associated with that **spoke-sdp-fec**. It allows a single-sided provisioning model where an incoming label mapping message with a TAII that matches the SAII of that spoke-SDP is automatically bound to that endpoint. This is useful in scenarios where a service provider wants to separate service configuration from the service activation phase.

Automatic endpoint configuration is supported for Epipe VLL **spoke-sdp-fec** endpoints bound to a VLL SAP. It is configured using the **spoke-sdp-fec auto-config** command, and excluding the TAII from the configuration. When auto-configuration is used, the node assumes passive behavior from a point of view of T-LDP signaling (see 2.8.6.3.1). Therefore, the far-end T-PE must be configured as the signaling master for that **spoke-sdp-fec.**

2.8.6.3.3 Selecting a Path for an MS-PW

Path selection for signaling occurs in the outbound direction (ST-PE to TT-PE) for an MS-PW. In the TT-PE to ST-PE direction, a label mapping message follows the reverse of the path already taken by the outgoing label mapping.

A node can use explicit paths, static routes, or BGP routes to select the next hop S-PE or T-PE. The order of preference used in selecting these routes is:

1. Explicit Path
2. Static route
3. BGP route

To use an explicit path for an MS-PW, an explicit path must have been configured in the **config>services>pw-routing>path** *path-name* context. The user must then configure the corresponding **path** *path-name* under **spoke-sdp-fec**.

If an explicit path name is not configured, the TT-PE or S-PE will perform a longest match lookup for a route (static if it exists, and BGP if not) to the next hop S-PE or T-PE to reach the TAII.

Pseudowire routing chooses the MS-PW path in terms of the sequence of S-PEs to use to reach a specified T-PE. It does not select the SDP to use on each hop, which is instead determined at signaling time. When a label mapping is sent for a specified pseudowire segment, an LDP SDP will be used to reach the next-hop S-PE/T-PE if such an SDP exists. If not, and an RFC 3107 labeled BGP SDP is available, then that will be used. Otherwise, the label mapping will fail and a label release will be sent.

2.8.6.3.4 Pseudowire Templates

Dynamic MS-PWs support the use of the pseudowire template for specifying generic pseudowire parameters at the T-PE. The pseudowire template to use is configured in the **spoke-sdp-fec>pw-template-bind** *policy-id* context. Dynamic MS-PWs do not support the provisioned SDPs specified in the pseudowire template.

2.8.6.4 Pseudowire Redundancy

Pseudowire redundancy is supported on dynamic MS-PWs used for VLLs. It is configured in a similar manner to pseudowire redundancy on VLLs using FEC128, whereby each spoke-sdp-fec within an endpoint is configured with a unique SAII/ TAII.

Figure 23 illustrates the use of pseudowire redundancy.

Figure 23 **Pseudowire Redundancy**

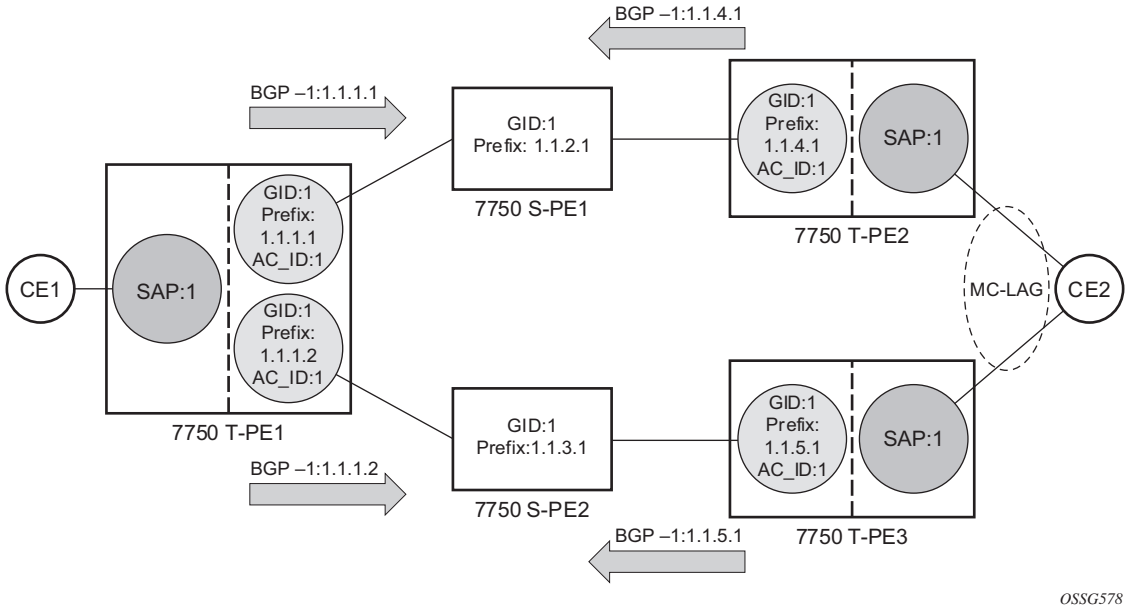

The following is a summary of the key points to consider in using pseudowire redundancy with dynamic MS-PWs:

- Each MS-PW in the redundant set must have a unique SAII/TAII set and is signaled separately. The primary pseudowire is configured in the **spoke-sdp-fec>primary** context.
- Each MS-PW in the redundant set should use a diverse path (from the point of view of the S-PEs traversed) from every other MS-PW in that set if path diversity is possible in a specific network topology. There are a number of possible ways to achieve this:
  - Configure an explicit path for each MS-PW.
  - Allow BGP routing to automatically determine diverse paths using BGP policies applied to different local prefixes assigned to the primary and standby MS-PWs.
  - Path diversity can be further provided for each primary pseudowire through the use of a BGP RD.

If the primary MS-PW fails, fail-over to a standby MS-PW occurs, as per the normal pseudowire redundancy procedures. A configurable retry timer for the failed primary MS-PW is then started. When the timer expires, attempts to reestablish the primary MS-PW using its original path occur, up to a maximum number of attempts as per the retry count parameter. On successful reestablishment the T-PE may then optionally revert back to the primary MS-PW.

Since the SDP ID is determined dynamically at signaling time, it cannot be used as a tie breaker to choose the primary MS-PW between multiple MS-PWs of the same precedence. The user should, therefore, explicitly configure the precedence values to determine which MS-PW is active in the final selection.

2.8.6.5 VCCV OAM for Dynamic MS-PWs

The primary difference between dynamic MS-PWs and those using FEC128 is support for FEC129 AII type 2. As in PW Switching, VCCV on dynamic MS-PWs requires the use of the VCCV control word on the pseudowire. Both the vccv-ping and vccv-trace commands support dynamic MS-PWs.

2.8.6.6 VCCV-Ping on Dynamic MS-PWs

VCCV-ping supports the use of FEC129 AII type 2 in the target FEC stack of the ping echo request message. The FEC to use in the echo request message is derived in one of two ways: Either the user can specify only the *spoke-sdp-fec-id* of the MS-PW in the **vccv-ping** command, or the user can explicitly specify the SAII and TAII to use.

If the SAII:TAII is entered by the user in the vccv-ping command, those values are used for the vccv-ping echo request, but their order is reversed before being sent so that they match the order for the downstream FEC element for an S-PE, or the locally configured SAII:TAII for a remote T-PE of that MS-PW. If SAII:TAII is entered as well as the *spoke-sdp-fec-id*, the system will verify the entered values against the values stored in the context for that *spoke-sdp-fec-id*.

Otherwise, if the SAII:TAII to use in the target FEC stack of the vccv-ping message is not entered by the user, and if a switching point TLV was previously received in the initial label mapping message for the reverse direction of the MS-PW (with respect to the sending PE), then the SAII:TAII to use in the target FEC stack of the vccv-ping echo request message is derived by parsing that switching point TLV based on the user-specified TTL (or a TTL of 255 if none is specified). In this case, the order of the SAII:TAII in the switching point TLV is maintained for the vccv-ping echo request message.

If no pseudowire switching point TLV was received, then the SAII:TAII values to use for the vccv-ping echo request are derived from the MS-PW context, but their order is reversed before being sent so that they match the order for the downstream FEC element for an S-PE, or the locally configured SAII:TAII for a remote T-PE of that MS-PW.

The use of *spoke-sdp-fec-id* in vccv-ping is only applicable at T-PE nodes, since it is not configured for a specified MS-PW at S-PE nodes.

2.8.6.7 VCCV-Trace on Dynamic MS-PWs

The 7750 SR, 7450 ESS, and 7950 XRS support the MS-PW path trace mode of operation for VCCV trace, as per pseudowire switching, but using FEC129 AII type 2. As in the case of vccv-ping, the SAII:TAII used in the VCCV echo request message sent from the T-PE or S-PE from which the VCCV trace command is executed is specified by the user or derived from the context of the MS-PW. The use of *spoke-sdp-fec-id* in vccv-trace is only applicable at T-PE nodes, since it is not configured for a specified MS-PW at S-PE nodes.

2.8.7 Example Dynamic MS-PW Configuration

This section describes an example of how to configure Dynamic MS-PWs for a VLL service between a set of Nokia nodes. The network consists of two T-PEs and two nodes, in the role of S-PEs, as shown in the following figure. Each 7750 SR, 7450 ESS, or 7950 XRS peers with its neighbor using LDP and BGP.

Figure 24 **Dynamic MS-PW Example**

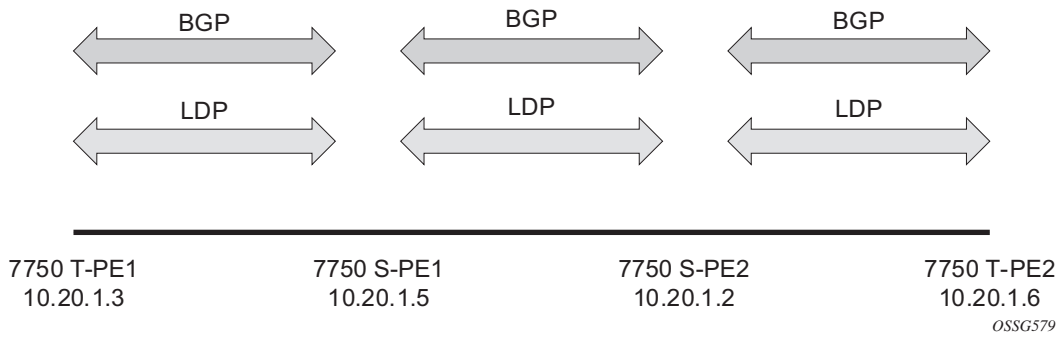

The example uses BGP to route dynamic MS-PWs and T-LDP to signal them. Therefore, each node must be configured to support the MS-PW address family under BGP, and BGP and LDP peerings must be established between the T-PEs/S-PEs. The appropriate BGP export policies must also be configured.

Next, pseudowire routing must be configured on each node. This includes an S-PE address for every participating node, and one or more local prefixes on the T-PEs. MS-PW paths and static routes may also be configured.

When this routing and signaling infrastructure is established, spoke-sdp-fecs can be configured on each of the T-PEs, as follows:

```
config
   router
      ldp
         targeted-session
            peer 10.20.1.5
            exit
         exit
      policy-options
         begin
         policy-statement "exportMsPw"
            entry 10
               from
                  family ms-pw
               exit
               action accept
               exit
            exit
         exit
         commit
      exit
      bgp
         family ms-pw
         connect-retry 1
         min-route-advertisement 1
         export "exportMsPw"
         rapid-withdrawal
         group "ebgp"
            neighbor 10.20.1.5
               multihop 255
               peer-as 200
            exit
         exit
     exit
config
   service
      pw-routing
         spe-address 3:10.20.1.3
         local-prefix 3:10.20.1.3 create
         exit
         path "path1_to_F" create
            hop 1 10.20.1.5
            hop 2 10.20.1.2
            no shutdown
        exit
     exit
     epipe 1 customer 1 vpn 1 create
        description "Default epipe
             description for service id 1"
        service-mtu 1400
        service-name "XYZ Epipe 1"
        sap 2/1/1:1 create
        exit
        spoke-sdp-fec 1 fec 129 aii-type 2 create
           retry-timer 10
           retry-count 10
```

```
            saii-type2 3:10.20.1.3:1
            taii-type2 6:10.20.1.6:1
            no shutdown
         exit
         no shutdown
      exit
config
   router
      ldp
         targeted-session
            peer 10.20.1.2
            exit
         exit
         …
      policy-options
         begin
         policy-statement "exportMsPw"
            entry 10
               from
                  family ms-pw
               exit
               action accept
               exit
            exit
         exit
         commit
      exit

      bgp
         family ms-pw
         connect-retry 1
         min-route-advertisement 1
         export "exportMsPw"
         rapid-withdrawal
         group "ebgp"
            neighbor 10.20.1.2
               multihop 255
               peer-as 300
            exit
         exit
      exit
config
   service
      pw-routing
         spe-address 6:10.20.1.6
         local-prefix 6:10.20.1.6 create
         exit
         path "path1_to_F" create
            hop 1 10.20.1.2
            hop 2 10.20.1.5
            no shutdown
         exit
      exit
      epipe 1 customer 1 vpn 1 create
         description "Default epipe
              description for service id 1"
service-mtu 1400
         service-name "XYZ Epipe 1"
         sap 1/1/3:1 create
```

```
        exit
        spoke-sdp-fec 1 fec 129 aii-type 2 create
           retry-timer 10
           retry-count 10
           saii-type2 6:10.20.1.6:1
           taii-type2 3:10.20.1.3:1
           no shutdown
        exit
        no shutdown
     exit


config
   router
      ldp
         targeted-session
            peer 10.20.1.3
            exit
            peer 10.20.1.2
            exit
         exit
         …
      bgp
         family ms-pw
         connect-retry 1
         min-route-advertisement 1
         rapid-withdrawal
         group "ebgp"
            neighbor 10.20.1.2
               multihop 255
               peer-as 300
            exit
            neighbor 10.20.1.3
               multihop 255
               peer-as 100
            exit
         exit
     exit

   service
      pw-routing
         spe-address 5:10.20.1.5
      exit


config
   router
      ldp
         targeted-session
            peer 10.20.1.5
            exit
            peer 10.20.1.6
            exit
         exit
         …
      bgp
         family ms-pw
         connect-retry 1
         min-route-advertisement 1
```

```
            rapid-withdrawal
            group "ebgp"
               neighbor 10.20.1.5
                  multihop 255
                  peer-as 200
               exit
               neighbor 10.20.1.6
                  multihop 255
                  peer-as 400
               exit
            exit
         exit
      service
          pw-routing
            spe-address 2:10.20.1.2
          exit
```

2.8.8 VLL Resilience with Two Destination PE Nodes

Figure 25 illustrates the application of pseudowire redundancy to provide Ethernet VLL service resilience for broadband service subscribers accessing the broadband service on the service provider BRAS.

Figure 25 **VLL Resilience**

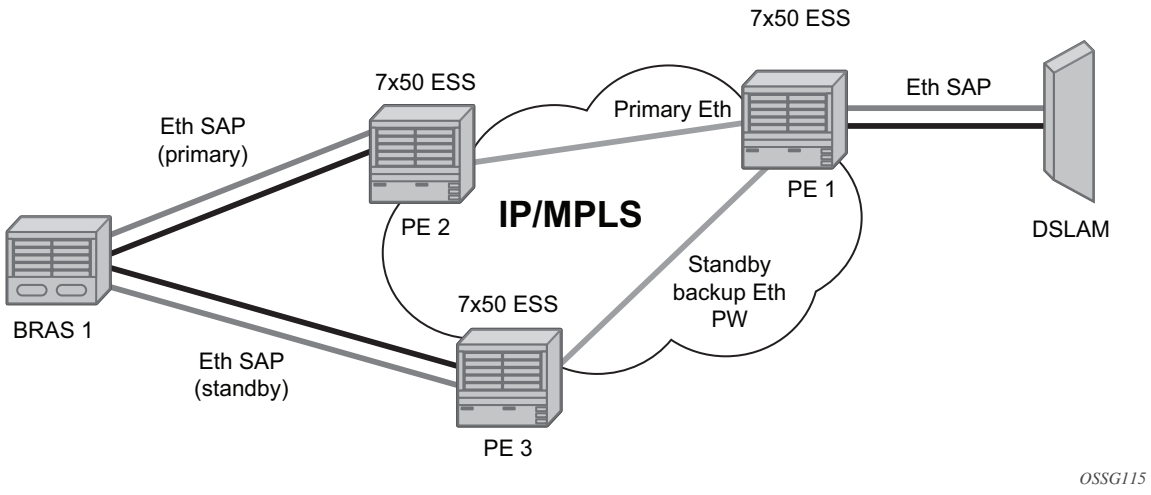

If the Ethernet SAP on PE2 fails, PE2 notifies PE1 of the failure by either withdrawing the primary pseudowire label it advertised or by sending a pseudowire status notification with the code set to indicate a SAP defect. PE1 will receive it and will immediately switch its local SAP to forward over the secondary standby spoke-SDP. To avoid black holing of packets during the switching of the path, PE1 will accept packets received from PE2 on the primary pseudowire while transmitting over the

backup pseudowire. However, in other applications such as those described in Access Node Resilience Using MC-LAG and Pseudowire Redundancy, it will be important to minimize service outage to end users.

When the SAP at PE2 is restored, PE2 updates the new status of the SAP by sending a new label mapping message for the same pseudowire FEC or by sending pseudowire status notification message indicating that the SAP is back up. PE1 then starts a timer and reverts back to the primary at the expiry of the timer. By default, the timer is set to 0, which means PE1 reverts immediately. A special value of the timer (infinity) will mean that PE1 should never revert back to the primary pseudowire.

The behavior of the pseudowire redundancy feature is the same if PE1 detects or is notified of a network failure that brought the spoke-SDP status to operationally down. The following are the events that will cause PE1 to trigger a switchover to the secondary standby pseudowire:

1. T-LDP peer (remote PE) node withdrew the pseudowire label.
2. T-LDP peer signaled a FEC status indicating a pseudowire failure or a remote SAP failure.
3. T-LDP session to peer node times out.
4. SDP binding and VLL service went down as a result of network failure condition such as the SDP to peer node going operationally down.

The SDP type for the primary and secondary pseudowires need not be the same. In other words, the user can protect a RSVP-TE based spoke-SDP with an LDP or GRE based one. This provides the ability to route the path of the two pseudowires over different areas of the network. All VLL service types, for example, Apipe, Epipe, Fpipe, and Ipipe, are supported on the 7750 SR.

Nokia routers support the ability to configure multiple secondary standby pseudowire paths. For example, PE1 uses the value of the user-configurable precedence parameter associated with each spoke-SDP to select the next available pseudowire path after the failure of the current active pseudowire (whether it is the primary or one of the secondary pseudowires). The revertive operation always switches the path of the VLL back to the primary pseudowire though. There is no revertive operation between secondary paths, meaning that the path of the VLL will not be switched back to a secondary pseudowire of higher precedence when the latter comes back up again.

Nokia routers support the ability for a user-initiated manual switchover of the VLL path to the primary or any of the secondary, be supported to divert user traffic in case of a planned outage such as in node upgrade procedures.

On the 7750 SR, this application can make use of all types of VLL supported on SR-series routers. However, if a SAP is configured on an MC-LAG instance, only the Epipe service type is allowed.

2.8.8.1 Master-Slave Operation

This section describes master-slave pseudowire redundancy. It adds the ability for the remote peer to react to the pseudowire standby status notification, even if only one spoke-SDP terminates on the VLL endpoint on the remote peer, by blocking the transmit (Tx) direction of a VLL spoke-SDP when the far-end PE signals standby. This solution enables the blocking of the Tx direction of a VLL spoke-SDP at both master and slave endpoints when standby is signaled by the master endpoint. This approach satisfies a majority of deployments where bidirectional blocking of the forwarding on a standby spoke-SDP is required.

Figure 26 illustrates the operation of master-slave pseudowire redundancy. In this scenario, an Epipe service is provided between CE1 and CE2. CE2 is dual-homed to PE2 and PE3; therefore, PE1 is dual-homed to PE2 and PE3 using Epipe spoke-SDPs. The objective of this feature is to ensure that only one pseudowire is used for forwarding in both directions by PE1, PE2, and PE3 in the absence of a native dual homing protocol between CE2 and PE2/PE3, such as MC-LAG. In normal operating conditions (the SAPs on PE2 and PE3 toward CE2 are both up and there are no defects on the ACs to CE2), PE2 and PE3 cannot choose which spoke-SDP to forward on, based on the status of the AC redundancy protocol.

Figure 26 **Master-Slave Pseudowire Redundancy**

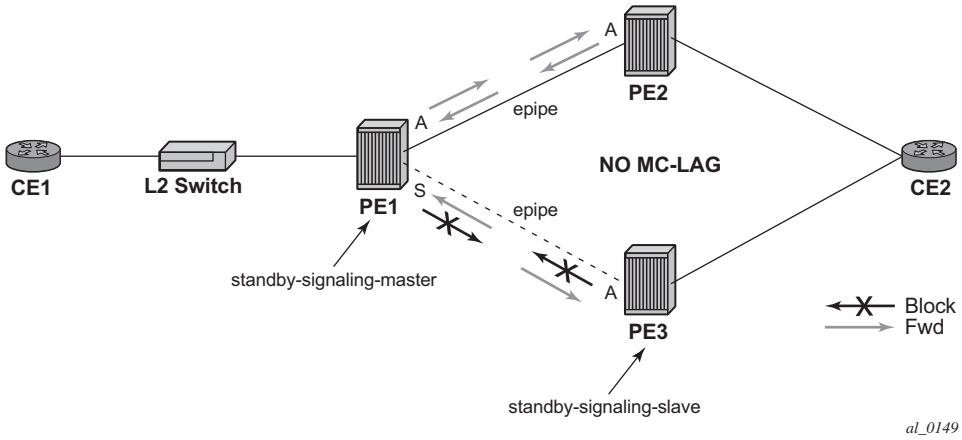

Master-slave pseudowire redundancy adds the ability for the remote peer to react to the pseudowire standby status notification, even if only one spoke-SDP terminates on the VLL endpoint on the remote peer. When the CLI command **standby-signaling-slave** is enabled at the spoke-SDP or explicit endpoint level in PE2 and PE3, then any spoke-SDP for which the remote peer signals PW FWD Standby will be blocked in the transmit direction.

This is achieved as follows. The **standby-signaling-master** state is activated on the VLL endpoint in PE1. In this case, a spoke-SDP is blocked in the transmit direction at this master endpoint if it is either in operDown state, or it has lower precedence than the highest precedence spoke-SDP, or the specific peer PE signals one of the following pseudowire status bits:

- Pseudowire not forwarding (0x01)
- SAP (ingress) receive fault (0x02)
- SAP (egress) transmit fault (0x04)
- SDP binding (ingress) receive fault (0x08)
- SDP binding (egress) transmit fault (0x10)

That the specified spoke-SDP has been blocked will be signaled to the LDP peer through the pseudowire status bit (PW FWD Standby (0x20)). This will prevent traffic being sent over this spoke-SDP by the remote peer, but only if that remote peer supports and reacts to pseudowire status notification. Previously, this applied only if the spoke-SDP terminates on an IES, VPRN, or VPLS. However, if standby-signaling-slave is enabled at the remote VLL endpoint, the Tx direction of the spoke-SDP will also be blocked, according to the rules in Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios.

Although master-slave operation provides bidirectional blocking of a standby spoke-SDP during steady-state conditions, it is possible that the Tx directions of more than one slave endpoint can be active for transient periods during a fail-over operation. This is due to slave endpoints transitioning a spoke-SDP from standby to active receiving and/or processing a pseudowire preferential forwarding status message before those endpoints transitioning a spoke-SDP to standby. This transient condition is most likely when a forced switchover is performed, or the relative preferences of the spoke-SDPs are changed, or the active spoke-SDP is shutdown at the master endpoint. During this period, loops of unknown traffic may be observed. Fail-overs due to common network faults that can occur during normal operation, or a failure of connectivity on the path of the spoke-SDP or the SAP, would not result in such loops in the data path.

2.8.8.1.1 Interaction with SAP-Specific OAM

If all of the spoke-SDPs bound to a SAP at a slave PE are selected as standby, then this should be treated from a SAP OAM perspective in the same manner as a fault on the service: an SDP binding down or remote SAP down. That is, a fault should be indicated to the service manager. If SAP-specific OAM is enabled toward the CE, such as Ethernet Continuity Check Message (CCM), Ethernet Link Management Interface (E-LMI), or FR LMI, then this should result in the appropriate OAM message being sent on the SAP. This can enable the remote CE to avoid forwarding traffic toward a SAP that will drop it.

Figure 27 shows an example for the case of Ethernet LMI.

Figure 27 **Example of SAP OAM Interaction with Master-Slave Pseudowire Redundancy**

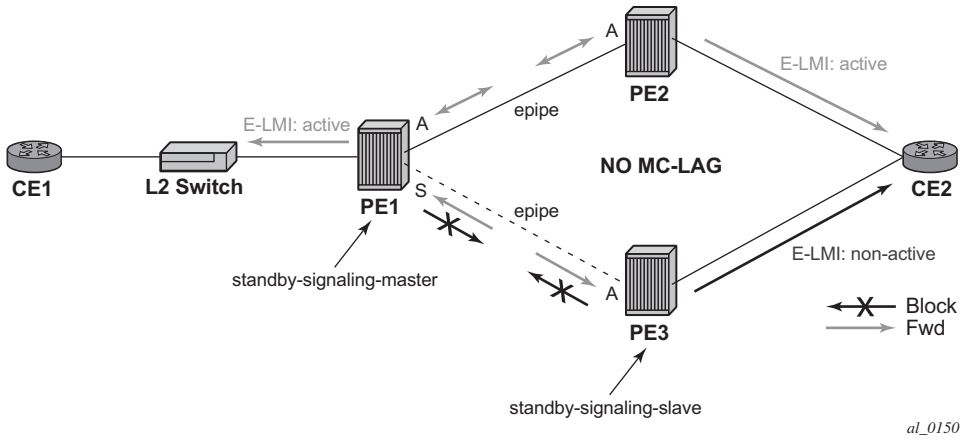

2.8.8.1.2 Local Rules at Slave VLL PE

It is not possible to configure a standby-signaling-slave on endpoints or spoke-SDPs that are bound to an IES, VPRN, ICB, or MC-EP, or that are part of an MC-LAG or MC-APS.

If **standby-signaling-slave** is configured on a specific spoke-SDP or explicit endpoint, then the following rules apply. The rules describe the case of several spoke-SDPs in an explicit endpoint. The same rules apply to the case of a single spoke-SDP outside of an endpoint where no endpoint exists:

- Rules for processing endpoint SAP active/standby status bits:
  - Since the SAP in endpoint X is never a part of an MC-LAG/MC-APS instance, a forwarding status of active is always advertised.

- Rules for processing and merging local and received endpoint objects with an up or down operational status:

1. Endpoint X is operationally up if at least one of its objects is operationally up. It is Down if all of its objects are operationally down.
2. If all objects in endpoint X did any or all of the following, the node must send status bits of SAP down over all Y endpoint spoke-SDPs:
   - transitioned locally to down state
   - received a SAP down notification via remote T-LDP or via SAP-specific OAM signal
   - received status bits of SDP-binding down
   - received states bits of PW not forwarding
3. Endpoint Y is operationally up if at least one of its objects is operationally up. It is down if all its objects are operationally down.
4. If a spoke-SDP in endpoint Y, including the ICB spoke-SDP, transitions locally to down state, the node must send T-LDP SDP-binding down status bits on this spoke-SDP.
5. If a spoke-SDP in endpoint Y received T-LDP SAP down status bits, and/or T-LDP SDP-binding down status bits, and/or status bits of PW not forwarding, the node saves this status and takes no further action. The saved status is used for selecting the active transmit endpoint object.
6. If all objects in endpoint Y, or a single spoke-SDP that exists outside of an endpoint (and no endpoint exists), the node must send a SAP down notification on the X endpoint SAP via the SAP-specific OAM signal, if applicable:
   - transitioned locally to down state
   - received status bits of T-LDP SAP down
   - received status bits of T-LDP SDP-binding down
   - received status bits of PW not forwarding
   - received status bits of PW FWD standby
7. If the peer PE for a specified object in endpoint Y signals PW FWD standby, the spoke-SDP must be blocked in the transmit direction and the spoke-SDP is not eligible for selection by the active transmit selection rules.
8. If the peer PE for a specified object in endpoint Y does not signal PW FWD standby, then spoke-SDP is eligible for selection.

2.8.8.1.3 Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios

This section discusses how master-slave pseudowire redundancy could operate.

VLL Resilience Path Example

Figure 28 shows VLL resilience path example. An sample configuration follows.

Figure 28 **VLL Resilience**

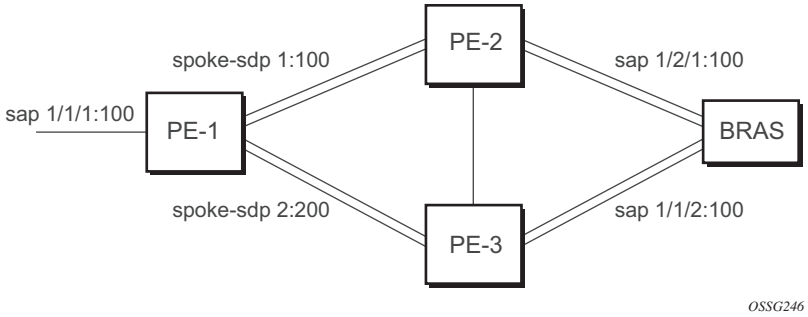

A **revert-time** value of zero (default) means that the VLL path will be switched back to the primary immediately after it comes back up.

```
PE-1
configure service epipe 1
endpoint X
exit
endpoint Y
   revert-time 0
   standby-signaling-master
   exit
   sap 1/1/1:100 endpoint X
   spoke-sdp 1:100 endpoint Y
precedence primary
   spoke-sdp 2:200 endpoint Y
precedence 1

PE-2
configure service epipe 1
   endpoint X
   exit
   sap 2/2/2:200 endpoint X
   spoke-sdp 1:100
      standby-signaling-slave

PE-3
configure service epipe 1
   endpoint X
   exit
   sap 3/3/3:300 endpoint X
   spoke-sdp 2:200
      standby-signaling-slave
```

2.8.8.1.4 VLL Resilience for a Switched Pseudowire Path

Figure 29 displays VLL resilience for a switched pseudowire path example. A sample configuration follows.

Figure 29 **VLL Resilience with Pseudowire Switching**

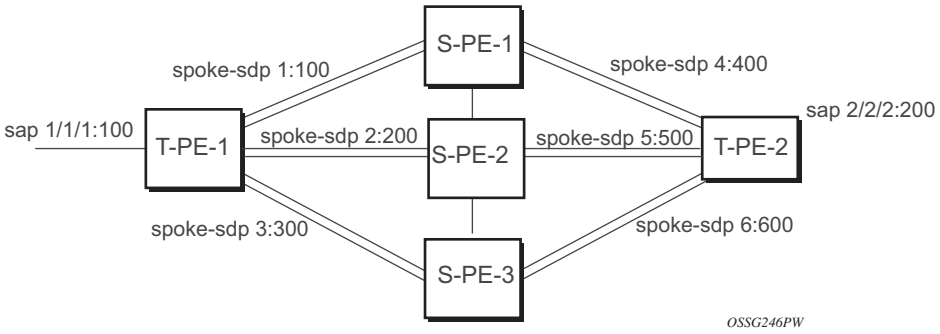


```
T-PE-1
configure service epipe 1
   endpoint X
   exit
   endpoint Y
   revert-time 100
   standby-signaling-master
   exit
   sap 1/1/1:100 endpoint X
   spoke-sdp 1:100 endpoint Y
      precedence primary
   spoke-sdp 2:200 endpoint Y
      precedence 1
   spoke-sdp 3:300 endpoint Y
      precedence 1

T-PE-2
configure service epipe 1
   endpoint X
   exit
   endpoint Y
   revert-time 100
   standby-signaling-slave
   exit
   sap 2/2/2:200 endpoint X
   spoke-sdp 4:400 endpoint Y
      precedence primary
   spoke-sdp 5:500 endpoint Y
      precedence 1
   spoke-sdp 6:600 endpoint Y
      precedence 1


S-PE-1
```

VC switching indicates a VC cross-connect so that the service manager does not signal the VC label mapping immediately but will put S-PE-1 into passive mode, as follows:

```
configure service epipe 1 vc-switching
   spoke-sdp 1:100
   spoke-sdp 4:400
```

2.8.9 Pseudowire SAPs

Refer to the *7450 ESS, 7750 SR, and 7950 XRS Layer 3 Services Guide: Internet Enhanced Services and Virtual Private Routed Network Services* for information about how to use pseudowire SAPs with Layer 2 services.

2.8.10 Epipe Using BGP-MH Site Support for Ethernet Tunnels

Using Epipe in combination with G.8031 and BGP multi-homing in the same manner as VPLS offers a multi-chassis resiliency option for Epipe services that is a non-learning and non-flooded service. MC-LAG (see Access Node Resilience Using MC-LAG and Pseudowire Redundancy) offers access node redundancy with active/stand-by links while Ethernet tunnels offer per service redundancy with all active links and active or standby services. G.8031 offers an end-to-end service resiliency for Epipe and VPLS services. BGP-MH site support for Ethernet tunnels offers Ethernet edge resiliency for Epipe services that integrates with MPLS pseudowire redundancy.

Figure 30 shows the BGP-MH site support for Ethernet tunnels, where a G.8031 edge device (A) is configure with two provider edge switches (B and C). G.8031 is configured on the Access devices (A and F). An Epipe endpoint service is configured along with BGP Multi-homing and Pseudowire Redundancy on the provider edge nodes (B/C and D/E). This configuration offers a fully redundant Epipe service.

Figure 30 **BGP-MH Site Support for Ethernet Tunnels**

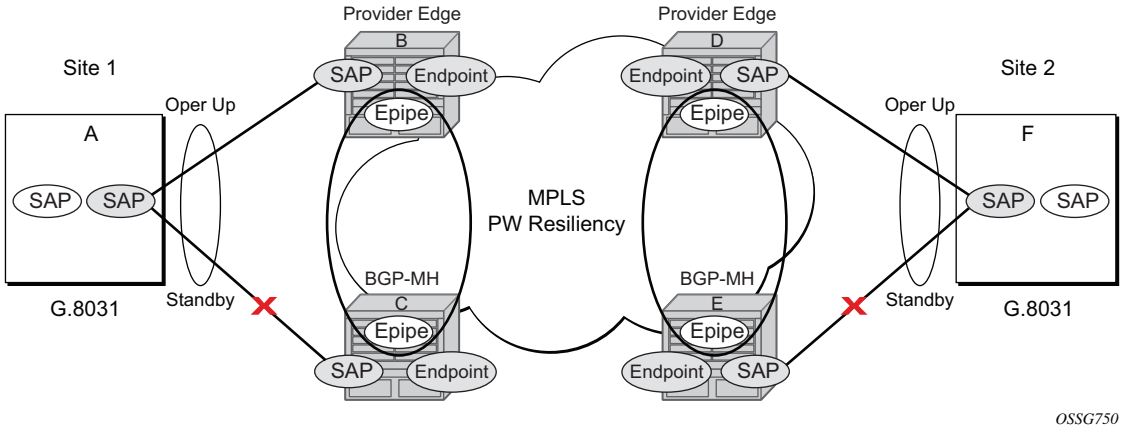

2.8.10.1 Operational Overview

G.8031 offers a number of redundant configurations. Normally, it offers the ability to control two independent paths for 1:1 protection. In the BGP-MH site support for Ethernet tunnels case, BGP drives G.8031 as a slave service. In this case, the provider edge operates using only standard 802.1ag MEPs with CCM to monitor the paths. Figure 31 shows an Epipe service on a Customer Edge (CE) device that uses G.8031 with two paths and two MEPs. The paths can use a single VLAN of dot1q or QinQ encapsulation.

Figure 31 **G.8031 for Slave Operation**

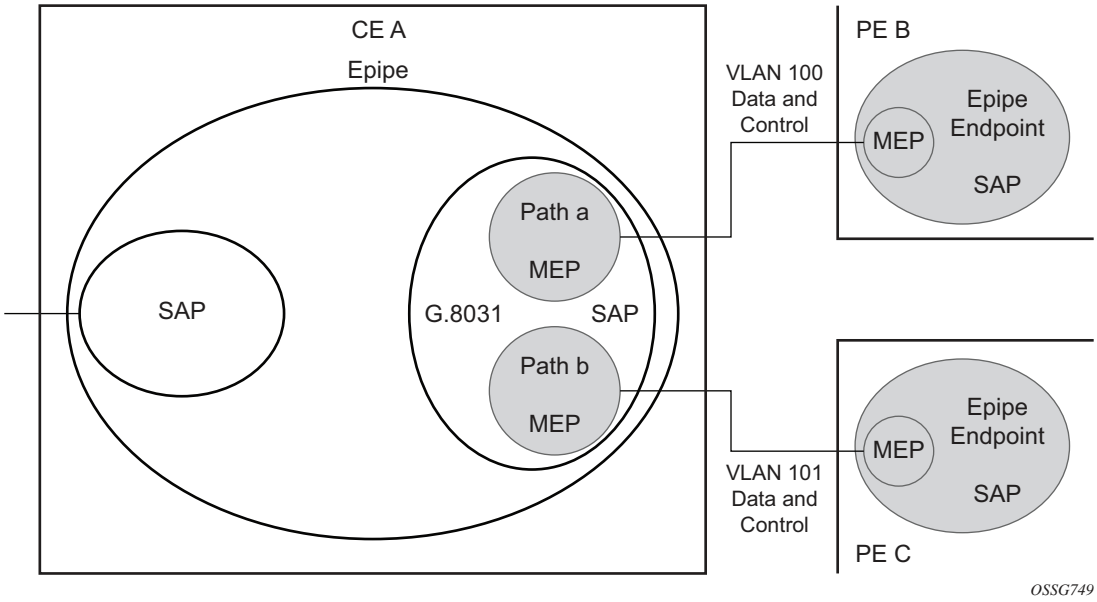

In a single-service deployment, the control (CFM) and data will share the same port and VID. For multiple services for scaling, fate sharing is allowed between multiple SAPs, but all SAPs within a group must be on the same physical port.

To get fate sharing for multiple services with this feature, a dedicated G.8031 CE-based service (one VLAN) is connected to a Epipe SAP on a PE, which uses BGP-MH and operational groups to control other G.8031 tunnels. This dedicated G.8031 service still has data control capabilities, but the data Epipe service is not bearing user data packets. On the CE, this G.8031 is only used for group control. Making this a dedicated control (CFM) for a set of G.8031 tunnels is to simplify operation and allow individual disabling of services. Using a dedicated G.8031 service to both control and to carry data traffic is allowed.

Fate sharing from the PE side is achieved using BGP and operational groups. G.8031 Epipe services can be configured on the CE as regular non-fate shared G.8031 services, but due to the configuration on the PE side, these Ethernet tunnels will be treated as a group following the one designated control service. The G.8031 control logic on the CE is a slave to the BGP-MH control.

On the CE, G.8031 allows independent configuration of VIDs on each path. On the PE, the Epipe or endpoint that connects to the G.8031 service must have a SAP with the corresponding VID. If the G.8031 service has a Maintenance End Point (MEP) for that VID, the SAP should be configured with a MEP. The MEPs on the paths on the CE signal standard interface status TLV (ifStatusTLV), No Fault (Up), and Fault

(Down). The MEPs on the PE (Epipe or endpoint) also use signaling of ifStatusTLV No Fault (Up), and Fault (Down) to control the G.8031 SAP. However, in the 7750 SR, 7450 ESS, and 7950 XRS model, fate shared Ethernet tunnels with no MEP are allowed. In this case, it is up to the CE to manage these CE-based fate shared tunnels.

Interface status signaling (ifStatusTLV) is used to control the G.8031 tunnel from the PE side. Normally the CE will signal No Fault (Up) in the path SAP MEP ifStatusTLV before the BGP-MH will cause the SAP MEP to become active by signaling No Fault (Up).

2.8.10.2 Detailed Operation

For this feature, BGP-MH is used as the master control and the Ethernet tunnel is a slave. The G.8031 on the CE is unaware that it is being controlled. While a single Epipe service is configured and will serve as the control for the CE connection, allowing fate sharing, all signaling to the CE is based on the ifStatusTLV per G.8031 tunnel. By controlling G.8031 with BGP-MH, the G.8031 CE is forced to be a slave to the PE BGP-MH election. BGP-MH election is controlled by the received VPLS preference or BGP local-preference, or the PE ID (IP address of provider edge) if local-preference is equal to VPLS preference. There may be traps generated on the CE side for some G.8031 implementations, but these can be suppressed or filtered to allow this feature to operate.

There are two configuration options:

- Every G.8031 service SAP terminates on a single Epipe that has BGP-MH. These Epipes may use endpoints with or without ICBs.
- A control Epipe service monitors a single SAP that is used for group control of fate shared CE services. In this case, the Epipe service has a SAP that serves as the control termination for one Ethernet tunnel connection. The group fate sharing SAPs may or may not have MEPs if they use shared fate. In this case, the Epipe may have endpoints but will not support ICBs.

The MEP ifStatusTlv and CCM are used for monitoring the PE to CE SAP. MEP ifStatusTlv is used to signal that the Ethernet tunnel inactive and CCM is used as an aliveness mechanism. There is no G.8031 logic on the PE; the SAP is controlling the corresponding CE SAP.

2.8.10.2.1 Sample Operation of G.8031 BGP-MH

Any Ethernet tunnel actions (force, lock) on the CE (single site) do not control the action to switch paths directly, but they may influence the outcome of BGP-MH if they are on a control tunnel. If a path is disabled on the CE, the result may force the SAP with an MEP on the PE to eventually take the SAP down; Nokia recommends running commands from the BGP-MH side to control these connections.

Figure 32 **Full Redundancy G.8031 Epipe and BGP-MH**

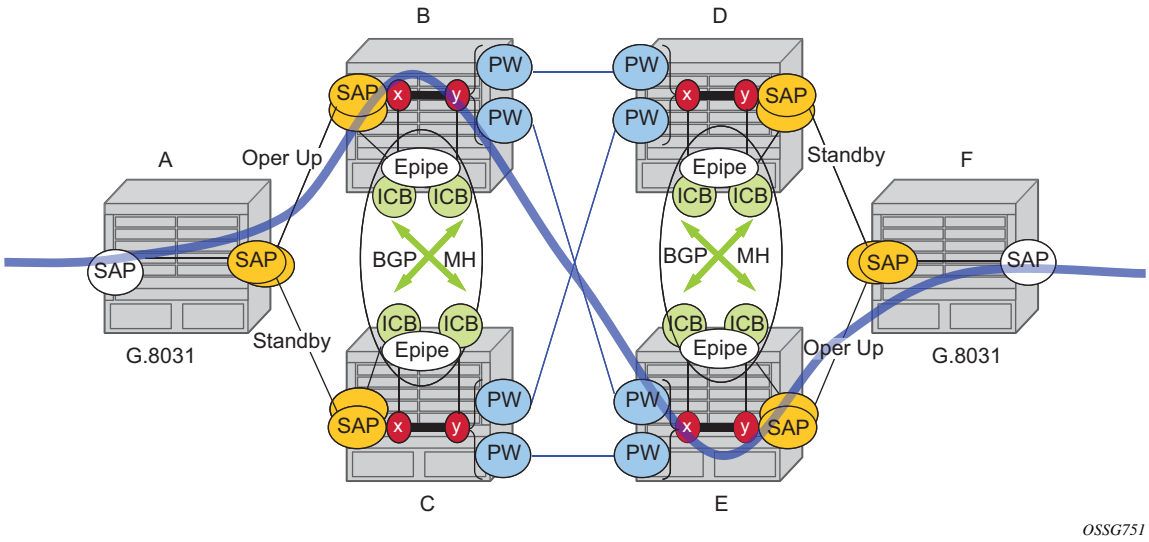

Table 10 lists the SAP MEP signaling shown in Figure 32. For a description of the events shown in this sample operation, see Events in Sample Operation.

Table 10 **SAP MEP Signaling**

	G.8031 ET on CE	Path A MEP Facing Node B Local ifStatus	Path B MEP Facing Node C Local ifStatus	Path B PE MEP ifStatus	Path B PE MEP ifStatus
1	Down (inactive)	No Fault [1]	No Fault	Fault	Fault
2	Up use Path A	No Fault	No Fault	No Fault	Fault
3	Up use Path B	No Fault	No Fault	Fault	No Fault
4	Down Path A fault	Fault [2]	No Fault	Fault	Fault
5	Down Path A and B fault at A	Fault	No Fault	Fault	Fault

Table 10 **SAP MEP Signaling (Continued)**

	G.8031 ET on CE	Path A MEP Facing Node B Local ifStatus	Path B MEP Facing Node C Local ifStatus	Path B PE MEP ifStatus	Path B PE MEP ifStatus
6	Partitioned Network Use Path Precedence Up use Path A	No Fault	No Fault	No Fault	No Fault

Notes:

1. No Fault = no ifStatusTlv transmit | CCM transmit normally
2. Fault = ifStatusTlv transmit down | no CCM transmit

Events in Sample Operation

The following describes the events for switchover in Figure 32. This configuration uses operational groups. The nodes of interest are A, B, and C listed in Table 10.

1. A single G.8031 SAP that represents the control for a group of G.8031 SAPs is configured on the CE.
   - The Control SAP does not normally carry any data; however, it can if needed.
   - An Epipe service is provisioned on each PE node (B/C), only for control (no customer traffic flows over this service).
   - On CE A, there is an Epipe Ethernet tunnel (G.8031) control SAP.
   - The Ethernet tunnel has two paths:
     - one facing B
     - one facing C
   - PE B has an Epipe control SAP that is controlled by the BGP-MH site and PE C also has the corresponding SAP that is controlled by the same BGP-MH site.
2. At node A, there are MEPs configured under each path that check connectivity on the A-B and A-C links. At nodes B and C, there is a MEP configured under their respective SAPs with fault propagation enabled with the use of ifStatusTlv.
3. Initially, assume there is no link failure:
   - SAPs on node A have ifStatusTLV No Fault to B and C (no MEP fault detected at A); see Table 10 row 1 (Fault is signaled in the other direction PE to CE).

- BGP-MH determines which is the master or Designated Forwarder (DF).
- Assume SAP on node B is picked as the DF.
- The MEP at Path A-B signals ifStatusTlv No Fault. Due to this signal, the MEP under the node A path facing node B detects the path to node B is usable by the path manager on A.

4. At the CE node A, Path A-C becomes standby and is brought down; see Table 10 row 2.
   - Since fault propagation is enabled under the SAP node C MEP, and ifStatusTLV is operationally Down, the Path remains in the present state.
   - Under these conditions, the MEP under the node A path facing node C detects the fault and informs Ethernet manager on node A.
   - Node A then considers bringing path A-C down.
   - ET port remains up since path A-B is operationally up. This is a stable state.
5. On nodes B and C, each Epipe-controlled SAP is the sole (controlling) member of an operational group.
   - Other data SAPs may be configured for fate shared VLANs (Ethernet tunnels) and to monitor the control SAP.
   - The SAPs facing the CE node A share the fate of the control SAP and follow the operation.
6. If there is a break in path A-B connectivity (CCM timeout or LOS on the port for link A-B), then on node A the path MEP detects connectivity failure and informs Ethernet tunnel manager; see Table 10 row 4.
7. At this point, the Ethernet tunnel is down since both path A-B and path A-C are down.
8. The CE node A Ethernet tunnel goes down.
9. At node B on the PE, the SAP also detects the failure and the propagation of fault status goes to BGP-MH; see Table 10 row 4.
10. This in turn feeds into BGP-MH, which deems the site non-DF and makes the site standby.
11. Since the SAP at node B is standby, the service manager feeds this to CFM, which then propagates a Fault toward node A. This is a cyclic fault propagation. However, since path A-B is broken, the situation is stable; see Table 10 row 5.
12. There is traffic loss during the BGP-MH convergence.
    - Load sharing mode is recommended when using a 7450 as a CE node A device.
    - BGP-MH signals that node C is now the DF; see Table 10 row 3.
13. BGP-MH on node C elects a SAP and brings it up.
14. ET port transitions to port A-C, and is operationally up. This is a stable state. The A-C SAPs monitoring the operational group on C transitions to operationally up.

Unidirectional failures: At point 6 the failure was detected at both ends. In the case of a unidirectional failure, CCM times out on one side.

1. In the case where the PE detects the failure, it propagates the failure to BGP-MH and the BGP-MH takes the site down causing the SAPs on the PE to signal a Fault to the CE.
2. In the case where G.8031 on the CE detects the failure, it takes the tunnel down and signals a fault to the PE, and then the SAP propagates that to BGP-MH.

2.8.10.3 BGP-MH Site Support for Ethernet Tunnels Operational Group Model

For operational groups, one or more services follow the controlling service. On node A, there is an ET SAP facing nodes B/C, and on nodes B/C there are SAPs of the Epipe on physical ports facing node A. Each of the PE data SAPs monitor their respective operational groups, meaning they are operationally up or down based on the operational status of the control SAPs. On node A, since the data SAP is on the ET logical port, it goes operationally down whenever the ET port goes down and similarly for going operationally up.

Alternatively, an Epipe service may be provisioned on each node for each G.8031 data SAP (one-for-one service with no fate sharing). On CE node A, there will be a G.8031 Ethernet tunnel. The Ethernet tunnel has two paths: one facing node B and one facing node C. This option is the same as the control SAP, but there are no operational groups. However, now there is a BGP-MH Site per service. For large sites, operational groups are more efficient.

2.8.10.4 BGP-MH Specifics for MH Site Support for Ethernet Tunnels

BGP Multi-Homing for VPLS describes the procedures for using BGP to control resiliency for VPLS. These procedures are the same except that an Epipe service can be configured for BGP-MH.

2.8.10.5 PW Redundancy for BGP MH Site Support for Ethernet Tunnels

Pseudowire Redundancy Service Models and VLL Resilience with Pseudowire Redundancy and Switching are used for the MPLS network resiliency. BGP MH site support for Ethernet tunnels reuses this model.

2.8.10.6 T-LDP Status Notification Handling Rules of BGP-MH Epipes

Using Figure 35 as a reference, the following are the rules for generating, processing, and merging T-LDP status notifications in VLL service with endpoints.

2.8.10.6.1 Rules for Processing Endpoint SAP Active/Standby Status Bits

1. The advertised admin forwarding status of Active/Standby reflects the status of the local Epipe SAP in BGP-MH instance. If the SAP is not part of an MC-LAG instance or a BGP-MH instance, the forwarding status of Active is always advertised.
2. When the SAP in endpoint X is part of a BGP-MH instance, a node must send T-LDP forwarding status bit of SAP Active/Standby over all Y endpoint spoke-SDPs, except the ICB spoke-SDP, whenever this (BGP-MH designated forwarder) status changes. The status bit sent over the ICB is always zero (Active by default).
3. When the SAP in endpoint X is not part of an MC-LAG instance or BGP-MH instance, then the forwarding status sent over all Y endpoint spoke-SDPs should always be set to zero (Active by default).
4. The received SAP Active/Standby status is saved and used for selecting the active transmit endpoint object Pseudowire Redundancy procedures.

2.8.10.6.2 Rules for Processing, Merging Local, and Received Endpoint Operational Status

1. Endpoint X is operationally up if at least one of its objects is operationally Up. It is Down if all its objects are operationally down.

2. If the SAP in endpoint X transitions locally to the down state, or received a SAP Down notification via SAP-specific OAM signal (SAP MEP), the node must send T-LDP SAP down status bits on the Y endpoint ICB spoke-SDP only. BGP-MH SAPs support MEPs for ifStatusTLV signaling. All other SAP types cannot exist on the same endpoint as an ICB spoke-SDP since non Ethernet SAP cannot be part of an MC-LAG instance or a BGP-MH Instance.
3. If the ICB spoke-SDP in endpoint X transitions locally to Down state, the node must send T-LDP SDP-binding down status bits on this spoke-SDP.
4. If the ICB spoke-SDP in endpoint X received T-LDP SDP-binding down status bits or PW not forwarding status bits, the node saves this status and takes no further action. The saved status is used for selecting the active transmit endpoint object as per Pseudowire Redundancy procedures.
5. If all objects in endpoint X did any or all of the following, the node must send status bits of SAP Down over all Y endpoint spoke-SDPs, including the ICB:
   - transitioned locally to the down state due to operator or BGP-MH DF election
   - received a SAP down notification via remote T-LDP status bits or via SAP-specific OAM signal (SAP MEP)
   - received status bits of SDP-binding down
   - received status bits of PW not forwarding
6. Endpoint Y is operationally up if at least one of its objects is operationally Up. It is Down if all its objects are operationally down.
7. If a spoke-SDP in endpoint Y, including the ICB spoke-SDP, transitions locally to down state, the node must send T-LDP SDP-binding down status bits on this spoke-SDP.
8. If a spoke-SDP in endpoint Y, including the ICB spoke-SDP, did any or all of the following, the node saves this status and takes no further action:
   - received T-LDP SAP down status bits
   - received T-LDP SDP-binding down status bits
   - received PW not forwarding status bits

   The saved status is used for selecting the active transmit endpoint object as per Pseudowire Redundancy procedures.
9. If all objects in endpoint Y, except the ICB spoke-SDP, did any or all of the following, the node must send status bits of SDP-binding down over the X endpoint ICB spoke-SDP only:
   - transitioned locally to the down state
   - received T-LDP SAP down status bits
   - received T-LDP SDP-binding down status bits
   - received PW not forwarding status bits

10. If all objects in endpoint Y did any or all of the following, the node must send status bits of SDP-binding down over the X endpoint ICB spoke-SDP only, and must send a SAP down notification on the X endpoint SAP via the SAP-specific OAM signal:
    - transitioned locally to down state
    - received T-LDP SAP down status bits
    - received T-LDP SDP-binding down status bits
    - received PW not forwarding status bits

    In this case the SAP MEP ifStatusTLV is operationally down and also signals the BGP-MH Site, if this SAP is part of a BGP Site.

2.8.10.6.3 Operation for BGP MH Site Support for Ethernet Tunnels

A multi-homed site can be configured on up to four PEs although two PEs are sufficient for most applications, with each PE having a single object SAP connecting to the multi-homed site. SR OS G.8031 implementation with load sharing allows multiple PEs as well. The designated forwarder election chooses a single connection to be operationally up, with the other placed in standby. Only revertive behavior is supported in release 14.0.

Fate sharing (the status of one site can be inherited from another site) is achievable using monitor-groups.

The following are supported:

- All Ethernet-tunnel G.8031 SAPs on CE:
  - 7750 SR, 7450 ESS, or 7950 XRS G.8031 in load sharing mode (recommended)
  - 7750 SR, 7450 ESS, or 7950 XRS G.8031 in non-load sharing mode
- Epipe and endpoint with SAPs on PE devices.
- Endpoints with PW.
- Endpoints with active/standby PWs.

There are the following constraints with this feature:

- Not supported with PBB Epipes.
- Spoke SDP (pseudowire).
  - BGP signaling is not supported.
  - Cannot use BGP MH for auto-discovered pseudowire. This is achieved in a VPLS service using SHGs, which are not available in Epipes.

- Other multi-chassis redundancy features are not supported on the multi-homed site object, as follows:
  - MC-LAG
  - MC-EP
  - MC-Ring
  - MC-APS
- Master and Slave pseudowire is not supported.

Figure 33 **Sample Topology Full Redundancy**

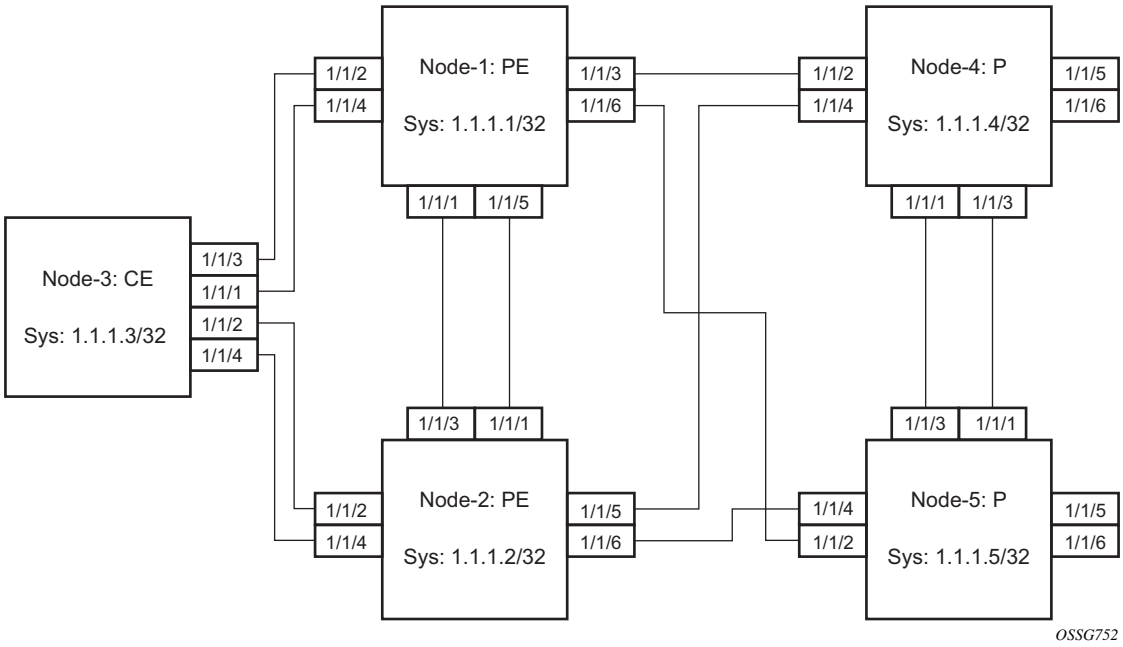

See the following Configuration Examples for configuration examples derived from Figure 33.

Configuration Examples

Node-1: Using operational groups and Ethernet CFM per SAP

```
#--------------------------------------------------
echo "Eth-CFM Configuration"
#--------------------------------------------------
    eth-cfm
        domain 100 format none level 3
            association 2 format icc-based name "node-3-site-1-0"
                bridge-identifier 1
                exit
```

```
                remote-mepid 310
            exit
            association 2 format icc-based name "node-3-site-1-1"
                bridge-identifier 100
                exit
                remote-mepid 311
            exit
        exit
    exit

#--------------------------------------------------
echo "Service Configuration"
#--------------------------------------------------
    service
        customer 1 create
            description "Default customer"
        exit
        sdp 2 mpls create
            far-end 1.1.1.4
            lsp "to-node-4-lsp-1"
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 3 mpls create  //  Etcetera

        pw-template 1 create
            vc-type vlan
        exit
        oper-group "og-name-et" create
        exit
        oper-group "og-name-et100" create
        exit
        epipe 1 customer 1 create
            service-mtu 500
            bgp
                route-distinguisher 65000:1
                route-target export target:65000:1 import target:65000:1
            exit
            site "site-1" create
                site-id 1
                sap 1/1/2:1.1
                boot-timer 100
                site-activation-timer 2
                no shutdown
            exit
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap 1/1/2:1.1 endpoint "x" create
                eth-cfm
                    mep 130 domain 100 association 2 direction down
                        fault-propagation-enable use-if-tlv
                        ccm-enable
                        no shutdown
                    exit
                exit
```

```
                    oper-group "og-name-et"
                exit
                spoke-sdp 2:1 endpoint "y" create
                    precedence primary
                    no shutdown
                exit
                spoke-sdp 3:1 endpoint "y" create
                    precedence 2
                    no shutdown
                exit
                no shutdown
            exit
            epipe 100 customer 1 create
                description "Epipe 100 in separate opergroup"
                service-mtu 500
                bgp
                    route-distinguisher 65000:2
                    route-target export target:65000:2 import target:65000:2
                exit
                site "site-name-et100" create
                    site-id 1101
                    sap 1/1/4:1.100
                    boot-timer 100
                    site-activation-timer 2
                    no shutdown
                exit

                endpoint "x" create
                exit
                endpoint "y" create
                exit
                sap 1/1/4:1.100 endpoint "x" create
                        eth-cfm
                        mep 131 domain 1 association 2 direction down
                            fault-propagation-enable use-if-tlv
                            ccm-enable
                            no shutdown
                        exit
                    exit
                    oper-group "og-name-et100"

                exit
                spoke-sdp 2:2 vc-type vlan endpoint "y" create
                    precedence 1
                    no shutdown
                exit
                spoke-sdp 3:2 vc-type vlan endpoint "y" create
                    precedence 2
                    no shutdown
                exit
                no shutdown
            exit

        exit
    #--------------------------------------------------
    echo "BGP Configuration"
    #--------------------------------------------------
            bgp
                rapid-withdrawal
```

```
            rapid-update l2-vpn
            group "internal"
                type internal
                neighbor 1.1.1.2
                    family l2-vpn
                exit
            exit
        exit
    exit
```

Node-3: Using operational groups and Ethernet CFM per SAP

```
#--------------------------------------------------
echo "Eth-CFM Configuration"
#--------------------------------------------------
    eth-cfm
        domain 100 format none level 3
            association 2 format icc-based name "node-3-site-1-0"
                bridge-identifier 1
                exit
                ccm-interval 1
                remote-mepid 130
            exit
            association 2 format icc-based name "node-3-site-1-1"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 131
            association 3 format icc-based name "node-3-site-2-0"
                bridge-identifier 1
                exit
                ccm-interval 1
                remote-mepid 120
            exit
            association 3 format icc-based name "node-3-site-2-1"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 121
            exit
        exit
    exit

#--------------------------------------------------
echo "Service Configuration"
#--------------------------------------------------

  eth-tunnel 1
        description "Eth Tunnel loadsharing mode QinQ example"
        protection-type loadsharing
        ethernet
            encap-type qinq
        exit
        path 1
            member 1/1/3
            control-tag 1.1
            eth-cfm
                mep 310 domain 100 association 2
```

```
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        path 2
            member 1/1/4
            control-tag 1.2
            eth-cfm
                mep 320 domain 100 association 3
                    ccm-enablepath
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        no shutdown
    exit
#--------------------------------------------------
echo "Ethernet Tunnel Configuration"
#--------------------------------------------------
    eth-tunnel 2
        description "Eth Tunnel QinQ"
        revert-time 10
        path 1
            precedence primary
            member 1/1/1
            control-tag 1.100
            eth-cfm
                mep 311 domain 100 association 2
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        path 2
            member 1/1/2
            control-tag 1.100
            eth-cfm
                mep 321 domain 100 association 3
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        no shutdown
    exit
#--------------------------------------------------
echo "Service Configuration"
#--------------------------------------------------
    service
        epipe 1 customer 1 create
```

```
            sap 2/1/2:1.1 create
            exit
            sap eth-tunnel-1 create
            exit
            no shutdown
        exit
        epipe 100 customer 1 create
            service-mtu 500
            sap 2/1/10:1.100 create
            exit
            sap eth-tunnel-2 create
            exit
            no shutdown
        exit
```

Configuration with Fate Sharing on Node-3

In this example, the SAPs monitoring the operational groups do not need CFM if the corresponding SAP on the CE side is using fate sharing.

Node-1:

```
#--------------------------------------------------
echo "Service Configuration"  Oper-groups
#--------------------------------------------------
    service
        customer 1 create
            description "Default customer"
        exit
        sdp 2 mpls create
          ...

        exit
        pw-template 1 create
            vc-type vlan
        exit
        oper-group "og-name-et" create
        exit
        epipe 1 customer 1 create
            service-mtu 500
            bgp
                route-distinguisher 65000:1
                route-target export target:65000:1 import target:65000:1
            exit
            site "site-1" create
                site-id 1
                sap 1/1/2:1.1
                boot-timer 100
                site-activation-timer 2
                no shutdown
            exit
            endpoint "x" create
            exit
            endpoint "y" create
            exit
```

```
            sap 1/1/2:1.1 endpoint "x" create
                eth-cfm
                    mep 130 domain 100 association 1 direction down
                        fault-propagation-enable use-if-tlv
                        ccm-enable
                        no shutdown
                    exit
                exit
                oper-group "og-name-et"
            exit
            spoke-sdp 2:1 endpoint "y" create
                precedence primary
                no shutdown
            exit
            spoke-sdp 3:1 endpoint "y" create
                precedence 2
                no shutdown
            exit
            no shutdown
        exit
        epipe 2 customer 1 create
            description "Epipe 2 in opergroup with Epipe 1"
            service-mtu 500
            bgp
                route-distinguisher 65000:2
                route-target export target:65000:2 import target:65000:2
            exit
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap 1/1/2:1.2 endpoint "x" create
                monitor-oper-group "og-name-et"
            exit
            spoke-sdp 2:2 vc-type vlan endpoint "y" create
                precedence 1
                no shutdown
            exit
            spoke-sdp 3:2 vc-type vlan endpoint "y" create
                precedence 2
                no shutdown
            exit
            no shutdown
        exit

    exit
```

Node-3:

```
#--------------------------------------------------
echo "Eth-CFM Configuration"
#--------------------------------------------------
    eth-cfm
        domain 100 format none level 3
            association 1 format icc-based name "node-3-site-1-0"
                bridge-identifier 1
                exit
                ccm-interval 1
```

```
                remote-mepid 130
            exit
            association 2 format icc-based name "node-3-site-2-0"
                bridge-identifier 2
                exit
                ccm-interval 1
                remote-mepid 120
            exit
        exit
    exit

#--------------------------------------------------
echo "Service Configuration"
#--------------------------------------------------

  eth-tunnel 2
        description "Eth Tunnel loadsharing mode QinQ example"
        protection-type loadsharing
        ethernet
            encap-type qinq
        exit
        path 1
            member 1/1/1
            control-tag 1.1
            eth-cfm
                mep 310 domain 100 association 1
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        path 2
            member 1/1/2
            control-tag 1.1
            eth-cfm
                mep 320 domain 100 association 2
                    ccm-enablepath
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        no shutdown
    exit


#--------------------------------------------------
echo "Service Configuration"
#--------------------------------------------------
    service
        epipe 1 customer 1 create
            sap 1/10/1:1 create
            exit
            sap eth-tunnel-1 create
            exit
            no shutdown
```

```
        exit
#--------------------------------------------------
echo "Service Configuration for a shared fate Ethernet Tunnel"
#--------------------------------------------------
        epipe 2 customer 1 create
            sap 1/10/2:3 create
            exit
            sap eth-tunnel-1:2 create
                eth-tunnel
                    path 1 tag 1.2
                    path 2 tag 1.2
                exit
            exit
            no shutdown
        exit
```

2.8.11 Access Node Resilience Using MC-LAG and Pseudowire Redundancy

Figure 34 shows the use of both Multi-Chassis Link Aggregation (MC-LAG) in the access network and pseudowire redundancy in the core network to provide a resilient end-to-end VLL service to the customers.

Figure 34 **Access Node Resilience**

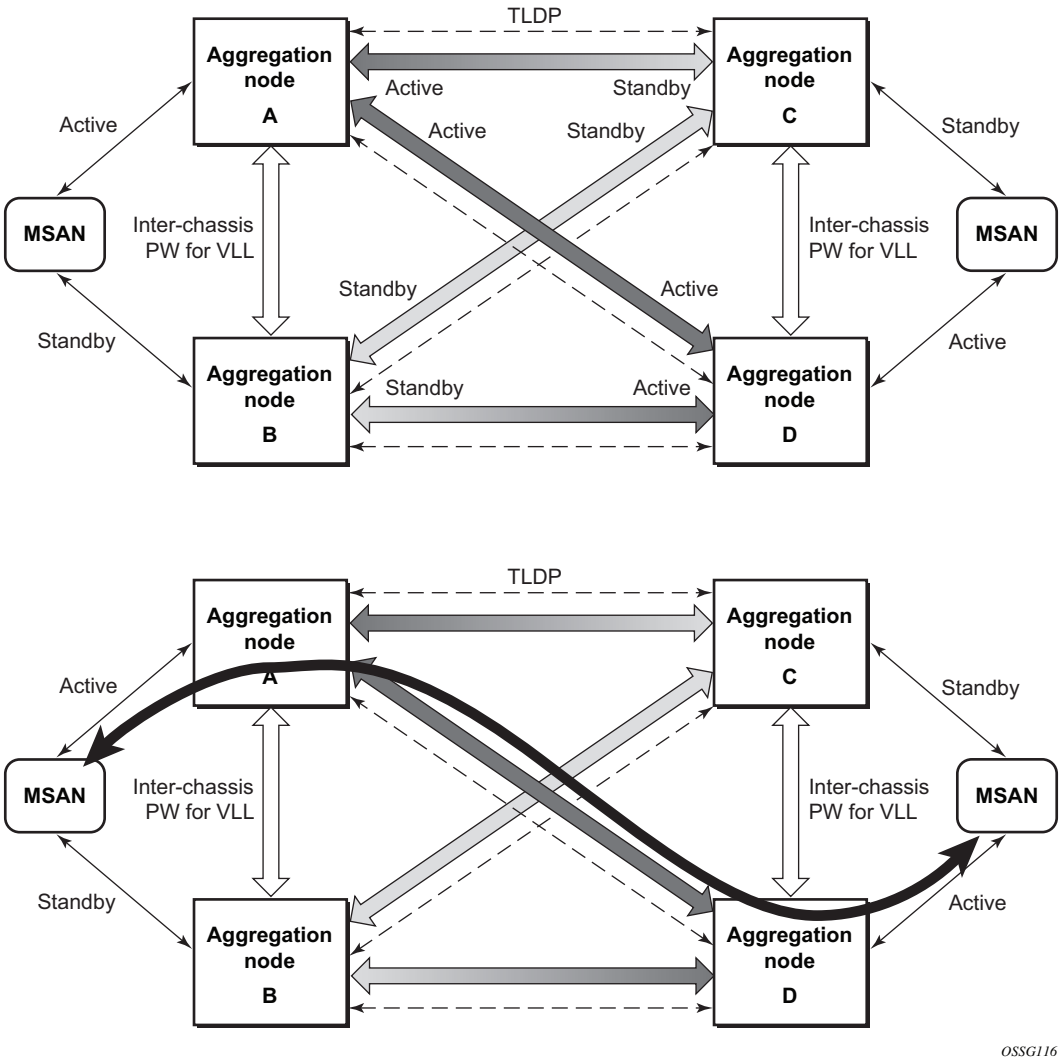

In this application, a new pseudowire status bit of active or standby indicates the status of the SAP in the MC-LAG instance in the SR-series aggregation node. All spoke-SDPs are of secondary type and there is no use of a primary pseudowire type in this mode of operation. Node A is in the active state according to its local MC-LAG instance and therefore advertises active status notification messages to both its peer pseudowire nodes; for example, nodes C and D. Node D performs the same operation. Node B is in the standby state according to the status of the SAP in its local MC-LAG instance, so advertises standby status notification messages to both nodes C and D. Node C performs the same operation.

An SR-series node selects a pseudowire as the active path for forwarding packets when both the local pseudowire status and the received remote pseudowire status indicate active status. However, an SR-series device in standby status according to the SAP in its local MC-LAG instance is capable of processing packets for a VLL service received over any of the pseudowires that are up. This is to avoid black holing of user traffic during transitions. The SR-series standby node forwards these packets to the active node by the Inter-Chassis Backup pseudowire (ICB pseudowire) for this VLL service. An ICB is a spoke-SDP used by an MC-LAG node to back up an MC-LAG SAP during transitions. The same ICB can also be used by the peer MC-LAG node to protect against network failures causing the active pseudowire to go down.

At configuration time, the user specifies a precedence parameter for each of the pseudowires that are part of the redundancy set, as described in the application in VLL Resilience with Two Destination PE Nodes. An SR-series node uses this to select which pseudowire to forward packets to in case both pseudowires show active/active for the local/remote status during transitions.

Only VLL service of type Epipe is supported in this application. Also, ICB spoke-SDP can only be added to the SAP side of the VLL cross-connect if the SAP is configured on an MC-LAG instance.

2.8.12 VLL Resilience for a Switched Pseudowire Path

Figure 35 illustrates the use of both pseudowire redundancy and pseudowire switching to provide a resilient VLL service across multiple IGP areas in a provider network.

Figure 35 **VLL Resilience with Pseudowire Redundancy and Switching**

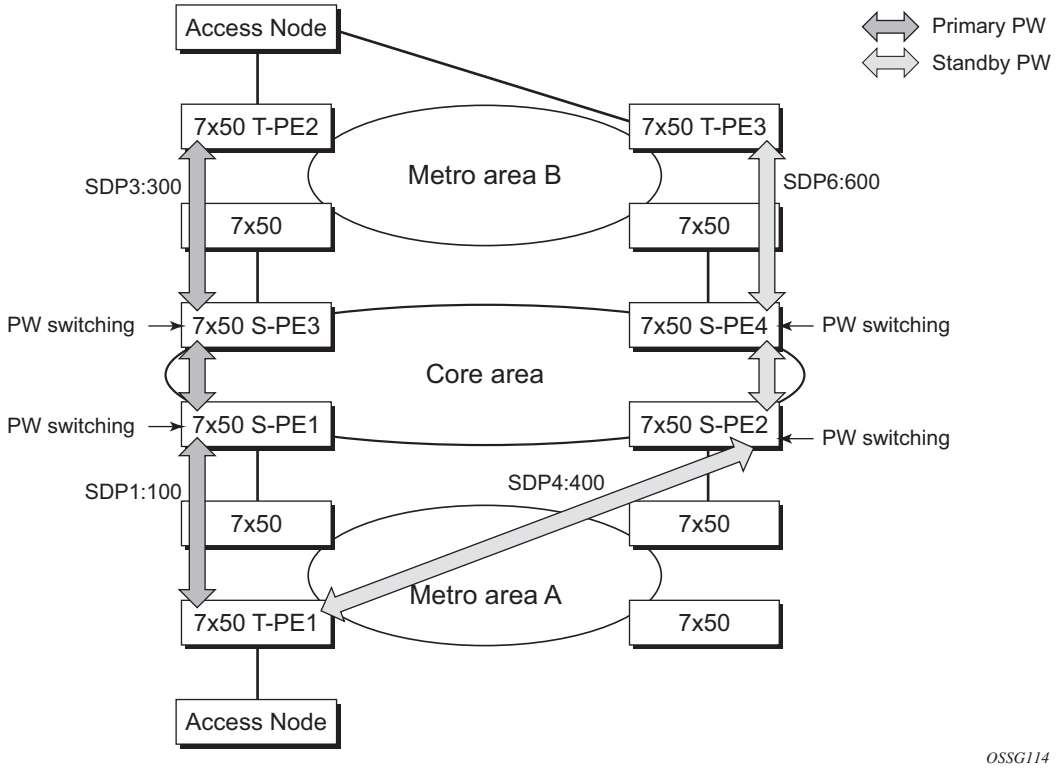

Pseudowire switching is a method for scaling a large network of VLL or VPLS services by removing the need for a full mesh of T-LDP sessions between the PE nodes as the number of these nodes grows over time.

Like in the application in VLL Resilience with Two Destination PE Nodes, the T-PE1 node switches the path of a VLL to a secondary standby pseudowire if a network side failure caused the VLL binding status to be operationally down or if T-PE2 notified it that the remote SAP went down. This application requires that pseudowire status notification messages generated by either a T-PE node or a S-PE node be processed and relayed by the S-PE nodes.

It is possible that the secondary pseudowire path terminates on the same target PE as the primary; for example, T-PE2. This provides protection against network side failures but not against a remote SAP failure. When the target destination PE for the primary and secondary pseudowires is the same, T-PE1 will normally not switch the VLL path onto the secondary pseudowire upon receipt of a pseudowire status notification indicating the remote SAP is down, because the status notification is sent over both the primary and secondary pseudowires. However, the status notification on the primary pseudowire may arrive earlier than the one on the secondary pseudowire due to the differential delay between the paths. This will cause T-PE1 to

switch the path of the VLL to the secondary standby pseudowire and remain there until the status notification is cleared. Then, the VLL path is switched back to the primary pseudowire due to the revertive behavior operation. The path will not switch back to a secondary path when it comes up, even if it has a higher precedence than the currently active secondary path.

For the 7750 SR, this application can make use of all types of VLL supported on the routers; for example, Apipe, Fpipe, Epipe, and Ipipe services. A SAP can be configured on a SONET/SDH port that is part of an APS group. However, if a SAP is configured on an MC-LAG instance, only the Epipe service type will be allowed.

2.9 Pseudowire Redundancy Service Models

This section describes the various MC-LAG and pseudowire redundancy scenarios as well as the algorithm used to select the active transmit object in a VLL endpoint.

The redundant VLL service model is described in the following section, Redundant VLL Service Model.

2.9.1 Redundant VLL Service Model

To implement pseudowire redundancy, a VLL service accommodates more than a single object on the SAP side and on the spoke-SDP side. Figure 36 illustrates the model for a redundant VLL service based on the concept of endpoints.

Figure 36 **Redundant VLL Endpoint Objects**

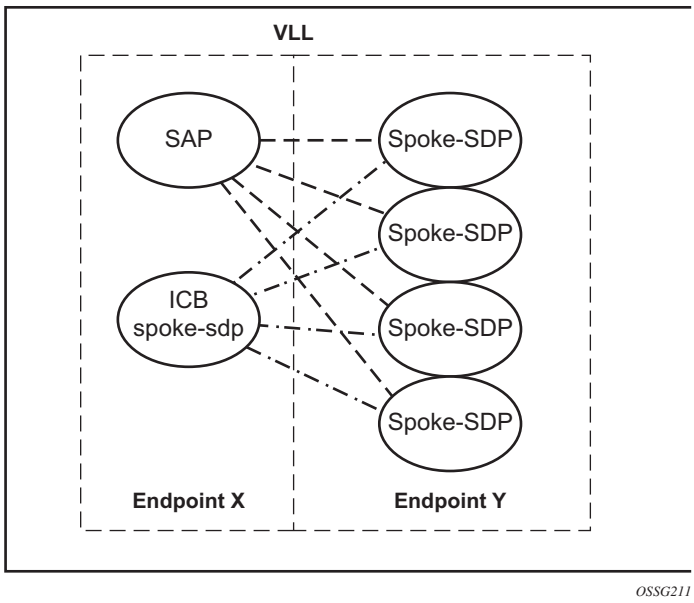

By default a VLL service supports two implicit endpoints managed internally by the system. Each endpoint can only have one object: a SAP or a spoke-SDP.

To add more objects, up to two explicitly named endpoints may be created per VLL service. The endpoint name is locally significant to the VLL service. They are referred to as endpoint X and endpoint Y as illustrated in Figure 36.

Figure 36 is just an example; the Y endpoint can also have a SAP and/or an ICB spoke-SDP. The following describes the four types of endpoint objects supported and the rules used when associating them with an endpoint of a VLL service:

- SAP — There can only be a maximum of one SAP per VLL endpoint.
- Primary spoke-SDP — The VLL service always uses this pseudowire and only switches to a secondary pseudowire when this primary pseudowire is down; the VLL service switches the path to the primary pseudowire when it is back up. The user can configure a timer to delay reverting back to primary or to never revert. There can only be a maximum of one primary spoke-SDP per VLL endpoint.
- Secondary spoke-SDP — There can be a maximum of four secondary spoke-SDPs per endpoint. The user can configure the precedence of a secondary pseudowire to indicate the order in which a secondary pseudowire is activated.
- Inter-Chassis Backup (ICB) spoke-SDP — This special pseudowire is used for MC-LAG and pseudowire redundancy applications. Forwarding between ICBs is blocked on the same node. The user has to explicitly indicate that the spoke-SDP is an ICB, at creation time. However, following are a few scenarios where the user can configure the spoke-SDP as ICB or as a regular spoke-SDP on a specified node. The CLI for those cases will indicate both options.

A VLL service endpoint can only use a single active object to transmit at any specified time but can receive from all endpoint objects.

An explicitly named endpoint can have a maximum of one SAP and one ICB. When a SAP is added to the endpoint, only one more object of type ICB spoke-SDP is allowed. The ICB spoke-SDP cannot be added to the endpoint if the SAP is not part of an MC-LAG instance. Conversely, a SAP that is not part of an MC-LAG instance cannot be added to an endpoint that already has an ICB spoke-SDP.

An explicitly named endpoint that does not have a SAP object, can have a maximum of four spoke-SDPs and can include any of the following:

- a single primary spoke-SDP
- one or many secondary spoke-SDPs with precedence
- a single ICB spoke-SDP

2.9.2 T-LDP Status Notification Handling Rules

Using Redundant VLL Endpoint Objects as a reference, the following are the rules for generating, processing, and merging T-LDP status notifications in VLL service with endpoints. Any allowed combination of objects as specified in Redundant VLL Service Model can be used on endpoints X and Y. The following sections see the specific combination objects in Figure 36 as an example to describe the more general rules.

2.9.2.1 Processing Endpoint SAP Active/Standby Status Bits

The advertised admin forwarding status of active/standby reflects the status of the local LAG SAP in MC-LAG application. If the SAP is not part of an MC-LAG instance, the forwarding status of active is always advertised.

When the SAP in endpoint X is part of an MC-LAG instance, a node must send a T-LDP forwarding status bit of SAP active/standby over all Y endpoint spoke-SDPs, except the ICB spoke-SDP, whenever this status changes. The status bit sent over the ICB is always zero (active by default).

When the SAP in endpoint X is not part of an MC-LAG instance, then the forwarding status sent over all Y endpoint spoke-SDPs should always be set to zero (active by default).

2.9.2.2 Processing and Merging

Endpoint X is operationally up if at least one of its objects is operationally up. It is down if all of its objects are operationally down.

If the SAP in endpoint X transitions locally to the down state, or received a SAP down notification by SAP-specific OAM signal, the node must send T-LDP SAP down status bits on the Y endpoint ICB spoke-SDP only. Ethernet SAP does not support SAP OAM protocol. All other SAP types cannot exist on the same endpoint as an ICB spoke-SDP because a non-Ethernet SAP cannot be part of an MC-LAG instance.

If the ICB spoke-SDP in endpoint X transitions locally to down state, the node must send T-LDP SDP-binding down status bits on this spoke-SDP.

If the ICB spoke-SDP in endpoint X received T-LDP SDP-binding down status bits or pseudowire not forwarding status bits, the node saves this status and takes no further action. The saved status is used for selecting the active transmit endpoint object.

If all objects in endpoint X did any or all of the following, the node must send status bits of SAP down over all “Y” endpoint spoke-SDPs, including the ICB:

- transitioned locally to down state
- received a SAP down notification by remote T-LDP status bits or by SAP-specific OAM signal
- received SDP-binding down status bits
- received PW not forwarding status bits

Endpoint Y is operationally up if at least one of its objects is operationally up. It is down if all its objects are operationally down.

If a spoke-SDP in endpoint Y, including the ICB spoke-SDP, transitions locally to down state, the node must send T-LDP SDP-binding down status bits on this spoke-SDP.

If a spoke-SDP in endpoint Y, including the ICB spoke-SDP, did any or all of the following, the node saves this status and takes no further action:

- received T-LDP SAP down status bits
- received T-LDP SDP-binding down status bits
- received PW not forwarding status bits

The saved status is used for selecting the active transmit endpoint object.

If all objects in endpoint Y, except the ICB spoke-SDP, did any or all of the following, the node must send status bits of SDP-binding down over the X endpoint ICB spoke-SDP only:

- transitioned locally to the down state
- received T-LDP SAP down status bits
- received T-LDP SDP-binding down status bits
- received PW not forwarding status bits

If all objects in endpoint Y did any or all of the following, he node must send status bits of SDP-binding down over the X endpoint ICB spoke-SDP, and must send a SAP down notification on the X endpoint SAP by the SAP-specific OAM signal if applicable:

- transitioned locally to down state

- received T-LDP SAP down status bits
- received T-LDP SDP-binding down status bits
- received PW not forwarding status bits

An Ethernet SAP does not support signaling status notifications.

2.10 High-Speed Downlink Packet Access (HSDPA) Off Load Fallback over ATM

For many Universal Mobile Telecommunications System (UMTS) networks planning to deploy High-Speed Downlink Packet Access (HSDPA), the existing mobile backhaul topology consists of a cell site that is partially backhauled over DSL (for the HSDPA portion) and partially over an existing TDM/ATM infrastructure (for UMTS voice traffic).

Figure 37 **HSDPA Off Load Fallback over ATM**

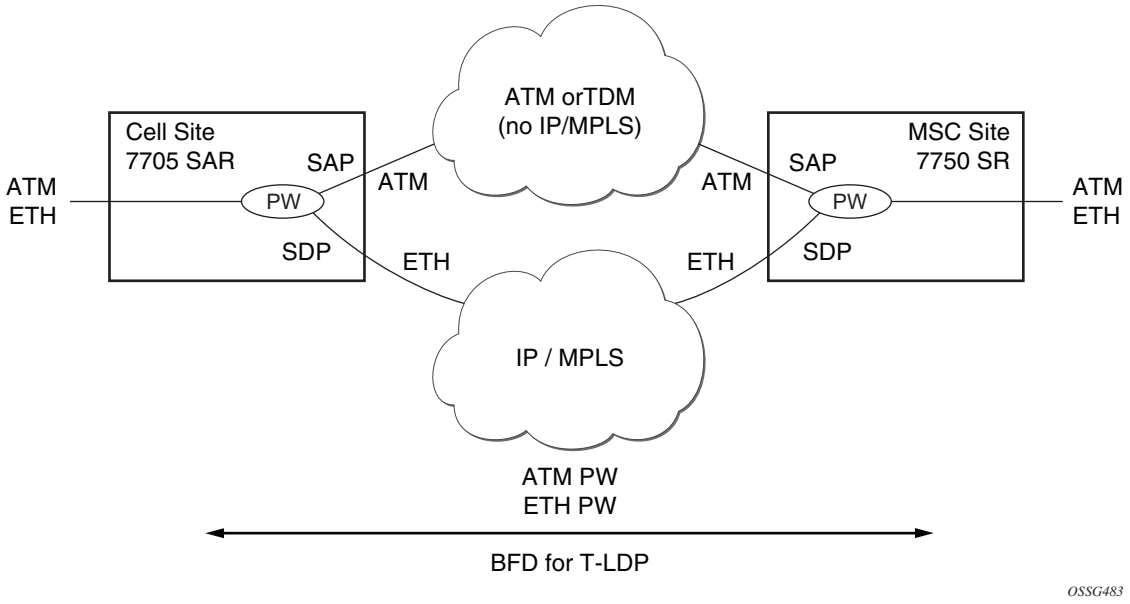

For example, the service pseudowires provider may use a 7705 SAR with one or two ATM E1 uplinks for real-time voice traffic and an Ethernet uplink connected to a DSL model for NRT data traffic. At the RNC site, a 7750 SR service router can be used, connected by ASAP (E1 IMA bundles) or STM-n ATM to the TDM/ATM network, and Ethernet to the DSL backhaul network.

On the MSC-located SR connected to the Radio Network Controller (RNC), there is a standard pseudowire (Ethernet or ATM) that has an active pseudowire by IP/MPLS, but the standby path is not IP/MPLS capable Therefore, the active/standby pseudowire concept is extended to allow standby to be an access SAP to an ATM network for ATM pseudowire or Ethernet (bridged over ATM) for ETH pseudowire.

Normally, if the MPLS pseudowire path is active, this path is used. If a failure happens on the IP/MPLS path, detected through BFD-TLDP or local notification, traffic needs to switch to the SAP that is connected to the ATM/TDM backhaul network. As soon as the MPLS pseudowire path becomes available again, reversion back to the pseudowire path is supported.

2.10.1 Primary Spoke SDP Fallback to Secondary SAP

For HSDPA, Apipe and Epipe service termination on the SR where an endpoint-X SAP connects to the mobile RNC (by ATM or Ethernet) and an endpoint Y has a primary spoke-SDP and a secondary SAP on an SR ATM or ASAP MDA (with bridged PDU encapsulation for Epipes). The secondary SAP has the same restrictions as the SAP in endpoint-X for Apipe and Epipe, respectively.

It is sufficient to have a single secondary SAP (without any precedence), which implies that it cannot be mixed with any secondary spoke-SDPs; 1+1 APS and MC-APS is supported on the secondary SAP interface.

Similar to the current pseudowire redundancy implementation, receive should be enabled on both objects even though transmit is only enabled on one.

It is expected that BFD for T-LDP will be used in most applications to decrease the fault detection times and minimize the outage times upon failure.

2.10.2 Reversion to Primary Spoke SDP Path

The **endpoint revert-time** reversion from secondary to primary paths in the **config**>**service**>**apipe**>**endpoint** and **config**>**service**>**epipe**>**endpoint** contexts are consistent with standard pseudowire redundancy. Various network configurations and equipment require different reversion configurations. The default revert-time is 0.

2.10.3 MC-APS and MC-LAG

In many cases, 7750 SRs are deployed in redundant pairs at the MSC. In this case, MC-APS is typically used for all ATM connections. Figure 38 illustrates this case, assuming that MC-APS is deployed on both the RNC connection and the ATM network connection. For MC-APS to be used, clear channel SONET or SDH connections should be used.

Figure 38 **HSDPA Off Load Fallback with MC-APS**

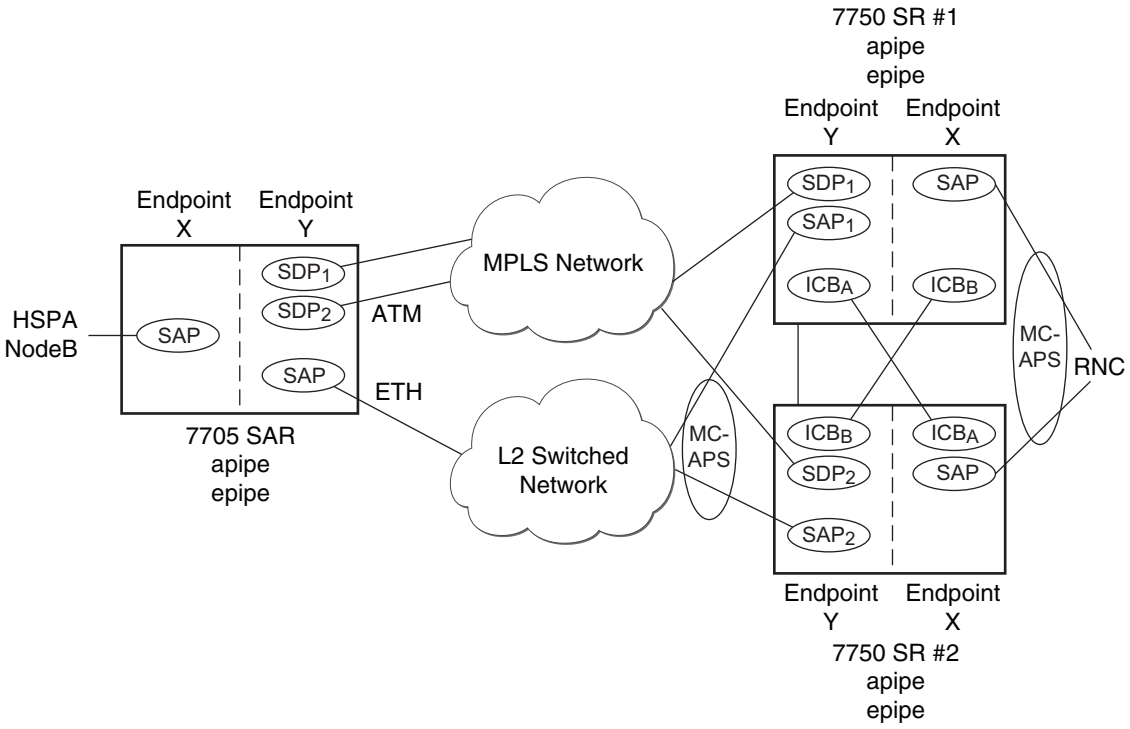

In this scenario, endpoint Y allows an ICB spoke-SDP as well as the primary spoke-SDP and secondary SAP. ICB operation is maintained as the current redundant pseudowire operation and the ICB spoke-SDP is always given an active status. The ICB spoke-SDP is only used if both the primary spoke-SDP and secondary SAP are not available. The secondary SAP is used if it is operationally up and the primary spoke-SDP pseudowire status is not active. Receive is enabled on all objects even though transmit is only enabled on one.

To allow correct operation in all failure scenarios, an ICB spoke-SDP must be added to endpoint X. The ICB spoke-SDP is only used if the SAP is operationally down.

The following is an example configuration of Epipes mapping to Figure 38. A SAP can be added to an endpoint with a non-ICB spoke-SDP only if the precedence of the spoke is **primary**.

7750 SR #1

```
*A:ALA-A>config>service#  epipe 1
------------------------------------------------
            endpoint X
            exit
            endpoint Y
            exit
```

```
            sap 1/1/2:0 endpoint X
              exit
            spoke-sdp 1:100 endpoint X icb
              exit
            spoke-sdp 10:500 endpoint Y
              precedence primary
              exit
            sap 1/1/3:0 endpoint Y
            exit
            spoke-sdp 1:200 endpoint Y icb
            exit
-----------------------------------------------
*A:ALA-A>config>service#
```

7750 SR #2

```
*A:ALA-B>config>service# epipe 1
-----------------------------------------------
            endpoint X
            exit
            endpoint Y
            exit
            sap 2/3/4:0 endpoint X
            exit
            spoke-sdp 1:200 endpoint X icb
            exit
            spoke-sdp 20:600 endpoint Y
              precedence primary
              exit
            sap 2/3/5:0 endpoint Y
            exit
            spoke-sdp 1:100 endpoint Y icb
            exit
-----------------------------------------------
*A:ALA-B>config>service#
```

2.10.3.1 Failure Scenario

Based on the previously mentioned rules, the following is an example of a failure scenario. Assuming both links are active on 7750 SR #1 and the Ethernet connection to the cell site fails (most likely failure scenario because the connection would not be protected), SDP1 would go down and the secondary SAP would be used in 7750 SR #1 and 7705 SAR, as shown in Figure 39.

Figure 39 **Ethernet Failure At Cell Site**

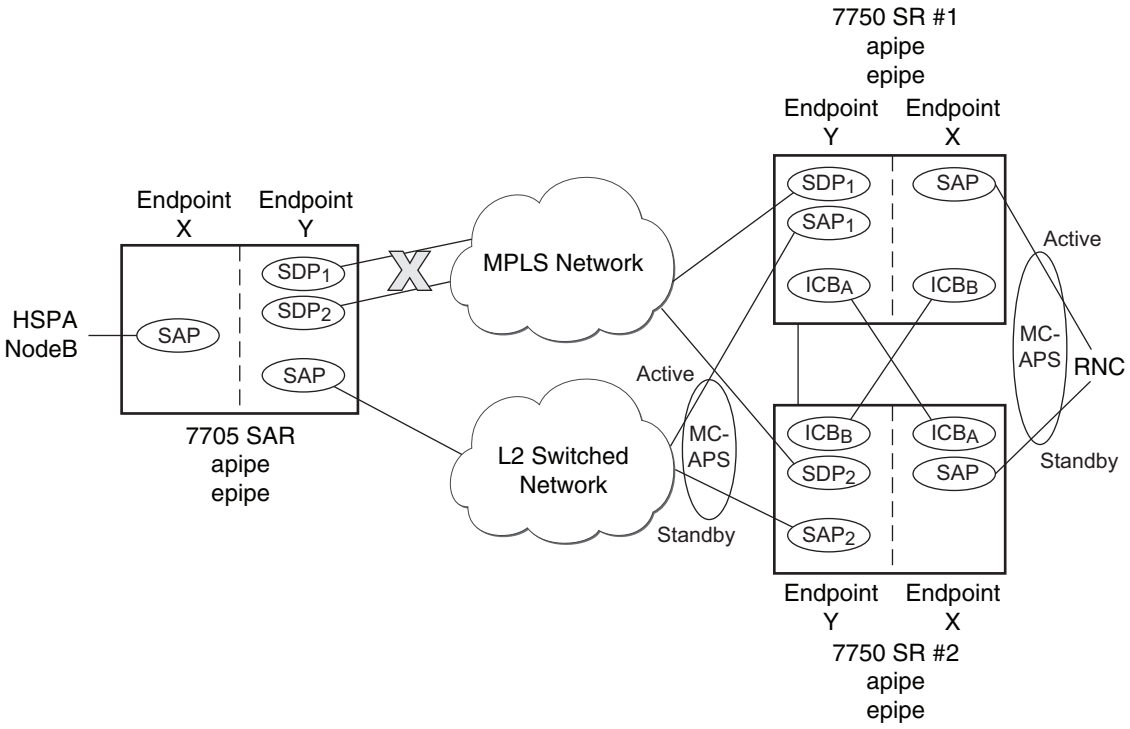

If the active link to the Layer 2 switched network was on 7750 SR #2 at the time of the failure, SAP1 would be operationally down (because the link is in standby) and ICBA would be used. Because the RNC SAP on 7750 SR #2 is on a standby APS link, ICBA would be active and it would connect to SAP2 because SDP2 is operationally down as well.

All APS link failures would be handled through the standard pseudowire status messaging procedures for the RNC connection and through standard ICB usage for the Layer 2 switched network connection.

2.11 VLL Using G.8031 Protected Ethernet Tunnels

The use of MPLS tunnels provides the 7450 ESS and 7750 SR OS a way to scale the core while offering fast failover times using MPLS FRR. In environments where Ethernet services are deployed using native Ethernet backbones, Ethernet tunnels are provided to achieve the same fast failover times as in the MPLS FRR case.

The Nokia VLL implementation offers the capability to use core Ethernet tunnels compliant with ITU-T G.8031 specification to achieve 50 ms resiliency for backbone failures. This is required to comply with the stringent SLAs provided by service providers. Epipe and Ipipe services are supported.

When using Ethernet tunnels, the Ethernet tunnel logical interface is created first. The Ethernet tunnel has member ports, which are the physical ports supporting the links. The Ethernet tunnel control SAPs carry G.8031 and 802.1ag control traffic and user data traffic. Ethernet service SAPs are configured on the Ethernet tunnel. Optionally, when tunnels follow the same paths, end-to-end services may be configured with fate shared Ethernet tunnel SAPs, which carry only user data traffic and share the fate of the Ethernet tunnel port (if correctly configured).

Ethernet tunnels provide a logical interface that VLL SAPs may use just as regular interfaces. The Ethernet tunnel provides resiliency by providing end-to-end tunnels. The tunnels are stitched together by VPLS or Epipe services at intermediate points. Epipes offer a more scalable option.

For further information, refer to *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide*.

2.12 MPLS Entropy Label and Hash Label

The router supports the MPLS entropy label (RFC 6790) and the Flow Aware Transport label, known as the hash label (RFC 6391). These labels allow LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by just hashing on the standard label stack. See the *7450 ESS, 7750 SR, and 7950 XRS MPLS Guide* for further information.

2.13 BGP Virtual Private Wire Service (VPWS)

BGP Virtual Private Wire Service (VPWS) is a point-to-point L2 VPN service based on RFC 6624 *Layer 2 Virtual Private Networks using BGP for Auto-Discovery and Signaling* which in turn uses the BGP pseudowire signaling concepts from RFC 4761, *Virtual Private LAN Service Using BGP for Auto-Discovery and Signaling*.

The BGP signaled pseudowires created can use either automatic or preprovisioned SDPs over LDP- or BGP-signaled tunnels (the choice of tunnel depends on the tunnel's preference in the tunnel table). Preprovisioned SDPs must be configured whenever GRE or RSVP signaled transport tunnels are used.

Inter-AS model C and dual-homing are supported.

2.13.1 Single-Homed BGP VPWS

A single-homed BGP VPWS service is implemented as an Epipe connecting a SAP or static GRE tunnel (a spoke-SDP using a GRE SDP configured with static MPLS labels) and a BGP signaled pseudowire, maintaining the Epipe properties such as no MAC learning. The pseudowire data plane uses a two-label stack; the inner label is derived from the BGP signaling and identifies the Epipe service while the outer label is the tunnel label of an LSP transporting the traffic between the two end systems.

Figure 40 shows how this service would be used to provide a virtual leased line service (VLL) across an MPLS network between two sites: A and B.

Figure 40 **Single-Homed BGP-VPWS Example**

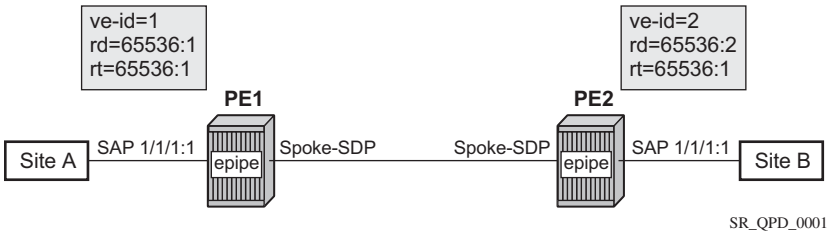

An Epipe is configured on PE1 and PE2 with BGP VPWS enabled. PE1 and PE2 are connected to site A and B, respectively, each using a SAP. The interconnection between the two PEs is achieved through a pseudowire that is signaled using BGP VPWS updates over a specific tunnel LSP.

2.13.2 Dual-Homed BGP VPWS

A BGP-VPWS service can benefit from dual-homing, as described in IETF Draft *draft-ietf-bess-vpls-multihoming-01*. When using dual-homing, two PEs connect to a site, with one PE being the designated forwarder for the site and the other blocking its connection to the site. On failure of the active PE, its pseudowire, or its connection to the site, the other PE becomes the designated forwarder and unblocks its connection to the site.

2.13.2.1 Single Pseudowire Example

A pseudowire is established between the designated forwarder of the dual-homed PEs and the remote PE. If a failure causes a change in the designated forwarder, the pseudowire is deleted and reestablished between the remote PE and the new designated forwarder. This topology requires that the VE IDs on the dual-homed PEs are set to the same value.

A dual-homed, single pseudowire topology example is shown in Figure 41.

Figure 41 **Dual-Homed BGP VPWS with Single Pseudowire**

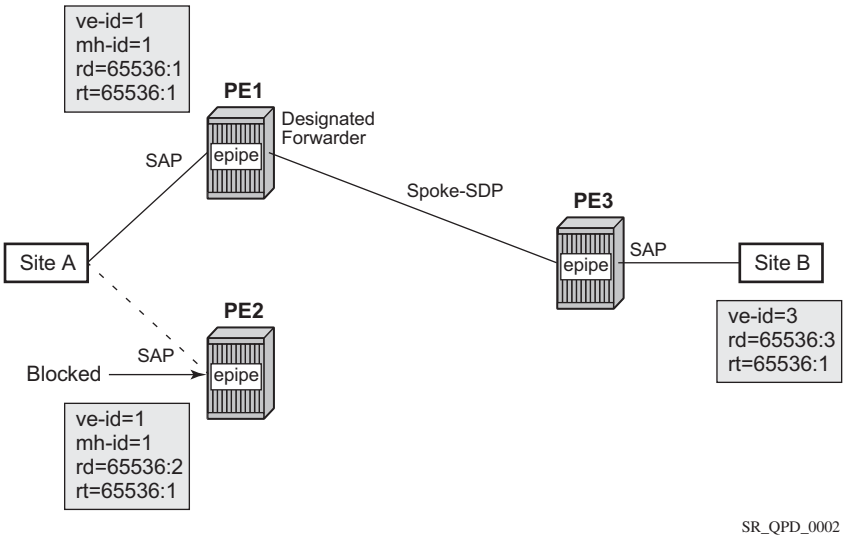

An Epipe with BGP VPWS enabled is configured on each PE. Site A is dual-homed to PE1 and PE2 with the remote PE (PE3) connecting to site B. An Epipe service is configured on each PE in which there is a SAP connecting to the local site.

The pair of dual-homed PEs perform a designated forwarder election, which is influenced by BGP route selection, the site state, and by configuring the site-preference. A site will only be eligible to be the designated forwarder if it is up (the site state will be down if there is no pseudowire established or if the pseudowire is in an oper down state). The winner, for example PE1, becomes the active switch for traffic sent to and from site A, while the loser blocks its connection to site A.

Pseudowires are signaled using BGP from PE1 and PE2 to PE3, but only from PE3 to the designated forwarder in the opposite direction (so only one bi-directional pseudowire is established). There is no pseudowire between PE1 and PE2; this is achieved by configuration.

Traffic is sent and received traffic on the pseudowire connected between PE3 and the designated forwarder, PE1.

If the site state is oper down, both the D and Circuit Status Vector (CSV) bits (see the following for more details) are set in the BGP-VPWS update which will cause the remote PE to use the pseudowire to the new designated forwarder.

2.13.2.2 Active/Standby Pseudowire Example

Pseudowires are established between the remote PE and each dual-homed PE. The remote PE can receive traffic on either pseudowire, but will only send on the one to the designated forwarder. This creates an active/standby pair of pseudowires. At most, one standby pseudowire will be established; this being determined using the tie-breaking rules defined in the multi-homing draft. This topology requires each PE to have a different VE ID.

A dual-homed, active/standby pseudowires topology example is shown in Figure 42.

Figure 42 **Dual-homed BGP VPWS with Active/Standby Pseudowires**

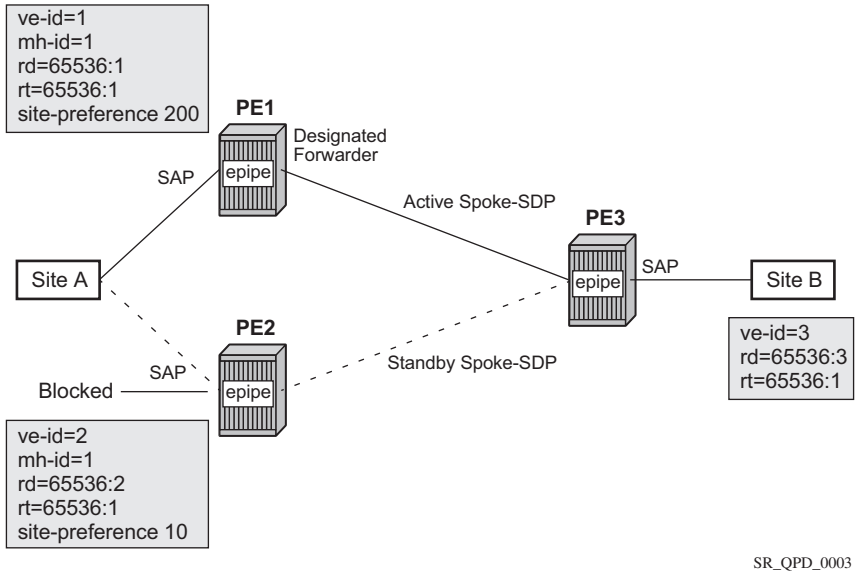

An Epipe with BGP VPWS enabled is configured on each PE. Site A is dual-homed to PE1 and PE2 with the remote PE (PE3) connecting to site B. An Epipe service is configured on each PE in which there is a SAP connecting to the local site.

The pair of dual-homed PEs perform a designated forwarder election, which is influenced by configuring the site-preference. The winner, PE1 (based on its higher site-preference) becomes the active switch for traffic sent to and from site A, while the loser, PE2, blocks its connection to site A. Pseudowires are signaled using BGP between PE1 and PE3, and between PE2 and PE3. There is no pseudowire between PE1 and PE2; this is achieved by configuration. The active/standby pseudowires on PE3 are part of an endpoint automatically created in the Epipe service.

Traffic is sent and received on the pseudowire connected to the designated forwarder, PE1.

2.13.3 BGP VPWS Pseudowire Switching

Pseudowire switching is supported with a BGP VPWS service allowing the cross connection between a BGP VPWS signaled spoke-SDP and a static GRE tunnel, the latter being a spoke-SDP configured with static MPLS labels using a GRE SDP. No other spoke-SDP types are supported. Support is not included for BGP multi-homing using an active and a standby pseudowire to a pair of remote PEs.

Operational state changes to the GRE tunnel are reflected in the state of the Epipe and propagated accordingly in the BGP VPWS spoke-SDP status signaling, specifically using the BGP update D and CSV bits.

The following configuration is required:

1. The Epipe service must be created using the **vc-switching** parameter.
2. The GRE tunnel spoke-SDP must be configured using a GRE SDP with **signaling off**, and have the ingress and egress vc-labels statically configured.

An example configuration is as follows:

```
configure
    service
        sdp 1 create
            signaling off
            far-end 192.168.1.1
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        pw-template 1 create
        exit
        epipe 1 customer 1 vc-switching create
            description "BGP VPWS service"
            bgp
                route-distinguisher 65536:1
                route-target export target:65536:1 import target:65536:1
                pw-template-binding 1
                exit
            exit
            bgp-vpws
                ve-name "PE1"
                    ve-id 1
                exit
                remote-ve-name "PE2"
                    ve-id 2
                exit
                no shutdown
            exit
            spoke-sdp 1:1 create
                ingress
                    vc-label 1111
                exit
                egress
                    vc-label 1122
                exit
                no shutdown
            exit
            no shutdown
        exit
```

2.13.3.1 Pseudowire Signaling

The BGP signaling mechanism used to establish the pseudowires is described in the BGP VPWS standards with the following differences:

- As stated in Section 3 of RFC 6624, there are two modifications of messages when compared to RFC 4761.
  - the Encaps Types supported in the associated extended community
  - the addition of a circuit status vector sub-TLV at the end of the VPWS NLRI
- The control flags and VPLS preference in the associated extended community are based on IETF Draft *draft-ietf-bess-vpls-multihoming-01*.

Figure 43 shows the format of the BGP VPWS update extended community.

Figure 43 **BGP VPWS Update Extended Community Format**

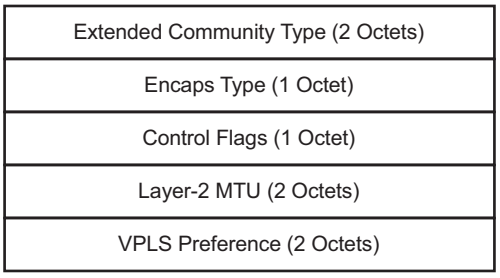

- Extended Community Type — The value allocated by IANA for this attribute is 0x800A.
- Encaps Type — Encapsulation type, identifies the type of pseudowire encapsulation. Ethernet VLAN (4) and Ethernet Raw mode (5), as described in RFC 4448, are the only values supported. If there is a mismatch between the Encaps Type signaled and the one received, the pseudowire is created but with the operationally down state.
- Control Flags — Control information regarding the pseudowires, see Figure 44 for more information.
- Layer 2 MTU — Maximum Transmission Unit to be used on the pseudowires. If the received Layer 2 MTU is zero, no MTU check is performed and the related pseudowire is established. If there is a mismatch between the local service-mtu and the received Layer 2 MTU, the pseudowire is created with the operationally down state and an MTU/Parameter mismatch indication.

- VPLS Preference – VPLS preference has a default value of zero for BGP-VPWS updates sent by the system, indicating that it is not in use. If the site-preference is configured, its value is used for the VPLS preference and is also used in the local designated forwarder election.

  On receipt of a BGP VPWS update containing a non-zero value, it will be used to determine to which system the pseudowire is established, as part of the VPWS update process tie-breaking rules. The BGP local preference of the BGP VPWS update sent by the system is set to the same value as the VPLS preference if the latter is non-zero, as required by the draft (as long as the D bit in the extended community is not set to 1). Consequently, attempts to change the BGP local preference when exporting a BGP VPWS update with a non-zero VPLS preference will be ignored. This prevents the updates being treated as malformed by the receiver of the update.

  For inter-AS, the preference information must be propagated between autonomous systems using the VPLS preference. Consequently, if the VPLS preference in a BGP-VPWS or BGP multi-homing update is zero, the local preference is copied by the egress ASBR into the VPLS preference field before sending the update to the eBGP peer. The adjacent ingress ASBR then copies the received VPLS preference into the local preference to prevent the update from being considered malformed.

The control flags are shown in Figure 44.

Figure 44 **Control Flags**

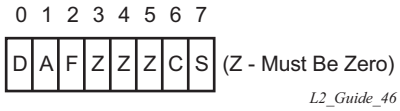

The following bits in the Control Flags are defined:

D — Access circuit down indicator from IETF Draft *draft-kothari-l2vpn-auto-site-id-01*. D is 1 if all access circuits are down, otherwise D is 0.

A — Automatic site ID allocation, which is not supported. This is ignored on receipt and set to 0 on sending.

F — MAC flush indicator. This is not supported because it relates to a VPLS service. This is set to 0 and ignored on receipt.

C — Presence of a control word. Control word usage is supported. When this is set to 1, packets will be sent and are expected to be received, with a control word. When this is set to 0, packets will be sent and are expected to be received, without a control word (by default).

S — Sequenced delivery. Sequenced delivery is not supported. This is set to 0 on sending (no sequenced delivery) and, if a non-zero value is received (indicating sequenced delivery required), the pseudowire will not be created.

The BGP VPWS NLRI is based on that defined for BGP VPLS, but is extended with a circuit status vector, as shown in Figure 45.

Figure 45 **BGP VPWS NLRI**

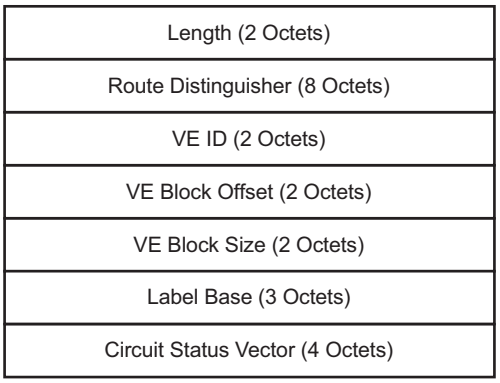

L2_Guide_43

The VE ID value is configured within each BGP VPWS service, the label base is chosen by the system, and the VE block offset corresponds to the remote VE ID because a VE block size of 1 is always used.

The circuit status vector is encoded as a TLV, as shown in Figure 46 and Figure 47.

Figure 46 **BGP VPWS NLRI TLV Extension Format**

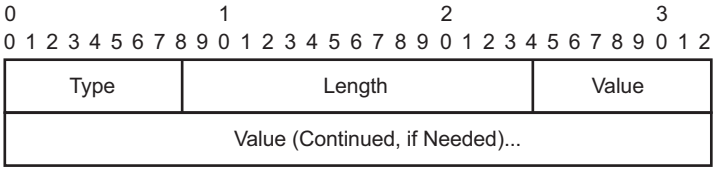

L2_Guide_44

Figure 47 **Circuit Status Vector TLV Type**

TLV Type	Description
1	Circuit Status Vector

L2_Guide_45

The circuit status vector is used to indicate the status of both the SAP/GRE tunnel and the status of the spoke-SDP within the local service. Because the VE block size used is 1, the most significant bit in the circuit status vector TLV value will be set to 1 if either the SAP/GRE tunnel or spoke-SDP is down, otherwise it will be set to 0. On receiving a circuit status vector, only the most significant byte of the CSV is examined for designated forwarder selection purposes.

If a circuit status vector length field of greater than 32 is received, the update will be ignored and not reflected to BGP neighbors. If the length field is greater than 800, a notification message will be sent and the BGP session will restart. Also, BGP VPWS services support a single access circuit, so only the most significant bit of the CSV is examined on receipt.

A pseudowire will be established when a BGP VPWS update is received that matches the service configuration, specifically the configured route-targets and remote VE ID. If multiple matching updates are received, the system to which the pseudowire is established is determined by the tie-breaking rules, as described in IETF Draft *draft-ietf-bess-vpls-multihoming-01*.

Traffic will be sent on the active pseudowire connected to the remote designated forwarder. Traffic can be received on either the active or standby pseudowire, although no traffic should be received on the standby pseudowire because the SAP/GRE tunnel on the non-designated forwarder should be blocked.

2.13.3.2 BGP-VPWS with Inter-AS Model C

BGP VPWS with inter-AS model C is supported both in a single-homed and dual-homed configuration.

When dual-homing is used, the dual-homed PEs must have different values configured for the **site-preference** (under the **site** within the Epipe service) to allow the PEs in a different AS to select the designated forwarder when all access circuits are up. The value configured for the **site-preference** is propagated between autonomous systems in the BGP VPWS and BGP multi-homing update extended community VPLS preference field. The receiving ingress ASBR copies the VPLS preference value into local preference of the update to ensure that the VPLS preference and local preference are equal, which prevents the update from being considered malformed.

2.13.3.3 BGP VPWS Configuration Procedure

In addition to configuring the associated BGP and MPLS infrastructure, the provisioning of a BGP VPWS service requires:

- Configuring the BGP Route Distinguisher, Route Target
  - Updates are accepted into the service only if they contain the configured import route-target
- Configuring a binding to the pseudowire template
  - Multiple pseudowire template bindings can be configured with their associated route-targets used to control which is applied
- Configuring the SAP or static GRE tunnel
- Configuring the name of the local VE and its associate VE ID
- Configuring the name of the remote VE and its associated VE ID
- For a dual-homed PE
  - Enabling the site
  - Configuring the site with non-zero site-preference
- For a remote PE
  - Configuring up to two remove VE names and associated VE IDs
- Enabling BGP VPWS

2.13.3.4 Use of Pseudowire Template for BGP VPWS

The pseudowire template concept used for BGP AD is re-used for BGP VPWS to dynamically instantiate pseudowires (SDP-bindings) and the related SDPs (provisioned or automatically instantiated).

The settings for the L2-Info extended community in the BGP Update sent by the system are derived from the pseudowire-template attributes. The following rules apply:

- If multiple pseudowire-template-bindings (with or without import-rt) are specified for the VPWS instance, the first (numerically lowest ID) pseudowire-template entry will be used.
- Both Ethernet VLAN and Ethernet Raw Mode Encaps Types are supported; these are selected by configuring the vc-type in the pseudowire template to be either vlan or ether, respectively. The default is ether.
  - The same value must be used by the remote BGP VPWS instance to ensure that the related pseudowire will come up.

- Layer 2 MTU – derived from service VPLS **service-mtu** parameter.
  - The same value must be used by the remote BGP VPWS instance to ensure that the related pseudowire will come up.
- Control Flag C – can be 0 or 1, depending on the setting of the *controlword* parameter in the pw-template 0.
- Control Flag S – always 0.

On reception, the values of the parameters in the L2-Info extended community of the BGP update are compared with the settings from the corresponding pseudowire-template. The following steps are used to determine the local pseudowire-template:

- The route-target values are matched to determine the pseudowire-template.
- If no matches are found from the previous step, the first (numerically lowest ID) pw-template-binding configured without an import-rt is used.
- If the values used for Encaps Type or Layer 2 MTU do not match, the pseudowire is created but with the operationally down state.
  - To interoperate with existing implementations, if the received MTU value = 0, then MTU negotiation does not take place; the related pseudowire is set up ignoring the MTU.
- If the value of the S flag is not zero, the pseudowire is not created.

The following pseudowire template parameters are supported when applied within a BGP VPWS service; the remainder are ignored:

```
configure service pw-template policy-id [use-provisioned-sdp |
       prefer-provisioned-sdp] [create]
   accounting-policy acct-policy-id
   no accounting-policy
   [no] collect-stats
   [no] controlword
   egress
      filter ipv6 ipv6-filter-id
      filter ip ip-filter-id
      filter mac mac-filter-id
      no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
      qos network-policy-id port-redirect-group queue-group-name instance instance
id
      no qos [network-policy-id]
   [no] force-vlan-vc-forwarding
   hash-label [signal-capability]
   no hash-label
   ingress
      filter ipv6 ipv6-filter-id
      filter ip ip-filter-id
      filter mac mac-filter-id
      no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
      qos network-policy-id fp-redirect-group queue-group-name instance instance-id
      no qos [network-policy-id]
   [no] sdp-exclude
   [no] sdp-include
```

```
vc-type {ether | vlan}
vlan-vc-tag vlan-id
no vlan-vc-tag
```

For more information about this command, refer to the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide*.

The **use-provisioned-sdp** option is permitted when creating the pseudowire template if a pre-provisioned SDP is to be used. Pre-provisioned SDPs must be configured whenever GRE- or RSVP-signaled transport tunnels are used.

When the **prefer-provisioned-sdp** option is specified, if the system finds an existing matching SDP that conforms to any restrictions defined in the pseudowire template (for example, **sdp-include**/**sdp-exclude** *group*), it uses this matching SDP (even if the existing SDP is operationally down); otherwise, it automatically creates an SDP.

The **tools perform** command can be used in the same way as for BGP-AD to apply changes to the pseudowire template using the following format:

```
tools perform service [id service-id] eval-pw-template
 policy-id [allow-service-impact]
```

If a user configures a service using a pseudowire template with the **prefer-provisioned-sdp** option, but without an applicable SDP being provisioned, and the system binds to an automatic SDP, and the user subsequently provisions an appropriate SDP, the system will not automatically switch to the new provisioned SDP. This will only occur if the pseudowire template is re-evaluated using the **tools perform service id** *service-id* **eval-pw-template** command.

2.13.3.5 Use of Endpoint for BGP VPWS

An endpoint is required on a remote PE connecting to two dual-homed PEs to associate the active/standby pseudowires with the Epipe service. An endpoint is automatically created within the Epipe service such that active/standby pseudowires are associated with that endpoint. The creation of the endpoint occurs when **bgp-vpws** is enabled (and deleted when it is disabled) and so will exist in both a single- and dual-homed scenario (this simplifies converting a single-homed service to a dual-homed service). The naming convention used is _tmnx_BgpVpws-x, where x is the service identifier. The automatically created endpoint has the default parameter values, although all are ignored in a BGP-VPWS service with the description field being defined by the system.

The command:

```
tools perform service id <service-id> endpoint <endpoint-name> force-switchover
```

will have no affect on an automatically created VPWS endpoint.

2.14 VLL Service Considerations

This section describes the general 7450 ESS, 7750 SR, and 7950 XRS service features and any special capabilities or considerations as they relate to VLL services.

2.14.1 SDPs

The most basic SDPs must have the following:

- A locally unique SDP identification (ID) number.
- The system IP address of the originating and far-end routers.
- An SDP encapsulation type, either GRE or MPLS.

The most basic Apipe and Fpipe SDP configurations for the 7750 SR must have the following:

- A locally unique SDP identification (ID) number and VC-ID.

2.14.1.1 SDP Statistics for VPLS and VLL Services

The three-node network in Figure 48 shows two MPLS SDPs and one GRE SDP defined between the nodes. These SDPs connect VPLS1 and VPLS2 instances that are defined in the three nodes. With this feature, the operator will have local CLI-based as well as SNMP-based statistics collection for each VC used in the SDPs. This will allow for traffic management of tunnel usage by the different services and with aggregation the total tunnel usage.

Figure 48 **SDP Statistics for VPLS and VLL Services**

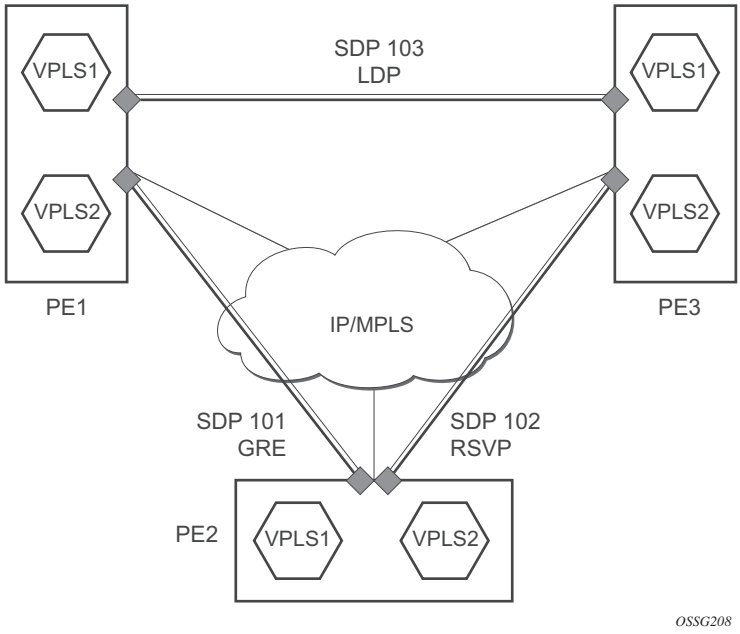

2.14.2 SAP Encapsulations and Pseudowire Types

The Epipe service is designed to carry Ethernet frame payloads, so it can provide connectivity between any two SAPs that pass Ethernet frames. The following SAP encapsulations are supported on the 7450 ESS, 7750 SR, and 7950 XRS Epipe service:

- Ethernet null
- Ethernet dot1q
- QinQ
- SONET/SDH BCP-null for the 7450 ESS and 7750 SR
- SONET/SDH BCP-dot1q for the 7450 ESS and 7750 SR
- ATM VC with RFC 2684 Ethernet bridged encapsulation (see Ethernet Interworking VLL) for the 7750 SR
- FR VC with RFC 2427 Ethernet bridged encapsulation (see Ethernet Interworking VLL) for the 7450 ESS and 7750 SR

While different encapsulation types can be used, encapsulation mismatching can occur if the encapsulation behavior is not understood by connecting devices, which are unable to send and receive the expected traffic. For example, if the encapsulation type on one side of the Epipe is dot1q and the other is null, tagged traffic received on the null SAP will be double-tagged when it is transmitted out of the dot1q SAP.

ATM VLLs can be configured with both endpoints (SAPs) on the same router or with the two endpoints on different 7750 SRs. In the latter case, Pseudowire Emulation Edge-to-Edge (PWE3) signaling is used to establish a pseudowire between the devices, allowing ATM traffic to be tunneled through an MPLS or GRE network:

Two pseudowire encapsulation modes, that is, SDP vc-type, are available:

- PWE3 N-to-1 Cell Mode Encapsulation
- PWE3 AAL5 SDU Mode Encapsulation

The endpoints of Fpipes must be Data-Link Connection Identifiers (DLCIs) on any port that supports Frame Relay. The pseudowire encapsulation, or SDP vc-type, supported is the 1-to-1 Frame Relay encapsulation mode.

2.14.2.1 PWE3 N-to-1 Cell Mode

The endpoints of an N-to-1 mode VLL on a 7750 SR can be:

- ATM VCs — VPI/VCI translation is supported (that is, the VPI/VCI at each endpoint does not need to be the same).
- ATM VPs — VPI translation is supported (that is, the VPI at each endpoint need not be the same, but the original VCI will be maintained).
- ATM VTs (a VP range) — No VPI translation is supported (that is, the VPI/VCI of each cell is maintained across the network).
- ATM ports — No translation is supported (that is, the VPI/VCI of each cell is maintained across the network).

For N-to-1 mode VLLs, cell concatenation is supported. Cells will be packed on ingress to the VLL and unpacked on egress. As cells are being packed, the concatenation process may be terminated by:

- Reaching a maximum number of cells per packet.
- Expiry of a timer.
- (Optionally) change of the CLP bit.
- (Optionally) change of the PTI bits indicating the end of the AAL5 packet.

In N-to-1 mode, OAM cells are transported through the VLL as any other cell. The PTI and CLP bits are untouched, even if VPI/VCI translation is carried out.

2.14.2.2 PWE3 AAL5 SDU Mode

The endpoints of an AAL5 SDU mode VLL on a 7750 SR must be ATM VCs specified by port/vpi/vci. VPI/ VCI translation is supported. The endpoint can also be a FR VC, specified by port/dlci. In this case, FRF.5 FR-ATM network interworking is performed between the ATM VC SAP or the SDP and the FR VC SAP.

In SDU mode, the mandatory PWE3 control word is supported. This allows the ATM VLL to transport OAM cells along with SDU frames, using the “T” bit to distinguish between them, to recover the original SDU length, and to carry CLP, EFCI, and UU information.

2.14.2.3 QoS Policies

When applied to 7450 ESS, 7750 SR, or 7950 XRS Epipe, Apipe, and Fpipe services, service ingress QoS policies only create the unicast queues defined in the policy. The multipoint queues are not created on the service.

With Epipe, Apipe, and Fpipe services, egress QoS policies function as with other services where the class-based queues are created as defined in the policy. Both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in a service. QoS policies on Apipes cannot perform any classification, and on Fpipes Layer 3 (IP), classification is performed.

2.14.2.4 Filter Policies

7450 ESS, 7750 SR, and 7950 XRS Epipe, Fpipe, and Ipipe services can have a single filter policy associated on both ingress and egress. Both MAC and IP filter policies can be used on Epipe services.

Filters cannot be configured on 7750 SR Apipe service SAPs.

2.14.2.5 MAC Resources

Epipe services are point-to-point Layer 2 VPNs capable of carrying any Ethernet payloads. Although an Epipe is a Layer 2 service, the 7450 ESS, 7750 SR, and 7950 XRS Epipe implementation does not perform any MAC learning on the service, so Epipe services do not consume any MAC hardware resources.

2.15 Configuring a VLL Service with CLI

This section provides information to configure Virtual Leased Line (VLL) services using the command line interface.

2.15.1 Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed, and the CLI commands to configure the VLL services, as follows:

1. Associate the service with a customer ID.
2. Define SAP parameters:
   - Optionally - configure ATM parameters on the 7750 SR
   - Optionally - select egress and ingress QoS and/or scheduler policies (configured in the **config>qos** context).
   - Optionally - select accounting policy (configured in the **config>log** context).
3. Define spoke-SDP parameters.
4. Enable the service.

2.15.2 Configuring VLL Components

This section provides VLL configuration examples for the VLL services.

2.15.2.1 Creating an Apipe Service

Use the following CLI syntax to create an Apipe service on a 7750 SR.

CLI Syntax:

```
config>service# apipe service-id [customer customer-id]
  [vpn vpn-id] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-
  cell}][vc-switching]
description description-string
interworking {frf-5}
service-mtu octets
no shutdown
```

The following example shows the command usage to create an Apipe service.

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5 customer 1 create
A:ALA-41config>service>apipe# description “apipe test”
A:ALA-41config>service>apipe# service-mtu 1400
A:ALA-41config>service>apipe# no shutdown
A:ALA-41config>service>apipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5 customer 1 create
A:ALA-42>config>service>apipe# description “apipe test”
A:ALA-42>config>service>apipe# service-mtu 1400
A:ALA-42>config>service>apipe# no shutdown
A:ALA-42>config>service>apipe#
```

The following example shows the Apipe service creation output.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

2.15.2.1.1 Configuring Basic Apipe SAP Parameters

Use the following CLI syntax to configure Apipe SAP parameters.

CLI Syntax:
```
config>service# apipe service-id [customer customer-id]
 [vpn vpn-id] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-
 cell}][vc-switching]
    sap sap-id
         accounting-policy acct-policy-id
         atm
              egress
                   traffic-desc traffic-desc-profile-id
              ingress
                   traffic-desc traffic-desc-profile-id
              oam
                   alarm-cells
                   terminate
         collect-stats
         description description-string
         egress
              qos policy-id
              scheduler-policy scheduler-policy-name
         ingress
              qos policy-id [shared-queuing]
              scheduler-policy scheduler-policy-name
         multi-service-site customer-site-name
         no shutdown
```

The following example shows the command usage to create Apipe SAPs:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 1/1/1:0/32 create
A:ALA-41>config>service>apipe>sap# ingress
A:ALA-41>config>service>apipe>sap>ingress# qos 102
A:ALA-41>config>service>apipe>sap>ingress# exit
A:ALA-41>config>service>apipe>sap# egress
A:ALA-41>config>service>apipe>sap>egress# qos 103
A:ALA-41>config>service>apipe>sap>egress# exit
A:ALA-41>config>service>apipe>sap# no shutdown
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe#
```

PE router 2 (A:ALA-42):

Example:
```
Example:A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# sap 2/2/2:0/32 create
A:ALA-42>config>service>apipe>sap# ingress
A:ALA-42>config>service>apipe>sap>ingress# qos 102
A:ALA-42>config>service>apipe>sap>ingress# exit
A:ALA-42>config>service>apipe>sap# egress
A:ALA-42>config>service>apipe>sap>egress# qos 103
```

```
A:ALA-42>config>service>apipe>sap>egress# exit
A:ALA-42>config>service>apipe>sap# no shutdown
A:ALA-42>config>service>apipe>sap# exit
A:ALA-42>config>service>apipe#
```

The following output shows the Apipe SAP configuration.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

2.15.2.1.2 Configuring an ATM SAP in the N-to-1 Mapping of ATM VPI/VCI to ATM Pseudowire

Users can configure an ATM-cell Apipe service with a new ATM SAP type. The SAP type refers to a preconfigured ATM connection profile name.

```
configure service apipe 100 vc-type atm-cell
      sap <port-id|aps-id>[:cp.<connection-profile-num>]
```

The ATM SAP connection profile is configured with the list of discrete VPI/VCI values, as follows:

```
configure connection-profile 2 {member vpi/vci...(up to 16)}
```

A connection profile can only be applied to a SAP that is part of an Apipe VLL service of vc-type atm-cell. The ATM SAP can be on a regular port or APS port. A connection profile can be applied to any number of ATM SAPs.

Up to a maximum of 16 discrete VPI/VCI values can be configured in a connection profile. After creation of the connection profile, the user can subsequently add, remove, or modify the VPI/VCI entries. This triggers a re-evaluation of all the ATM SAPs that are currently using that profile.

The user must also override the PW type signaled to type '0x0009 N:1 VCC cell' by using the following command:

```
config>service>apipe>signaled-vc-type-override atm-vcc
```

This command is not allowed in an Apipe VLL of vc-type value atm-cell if a configured ATM SAP is not using a connection profile. Conversely, if the signaling override command is enabled, only an ATM SAP with a connection profile assigned will be allowed.

The override command is not allowed on an Apipe VLL service of vc-type value other than atm-cell. It is also not allowed on a VLL service with the **vc-switching** option enabled because signaling of the pseudowire FEC in a Multi-Segment Pseudowire (MS-PW) is controlled by the T-PE nodes. Therefore, for this feature to be used on a MS-PW, configure an Apipe service of vc-type atm-cell at the T-PE nodes with the **signaled-vc-type-override** command enabled, and configure an Apipe VLL service of vc-type atm-vcc at the S-PE node with the **vc-switching** option enabled.

The following are the restrictions of this feature:

- A SAP-to-SAP VLL service is not supported using ATM SAP with a connection profile assigned. The user must configure each VPI/VCI into a separate SAP and create as many Apipe VLL services of type atm-vcc as required.

- An ATM SAP with a connection profile assigned cannot be configured on a port that is part of an MC-APS protection group.
- It is strongly recommended to not apply a VCI-based QoS Filter to the ingress of an ATM SAP with a connection profile. Because the filter matches the VCI value of the first cell of a concatenated packet, the entire packet will be treated the same way based on the action of the entry of the criteria; all cells of the concatenated packet will be mapped to the same FC and profile, based on the VCI value of the first cell.

This feature is supported on the 4-port OC-3/STM-1:OC-12/STM-4 ATM MDA and on the 16-port OC-3/STM-1 ATM MDA, and is supported on IOM3/IMM on the 7750 SR-7 and 7750 SR-12, as well as the 7750-C4 and C12 chassis.

2.15.2.1.3 Configuring Apipe SDP Bindings

Use the following CLI syntax to create a spoke-SDP binding with an Apipe service.

CLI Syntax:
```
config>service# apipe service-id [customer customer-id]
  [vpn vpn-id] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-
  cell}] [vc-switching]
spoke-sdp sdp-id:vc-id
     cell-concatenation
          aal5-frame-aware
          clp-change
          max-cells cell-count
          max-delay delay-time
     egress
          vc-label egress-vc-label
     ingress
          vc-label ingress-vc-label
     no shutdown
```

The following example displays the command usage to create Apipe spoke-SDPs:

PE router 1 (A:ALA-41):

Example:
```
Example:A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# spoke-sdp 1:5 create
A:ALA-41>config>service>apipe>spoke-sdp# no shutdown
A:ALA-41>config>service>apipe>spoke-sdp# exit
```

PE router 2 (A:ALA-42):

Example:
```
Example:A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# spoke-sdp 1:5 create
```

```
A:ALA-42>config>service>apipe>spoke-sdp# no shutdown
A:ALA-42>config>service>apipe>spoke-sdp# exit
```

The following output shows the Apipe spoke-SDP configurations.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:5 create
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:5 create
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

2.15.2.2 Creating a Cpipe Service

2.15.2.2.1 Basic Configuration

Use the following CLI syntax to create a Cpipe service on a 7750 SR. A route distinguisher must be defined in order for Cpipe to be operationally active.

CLI Syntax:
```
config>service# cpipe service-id [customer customer-id]
  [vpn vpn-id] [vc-type {satop-e1 | satop-t1 | cesopsn |
  cesopsn-cas}] [vc-switching] [create]
```

For the 7450 ESS platforms, the vc-switching option must be configured for Cpipe functionality, as follows:

```
cpipe 1 customer 1 vc-switching vc-type cesopsn create
            service-name "XYZ Cpipe 1"
            spoke-sdp 20:1 create
                description "Description for Sdp Bind 20 for Svc ID 1"
                ingress
                    vc-label 10002
                exit
                egress
                    vc-label 10001
                exit
            exit
            spoke-sdp 50:1 create
                description "Description for Sdp Bind 50 for Svc ID 1"
            exit
            no shutdown
        exit
```

The following displays a Cpipe service configuration example:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        cpipe 210 customer 1 vc-type cesopsn create
            service-mtu 1400
            sap 1/5/1.1.3.1 create
            exit
            spoke-sdp 1:210 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

2.15.2.2.2 Configuration Requirements

Before a Cpipe service can be provisioned, the following tasks must be completed:

- Configuring a DS1 Port
- Configuring a Channel Group

Configuring a DS1 Port

The following example shows a DS1 port configured for CES:

```
A:sim216#  show port 1/5/1.1.3.1
===============================================================================
TDM DS1 Interface
===============================================================================
Description        : DS1
Interface          : 1/5/1.1.3,1
Type               : ds1                      Framing             : esf
Admin Status       : up                       Oper Status         : up
Physical Link      : yes                      Clock Source        : loop-timed
Signal Mode        : none
Last State Change  : 10/31/2006 14:23:12      Channel IfIndex     : 580943939
Loopback           : none                     Invert Data         : false
Remote Loop respond: false                    In Remote Loop      : false
Load-balance-algo  : default                  Egr. Sched. Pol     : n/a
BERT Duration      : N/A                      BERT Pattern        : none
BERT Synched       : 00h00m00s                Err Insertion Rate  : 0
BERT Errors        : 0                        BERT Status         : idle
BERT Total Bits    : 0
Cfg Alarm          : ais los
Alarm Status       :
===============================================================================
A:sim216#
```

Configuring a Channel Group

The following example shows a DS1 channel group configured for CES:

```
A:sim216#  show port 1/5/1.1.3.1
===============================================================================
TDM DS0 Chan Group
===============================================================================
Description        : DS0GRP
Interface          : 1/5/1.1.3.1
TimeSlots          : 1-12
Speed              : 64                       CRC                 : 16
Admin Status       : up                       Oper Status         : up
Last State Change  : 10/31/2006 14:23:12      Chan-Grp IfIndex    : 580943940
Configured mode    : access                   Encap Type          : cem
Admin MTU          : 4112                     Oper MTU            : 4112
Physical Link      : Yes                      Bundle Number       : none
```

```
Idle Cycle Flags   : flags                   Load-balance-algo    : default
Egr. Sched. Pol    : n/a
===============================================================================
A:sim216#
```

2.15.2.2.3 Configuring Cpipe SAPs and Spoke-SDPs

The following examples show Cpipe SAP and spoke-SDP configurations:

```
*A:ALA-49>config>service# info
#------------------------------------------------
echo "Service Configuration"
#------------------------------------------------
...
        cpipe 100 customer 1 vc-type cesopsn create
            service-mtu 1400
            sap 1/5/1.1.1.1 create
            exit
            spoke-sdp 1:100 create
            exit
            no shutdown
        exit
        cpipe 200 customer 1 vc-type cesopsn-cas create
            sap 1/5/1.2.1.1 create
            exit
            sap 1/5/1.2.2.1 create
            exit
            no shutdown
        exit
        cpipe 210 customer 1 vc-type cesopsn-cas create
            service-mtu 1400
            sap 1/5/1.1.3.1 create
            exit
            spoke-sdp 1:210 create
            exit
            no shutdown
        exit
        cpipe 300 customer 1 vc-type cesopsn create
            sap 1/5/1.3.4.1 create
            exit
            sap 1/5/1.3.6.1 create
            exit
            no shutdown
        exit
        cpipe 400 customer 1 vc-type satop-e1 create
            sap 1/5/1.2.3.1 create
            exit
            spoke-sdp 1:400 create
            exit
            no shutdown
        exit
...
#------------------------------------------------
*A:ALA-49>config>service#
```

```
A:sim213>config>service>cpipe# info
----------------------------------------------
            description "cpipe-100"
            sap 1/5/1.1.1.1 create
                cem
                    packet jitter-buffer 16 payload-size 384
                    report-alarm rpktloss
                    no report-alarm stray
                    rtp-header
                exit
            exit
            spoke-sdp 1:100 create
            exit
            no shutdown
----------------------------------------------
A:sim213>config>service>cpipe#
```

2.15.2.3 Creating an Epipe Service

Use the following CLI syntax to create an Epipe service.

CLI Syntax:
```
config>service# epipe service-id [customer customer-id]
  [vpn vpn-id] [vc-switching]
description description-string
no shutdown
```

The following example shows an Epipe configuration:

```
A:ALA-1>config>service# info
----------------------------------------
...
        epipe 500 customer 5 vpn 500 create
            description "Local epipe service"
            no shutdown
        exit
----------------------------------------
A:ALA-1>config>service#
```

2.15.2.3.1 Configuring Epipe SAP Parameters

A default QoS policy is applied to each ingress and egress SAP. Additional QoS policies can be configured in the **config>qos** context. Filter policies are configured in the config>filter context and explicitly applied to a SAP. There are no default filter policies.

Use the following CLI syntax to create:

- Local Epipe SAPs

- Distributed Epipe SAPs

The following example shows a configuration for the 7950 XRS.

CLI Syntax:

```
config>service# epipe service-id [customer customer-id]
sap sap-id [endpoint endpoint-name]
sap sap-id [no-endpoint]
     accounting-policy policy-id
     collect-stats
     description description-string
     no shutdown
     egress
          filter {ip ip-filter-name | mac mac-filter-
            name}
          qos sap-egress-policy-id
          scheduler-policy scheduler-policy-name
     ingress
          filter {ip ip-filter-name | mac mac-filter-
            name}
          match-qinq-dot1p {top | bottom}
          qos policy-id
          scheduler-policy scheduler-policy-name
```

The following example shows a configuration for the 7450 ESS and 7750 SR.

CLI Syntax:

```
config>service# epipe service-id [customer customer-id]
sap sap-id [endpoint endpoint-name]
sap sap-id [no-endpoint]
     accounting-policy policy-id
     collect-stats
     description description-string
     no shutdown
     egress
          filter {ip ip-filter-name | mac mac-filter-
            name}
          qos sap-egress-policy-id
          scheduler-policy scheduler-policy-name
     ingress
          filter {ip ip-filter-name | mac mac-filter-
            name}
          match-qinq-dot1p {top|bottom}
          qos policy-id [shared-queuing]
          scheduler-policy scheduler-policy-name
```

Local Epipe SAPs

To configure a basic local Epipe service, enter the **sap** *sap-id* command twice with different port IDs in the same service configuration.

By default, QoS policy ID 1 is applied to ingress and egress service SAPs. Existing filter policies or other existing QoS policies can be associated with service SAPs on ingress and egress ports. Table 11 shows supported SAP types.

Table 11 **Supported SAP Types**

Uplink Type	Svc SAP Type	Cust. VID	Access SAPs	Network SAPs
L2	Null-star	N/A	Null, dot1q *	Q.*
L2	Dot1q	N/A	Dot1q	Q.*
L2	Dot1q-preserve	X	Dot1q (encap = X)	Q1.Q2 (where Q2 = X)

An existing scheduler policy can be applied to ingress and egress SAPs to be used by the SAP queues and, at egress only, policers. The schedulers comprising the policy are created when the scheduler policy is applied to the SAP. If any policers or orphaned queues (with a non-existent local scheduler defined) exist on a SAP and the policy application creates the required scheduler, the status on the queue becomes non-orphaned at this time.

Ingress and egress SAP parameters can be applied to local and distributed Epipe service SAPs.

The following example shows the SAP configurations for local Epipe service 500 on SAP 1/1/2 and SAP 1/1/3 on ALA-1:

```
A:ALA-1>config>service# epipe 500 customer 5 create
    config>service>epipe$ description "Local epipe service"
    config>service>epipe# sap 1/1/2:0 create
    config>service>epipe>sap? ingress
    config>service>epipe>sap>ingress# qos 20
    config>service>epipe>sap>ingress# filter ip 1
    config>service>epipe>sap>ingress# exit
    config>service>epipe>sap# egress
    config>service>epipe>sap>egress# qos 20
    config>service>epipe>sap>egress# scheduler-policy test1
    config>service>epipe>sap>egress# exit
    config>service>epipe>sap# no shutdown
    config>service>epipe>sap# exit

    config>service>epipe# sap 1/1/3:0 create
    config>service>epipe>sap# ingress
    config>service>epipe>sap>ingress# qos 555
```

```
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# egress
config>service>epipe>sap>egress# qos 627
config>service>epipe>sap>egress# scheduler-policy alpha
config>service>epipe>sap>egress# exit
config>service>epipe>sap# no shutdown
config>service>epipe>sap# exit
```

The following example shows the local Epipe configuration:

```
A:ALA-1>config>service# info
----------------------------------------------
...
        epipe 500 customer 5 vpn 500 create
            description "Local epipe service"
            sap 1/1/2:0 create
                ingress
                    qos 20
                    filter ip 1
                exit
                egress
                    scheduler-policy "test1"
                    qos 20
                exit
            exit
            sap 1/1/3:0 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
            no shutdown
        exit
----------------------------------------------
A:ALA-1>config>service#
```

2.15.2.3.2 Distributed Epipe SAPs

To configure a distributed Epipe service, you must configure service entities on the originating and far-end nodes. You should use the same service ID on both ends (for example, Epipe 5500 on ALA-1 and Epipe 5500 on ALA-2). The **spoke-sdp** *sdp-id:vc-id* must match on both sides. A distributed Epipe consists of two SAPs on different nodes.

By default, QoS policy ID 1 is applied to ingress and egress service SAPs. Existing filter policies or other existing QoS policies can be associated with service SAPs on ingress and egress.

An existing scheduler policy can be applied to ingress and egress SAPs to be used by the SAP queues and, at egress only, policers. The schedulers comprising the policy are created when the scheduler policy is applied to the SAP. If any policers or orphaned queues (with a non-existent local scheduler defined) exist on a SAP and the policy application creates the required scheduler, the status on the queue becomes non-orphaned at this time.

Ingress and egress SAP parameters can be applied to local and distributed Epipe service SAPs.

For SDP configuration information, refer to the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide*. For SDP binding information, see Configuring SDP Bindings.

The following example shows a configuration of a distributed service between ALA-1 and ALA-2:

```
A:ALA-1>epipe 5500 customer 5 create
    config>service>epipe$ description "Distributed epipe service to east coast"
    config>service>epipe# sap 221/1/3:21 create
    config>service>epipe>sap# ingress
    config>service>epipe>sap>ingress# qos 555
    config>service>epipe>sap>ingress# filter ip 1
    config>service>epipe>sap>ingress# exit
    config>service>epipe>sap# egress
    config>service>epipe>sap>egress# qos 627
    config>service>epipe>sap>egress# scheduler-policy alpha
    config>service>epipe>sap>egress# exit
    config>service>epipe>sap# no shutdown
    config>service>epipe>sap# exit
    config>service>epipe#

A:ALA-2>config>service# epipe 5500 customer 5 create
    config>service>epipe$ description "Distributed epipe service to west coast"
    config>service>epipe# sap 441/1/4:550 create
    config>service>epipe>sap# ingress
    config>service>epipe>sap>ingress# qos 654
    config>service>epipe>sap>ingress# filter ip 1020
    config>service>epipe>sap>ingress# exit
    config>service>epipe>sap# egress
    config>service>epipe>sap>egress# qos 432
    config>service>epipe>sap>egress# filter ip 6
    config>service>epipe>sap>egress# scheduler-policy test1
    config>service>epipe>sap>egress# exit
    config>service>epipe>sap# no shutdown
    config>service>epipe#
```

The following example shows the SAP configurations for ALA-1 and ALA-2:

```
A:ALA-1>config>service# info
----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
```

```
            sap 221/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
        exit
...
-----------------------------------------------
A:ALA-1>config>service#


A:ALA-2>config>service# info
-----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
            sap 441/1/4:550 create
                ingress
                    qos 654
                    filter ip 1020
                exit
                egress
                    scheduler-policy "test1"
                    qos 432
                    filter ip 6
                exit
            exit
        exit
...
-----------------------------------------------
A:ALA-2>config>service#
```

PBB Epipe Configuration

The following example shows the PBB Epipe configuration:

```
*A:Wales-1>config>service>epipe# info
----------------------------------------------------------------------
...
    description "Default epipe description for service id 20000"
    pbb-tunnel 200 backbone-dest-mac 00:03:fa:15:d3:a8 isid 20000
    sap 1/1/2:1.1 create
        description "Default sap description for service id 20000"
        ingress
        filter mac 1
        exit
    exit
    no shutdown
----------------------------------------------------------------------
*A:Wales-1>config>service>epipe#
```

CLI Syntax: `configure service vpls 200 customer 1 b-vpls create`

```
*A:Wales-1>config>service>vpls# info
----------------------------------------------------------------------
...
    service-mtu 2000
    fdb-table-size 131071
    stp
    no shutdown
    exit
    sap 1/1/8 create
    exit
    sap 1/2/3:200 create
    exit
    mesh-sdp 1:200 create
    exit
    mesh-sdp 100:200 create
    exit
    mesh-sdp 150:200 create
    exit
    mesh-sdp 500:200 create
    exit
    no shutdown
----------------------------------------------------------------------
*A:Wales-1>config>service>vpls#
```

Configuring Ingress and Egress SAP Parameters

By default, QoS policy ID 1 is applied to ingress and egress service SAPs. Existing filter policies or other existing QoS policies can be associated with service SAPs on ingress and egress ports.

An existing scheduler policy can be applied to ingress and egress SAPs to be used by the SAP queues and, at egress only, policers. The schedulers comprising the policy are created when the scheduler policy is applied to the SAP. If any policers or orphaned queues (with a non-existent local scheduler defined) exist on a SAP and the policy application creates the required scheduler, the status on the queue becomes non-orphaned at this time.

Ingress and egress SAP parameters can be applied to local and distributed Epipe service SAPs.

The following example shows the SAP ingress and egress parameters:

```
ALA-1>config>service# epipe 5500
    config>service>epipe# sap 2/1/3:21
    config>service>epipe>sap# ingress
    config>service>epipe>sap>ingress# qos 555
    config>service>epipe>sap>ingress# filter ip 1
    config>service>epipe>sap>ingress# exit
    config>service>epipe>sap# egress
    config>service>epipe>sap>egress# qos 627
```

```
    config>service>epipe>sap>egress# scheduler-policy alpha
    config>service>epipe>sap>egress# exit
    config>service>epipe>sap#
```

The following example shows the Epipe SAP ingress and egress configuration:

```
A:ALA-1>config>service#
------------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
            sap 2/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
            spoke-sdp 2:123 create
                ingress
                    vc-label 6600
                exit
                egress
                    vc-label 5500
                exit
            exit
            no shutdown
        exit
------------------------------------------------
A:ALA-1>config>service#
```

2.15.2.3.3 Configuring SDP Bindings

Figure 49 shows an example of a distributed Epipe service configuration between two routers, identifying the service and customer IDs, and the uni-directional SDPs required to communicate to the far-end routers.

A spoke-SDP is treated like the equivalent of a traditional bridge “port”, where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

Figure 49 **SDPs — Uni-Directional Tunnels**

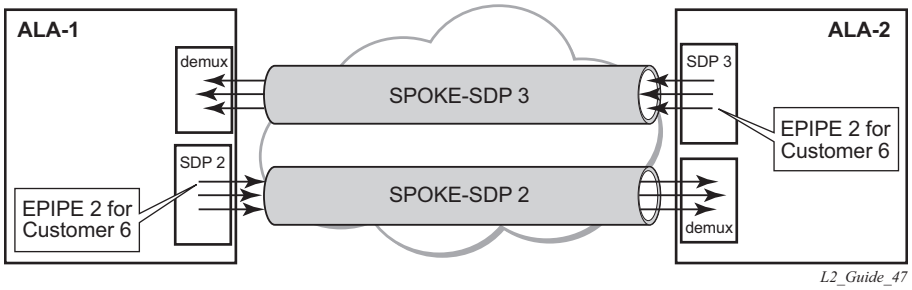

Use the following CLI syntax to create a spoke-SDP binding with an Epipe service.

CLI Syntax:

```
config>service# epipe service-id [customer customer-id]
spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}]
     vlan-vc-tag 0..4094
     egress
          filter {ip ip-filter-id}
          vc-label egress-vc-label
     ingress
          filter {ip ip-filter-id}
          vc-label ingress-vc-label
     no shutdown
```

The following example shows the command usage to bind an Epipe service between ALA-1 and ALA-2. This example assumes the SAPs have already been configured (see Distributed Epipe SAPs).

```
A:ALA-1>config>service# epipe 5500
    config>service>epipe# spoke-sdp 2:123
    config>service>epipe>spoke-sdp# egress
    config>service>epipe>spoke-sdp>egress# vc-label 5500
    config>service>epipe>spoke-sdp>egress# exit
    config>service>epipe>spoke-sdp# ingress
    config>service>epipe>spoke-sdp>ingress# vc-label 6600
    config>service>epipe>spoke-sdp>ingress# exit
    config>service>epipe>spoke-sdp# no shutdown

    ALA-2>config>service# epipe 5500
    config>service>epipe# spoke-sdp 2:456
    config>service>epipe>spoke-sdp# egress
    config>service>epipe>spoke-sdp>egress# vc-label 6600
    config>service>epipe>spoke-sdp>egress# exit
    config>service>epipe>spoke-sdp# ingress
    config>service>epipe>spoke-sdp>ingress# vc-label 5500
    config>service>epipe>spoke-sdp>ingress# exit
    config>service>epipe>spoke-sdp# no shutdown
```

The following example shows the SDP binding for the Epipe service between ALA-1 and ALA-2:

```
A:ALA-1>config>service# info
----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
            sap 2/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
            spoke-sdp 2:123 create
                ingress
                    vc-label 6600
                exit
                egress
                    vc-label 5500
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-1>config>service#


A:ALA-2>config>service# info
----------------------------------------------
...
exit
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
            sap 441/1/4:550 create
                ingress
                    qos 654
                    filter ip 1020
                exit
                egress
                    scheduler-policy "test1"
                    qos 432
                    filter ip 6
                exit
            exit
            spoke-sdp 2:456 create
                ingress
                    vc-label 5500
                exit
                egress
                    vc-label 6600
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
```

```
A:ALA-2>config>service#
```

2.15.2.4 Creating an Fpipe Service

Use the following CLI syntax to create an Fpipe service on a 7750 SR.

CLI Syntax:
```
config>service# fpipe service-id [customer customer-id]
  [vpn vpn-id] [vc-type {fr-dlci}] [vc-switching]
description description-string
service-mtu octets
no shutdown
```

The following example shows the command usage to create an Fpipe service:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1 customer 1 create
A:ALA-41config>service>fpipe# description "fpipe test"
A:ALA-41config>service>fpipe# service-mtu 1400
A:ALA-41config>service>fpipe# no shutdown
A:ALA-41config>service>fpipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1 customer 1 create
A:ALA-42>config>service>fpipe# description "fpipe test"
A:ALA-42>config>service>fpipe# service-mtu 1400
A:ALA-42>config>service>fpipe# no shutdown
A:ALA-42>config>service>fpipe#
```

The following example shows the Fpipe service creation output:

PE router 1 (A:ALA-41):

A:ALA-41>config>service# info

```
----------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            no shutdown
        exit
...
----------------------------------------
A:ALA-41>config>service#
```

PE router 2 (A:ALA-42):

A:ALA-42>config>service# info

```
----------------------------------------

...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            no shutdown
        exit
...
----------------------------------------------
A:ALA-42>config>service#
```

2.15.2.4.1 Configuring Fpipe SAP Parameters

Use the following CLI syntax to configure Fpipe SAP parameters.

CLI Syntax:

```
config>service# fpipe service-id [customer customer-id]
  [vpn vpn-id] [vc-type {fr-dlci}] [vc-switching]
      sap sap-id
           accounting-policy acct-policy-id
           collect-stats
           description description-string
           egress
                filter [ip ip-filter-id]
                qos policy-id
                scheduler-policy scheduler-policy-name
           ingress
                filter [ip ip-filter-id]
                qos policy-id [shared-queuing]
                scheduler-policy scheduler-policy-name
           multi-service-site customer-site-name
           no shutdown
```

The following example shows the command usage to create an Fpipe SAP:

PE router 1 (A:ALA-41):

Example:

```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 1/2/1:16 create
A:ALA-41>config>service>fpipe>sap# ingress
A:ALA-41>config>service>fpipe>sap>ingress# qos 101
A:ALA-41>config>service>fpipe>sap>ingress# exit
A:ALA-41>config>service>fpipe>sap# egress
A:ALA-41>config>service>fpipe>sap>egress# qos 1020
```

```
A:ALA-41>config>service>fpipe>sap>egress# exit
A:ALA-41>config>service>fpipe>sap# no shutdown
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe#
```

PE router 2 (A:ALA-42):

Example:

```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# sap 2/1/1.1:16 create
A:ALA-42>config>service>fpipe>sap# ingress
A:ALA-42>config>service>fpipe>sap>ingress# qos 101
A:ALA-42>config>service>fpipe>sap>ingress# exit
A:ALA-42>config>service>fpipe>sap# egress
A:ALA-42>config>service>fpipe>sap>egress# qos 1020
A:ALA-42>config>service>fpipe>sap>egress# exit
A:ALA-42>config>service>fpipe>sap# no shutdown
A:ALA-42>config>service>fpipe>sap# exit
A:ALA-42>config>service>fpipe#
```

The following example shows the Fpipe SAP configuration:

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               ingress
```

```
                    qos 101
                exit
                egress
                    qos 1020
                exit
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

2.15.2.4.2 Configuring Fpipe SDP Bindings

Use the following CLI syntax to create a spoke-SDP binding with an Fpipe service.

CLI Syntax:
```
config>service# fpipe service-id [customer customer-id]
  [vpn vpn-id] [vc-type {fr-dlci}] [vc-switching]
spoke-sdp sdp-id:vc-id
     egress
          filter ip ip-filter-id
          vc-label egress-vc-label
     ingress
          filter ip ip-filter-id
          vc-label ingress-vc-label
     no shutdown
```

The following example shows the command usage to create an Fpipe spoke-SDP:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# spoke-sdp 1:1 create
A:ALA-41>config>service>spoke-sdp# no shutdown
A:ALA-41>config>service>spoke-sdp# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# spoke-sdp 1:1 create
A:ALA-42>config>service>spoke-sdp# no shutdown
A:ALA-42>config>service>spoke-sdp# exit
```

The following example shows the Fpipe spoke-SDP configuration:

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
```

```
-----------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
            exit
            no shutdown
        exit
...
-----------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
-----------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
            exit
            no shutdown
        exit
...
-----------------------------------
A:ALA-42>config>service#
```

2.15.2.5 Creating an Ipipe Service

Use the following CLI syntax to create an Ipipe service on a 7450 ESS or 7750 SR.

CLI Syntax:

```
config>service# ipipe service-id [customer customer-id]
  [vpn vpn-id][vc-switching]
description description-string
no shutdown
```

The following example shows an Ipipe configuration:

```
A:ALA-1>config>service# info
----------------------------------------
...
        ipipe 202 customer 1 create
            description "eth_ipipe"
            no shutdown
        exit
----------------------------------------
A:ALA-1>config>service#
```

2.15.2.5.1 Configuring Ipipe SAP Parameters

The following example shows an Ipipe SAP configuration:

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 202 customer 1 create
            sap 1/1/2:444 create
                description "eth_ipipe"
                ce-address 31.31.31.1
            exit
            sap 1/3/2:445 create
                description "eth_ipipe"
                ce-address 31.31.31.2
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

The following example shows a Frame Relay to Ethernet local Ipipe configuration.

Example:

```
config>service# ipipe 204 customer 1 create
config>service>ipipe$ sap 1/1/2:446 create
config>service>ipipe>sap$ description "eth_fr_ipipe"
config>service>ipipe>sap$ ce-address 32.32.32.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# sap 2/2/2:16 create
config>service>ipipe>sap$ ce-address 32.32.32.2
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# no shutdown
config>service>ipipe# exit
config>service#
```

The following example shows the output:

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 204 customer 1 create
            sap 1/1/2:446 create
                description "eth_fr_ipipe"
                ce-address 32.32.32.1
            exit
            sap 2/2/2:16 create
                ce-address 32.32.32.2
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

The following example shows a PPP to Ethernet local Ipipe configuration.

Example:

```
config>service# ipipe 206 customer 1 create
config>service>ipipe$ sap 1/1/2:447 create
config>service>ipipe>sap$ description "eth_ppp_ipipe"
config>service>ipipe>sap$ ce-address 33.33.33.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# sap 2/2/2 create
config>service>ipipe>sap$ description "ppp_eth_ipipe"
config>service>ipipe>sap$ ce-address 33.33.33.2
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# no shutdown
config>service>ipipe# exit
config>service#
```

The following example shows the output:

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 206 customer 1 create
            sap 1/1/2:447 create
                description "eth_ppp_ipipe"
                ce-address 33.33.33.1
            exit
            sap 2/2/2 create
                description "ppp_eth_ipipe"
                ce-address 33.33.33.2
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

2.15.2.5.2 Configuring Ipipe SDP Bindings

The following example shows an Ipipe SDP configuration:

```
A:ALA-48>config>service# info
----------------------------------------------
...
        sdp 16 mpls create
            far-end 4.4.4.4
            ldp
            path-mtu 1600
            keep-alive
                shutdown
            exit
            no shutdown
        exit
...
        ipipe 207 customer 1 create
            shutdown
            sap 1/1/2:449 create
                description "Remote_Ipipe"
                ce-address 34.34.34.1
            exit
            spoke-sdp 16:516 create
                ce-address 31.31.31.2
            exit
        exit
...

----------------------------------------------
A:ALA-48>config>service#
```

2.15.3 Using Spoke-SDP Control Words

The control word command provides the option to add a control word as part of the packet encapsulation for PW types for which the control word is optional. These are Ethernet pseudowire (Epipe), ATM N:1 cell mode pseudowires (Apipe vc-types atm-vcc and atm-vpc), and VT pseudowire (Apipe vc-type atm-cell). The control word might be needed because when ECMP is enabled on the network, packets of a specific pseudowire may be spread over multiple ECMP paths if the hashing router mistakes the PW packet payload for an IPv4 or IPv6 packet. This occurs when the first nibble following the service label corresponds to a value of 4 or 6.

The control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported and, therefore, the service will only come up if the same C-bit value is signaled in both directions. If a spoke-SDP is configured to use the control word, but the node receives a label mapping message with a C-bit clear, the node releases the label with an “Illegal C-bit” status code per Section 6.1 of RFC 4447. As soon as the user enables control of the remote peer, the remote peer withdraws its original label and sends a label mapping with the C-bit set to 1 and the VLL service is up in both nodes.

When the control word is enabled, VCCV packets also include the VCCV control word. In that case, the VCCV CC type 1 (OAM CW) is signaled in the VCCV parameter in the FEC. If the control word is disabled on the spoke-SDP, the Router Alert label is used. In that case, VCCV CC type 2 is signaled. For a multi-segment pseudowire (MS-PW), the CC type 1 is the only type supported; therefore, the control word must be enabled on the spoke-SDP to be able to use VCCV-ping and VCCV-trace.

The following example shows a spoke-SDP control word configuration:

```
-Dut-B>config>service>epipe# info
----------------------------------------------
    description "Default epipe description for service id 2100"
    sap 1/2/7:4 create
        description "Default sap description for service id 2100"
    exit
    spoke-sdp 1:2001 create
        control-word
    exit
    no shutdown
----------------------------------------------
*A:ALA-Dut-B>config>service>epipe#
To disable the control word on spoke-sdp 1:2001:
*A:ALA-Dut-B>config>service>epipe# info
----------------------------------------------
    description "Default epipe description for service id 2100"
    sap 1/2/7:4 create
        description "Default sap description for service id 2100"
    exit
    spoke-sdp 1:2001 create
    exit
    no shutdown
----------------------------------------------
*A:ALA-Dut-B>config>service>epipe#
```

2.15.4 Same-Fate Epipe VLANs Access Protection

The following example shows a G.8031 Ethernet tunnel for Epipe protection configuration for 7450 ESS or 7750 SR using same-fate SAPs for each Epipe access (two Ethernet member ports 1/1/1 and 2/1/1/1 are used):

```
*A:7750_ALU>config>eth-tunnel 1
-----------------------------------------------
        description "Protection is APS"
        protection-type 8031_1to1
        ethernet
            mac 00:11:11:11:11:12
            encap-type dot1q
        exit
        ccm-hold-time down 5 up 10 // 50 ms down, 1 second up
        path 1
            member 1/1/1
            control-tag 5 // primary control vlan 5
            precedence primary
            eth-cfm
                mep 2 domain 1 association 1
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        path 2
            member 2/1/1
            control-tag 105 //secondary control vlan 105
            eth-cfm
                mep 2 domain 1 association 2
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        no shutdown
-------------------------------------------------
# Configure Ethernet tunnel SAPs
-------------------------------------------------
*A:7750_ALU>config>service epipe 10 customer 5 create
        sap eth-tunnel-1 create // Uses control tags from the Ethernet tunnel port
            description "g8031-protected access ctl/data SAP for eth-tunnel 1"

        exit
        no shutdown
-----------------------------------------------
*A:7750_ALU>config>service epipe 11 customer 5 create
    sap eth-tunnel-1:1 create
        description "g8031-protected access same-fate SAP for eth-tunnel 1"

        // must specify tags for each corresponding path in Ethernet tunnel port
        eth-tunnel path 1 tag 6
        eth-tunnel path 2 tag 106
    exit
    …
-----------------------------------------------
*A:7750_ALU>config>service epipe 10 customer 5 create
    sap eth-tunnel-1:3 create
        description "g8031-protected access same-fate SAP for eth-tunnel 1"
        // must specify tags for each path for same-fate SAPs
```

```
        eth-tunnel path 1 tag 10
        eth-tunnel path 2 tag 110
    exit
    …
-----------------------------------------------
```

2.15.5 Pseudowire Configuration Notes

The **vc-switching** parameter must be specified when the VLL service is created. When the **vc-switching** parameter is specified, you are configuring an S-PE. This is a pseudowire switching point (switching from one pseudowire to another). Therefore, you cannot add a SAP to the configuration.

The following example shows the configuration when a SAP is added to a pseudowire. The CLI generates an error response if you attempt to create a SAP. VC switching is only needed on the pseudowire at the S-PE.

```
*A:ALA-701>config>service# epipe 28 customer 1 create vc-switching
*A:ALA-701>config>service>epipe$ sap 1/1/3 create
MINOR: SVCMGR #1311 SAP is not allowed under PW switching service
*A:ALA-701>config>service>epipe$
```

Use the following CLI syntax to create pseudowire switching VLL services. These are examples only. Different routers support different pseudowire switching VLL services.

CLI Syntax:
```
config>service# apipe service-id [customer customer-id]
  [vpn vpn-id] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-
  cell}] [vc-switching]
description description-string
spoke-sdp sdp-id:vc-id
```

CLI Syntax:
```
config>service# epipe service-id [customer customer-
  id][vpn vpn-id][vc-switching]
description description-string
spoke-sdp sdp-id:vc-id
```

CLI Syntax:
```
config>service# fpipe service-id [customer customer-
  id][vpn vpn-id] [vc-type {fr-dlci}] [vc-switching]
description description-string
spoke-sdp sdp-id:vc-id
```

CLI Syntax:
```
config>service# ipipe service-id [customer customer-
  id][vpn vpn-id] [vc-switching]
description description-string
spoke-sdp sdp-id:vc-id
```

The following example shows the command usage to configure VLL pseudowire switching services:

Example:

```
config>service# apipe 1 customer 1 vpn 1 vc-switching
 create
config>service>apipe$ description “Default apipe
 description for service id 100”
config>service>apipe# spoke-sdp 3:1 create
config>service>apipe>spoke-sdp# exit
config>service>apipe# spoke-sdp 6:200 create
config>service>apipe>spoke-sdp# exit
config>service>apipe# no shutdown
```

The following example shows configurations for each service:

```
*A:ALA-48>config>service# info
----------------------------------------------
...
        apipe 100 customer 1 vpn 1 vc-switching create
            description "Default apipe description for service id 100"
            spoke-sdp 3:1 create
            exit
            spoke-sdp 6:200 create
            exit
            no shutdown
        exit
...
        epipe 107 customer 1 vpn 107 vc-switching create
            description "Default epipe description for service id 107"
            spoke-sdp 3:8 create
            exit
            spoke-sdp 6:207 create
            exit
            no shutdown
        exit
...
        ipipe 108 customer 1 vpn 108 vc-switching create
            description "Default ipipe description for service id 108"
            spoke-sdp 3:9 create
            exit
            spoke-sdp 6:208 create
            exit
            no shutdown
        exit
...
        fpipe 109 customer 1 vpn 109 vc-switching create
            description "Default fpipe description for service id 109"
            spoke-sdp 3:5 create
            exit
            spoke-sdp 6:209 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-48>config>service#
```

2.15.6 Configuring Two VLL Paths Terminating on T-PE2

Figure 50 **VLL Resilience with Pseudowire Redundancy and Switching**

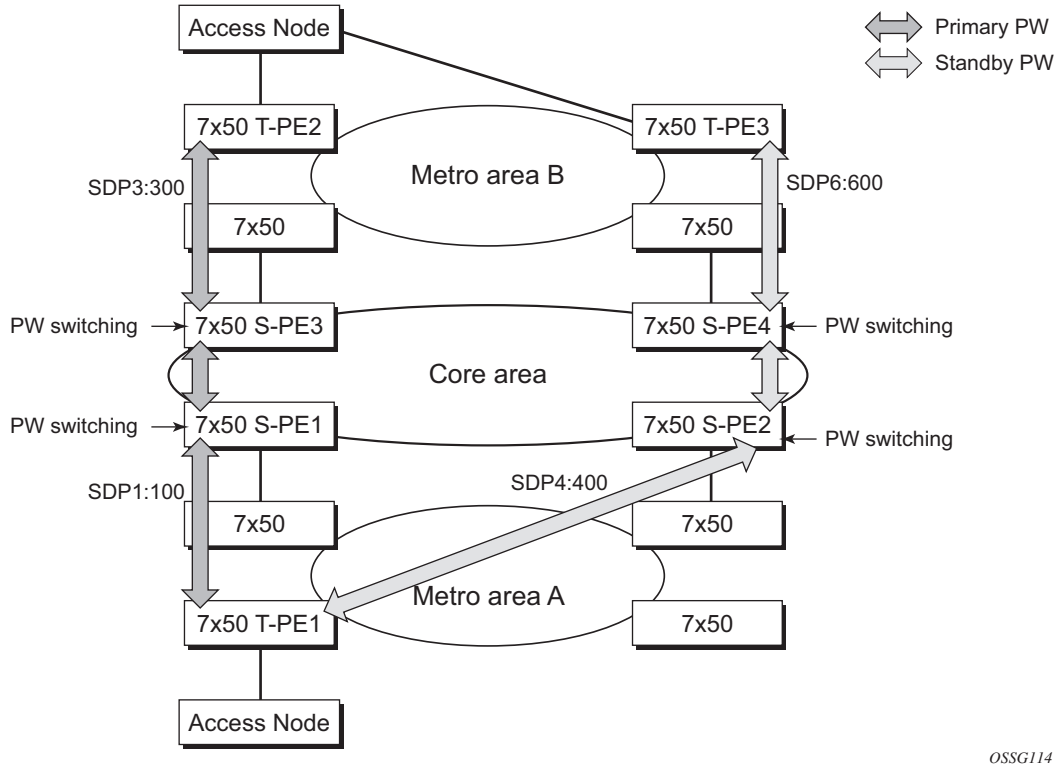

T-PE1

The following shows an example of the T-PE1 configuration:

```
*A:ALA-T-PE1>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            spoke-sdp 1:100 endpoint "y" create
                precedence primary
                revert-time 0
            exit
            spoke-sdp 4:400 endpoint "y" create
                precedence 0
            exit
            no shutdown
----------------------------------------------
*A:ALA-T-PE1>config>service>epipe#
```

The following shows an example of the T-PE2 configuration for 7950 XRS.

T-PE2

```
*A:ALA-T-PE2>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap endpoint "x" create
            exit
            spoke-sdp 3:300 endpoint "y" create
                precedence primary
                revert-time 0
            exit
            spoke-sdp 6:600 endpoint "y" create
                precedence 0
            exit
            no shutdown
----------------------------------------------
*A:ALA-T-PE2>config>service>epipe#
```

The following shows and example of the T-PE2 configuration for 7450 ESS and 7750 SR.

T-PE2

```
*A:ALA-T-PE2>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap 2/2/2:200 endpoint "x" create
            exit
            spoke-sdp 3:300 endpoint "y" create
                precedence primary
                revert-time 0
            exit
            spoke-sdp 6:600 endpoint "y" create
                precedence 0
            exit
            no shutdown
----------------------------------------------
*A:ALA-T-PE2>config>service>epipe#
```

S-PE1: Specifying the **vc-switching** parameter enables a VC cross-connect, so the service manager does not signal the VC label mapping immediately, but will put this into passive mode.

The following example shows the configuration:

```
*A:ALA-S-PE1>config>service>epipe# info
----------------------------------------------
```

```
...
            spoke-sdp 2:200 create
            exit
            spoke-sdp 3:300 create
            exit
            no shutdown
-----------------------------------------------
*A:ALA-S-PE1>config>service>epipe#
```

S-PE2: Specifying the **vc-switching** parameter enables a VC cross-connect, so the service manager does not signal the VC label mapping immediately, but will put this into passive mode.

The following example shows the configuration:

```
*A:ALA-S-PE2>config>service>epipe# info
-----------------------------------------------
...
            spoke-sdp 2:200 create
            exit
            spoke-sdp 3:300 create
            exit
            no shutdown
-----------------------------------------------
*A:ALA-S-PE2>config>service>epipe#
```

2.15.7 Configuring VLL Resilience

Figure 51 shows an example to create VLL resilience. The zero revert-time value means that the VLL path will be switched back to the primary immediately after it comes back up.

Figure 51 **VLL Resilience**

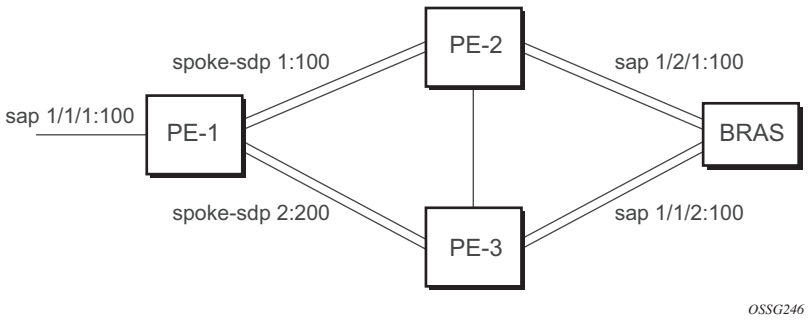

PE-1:

The following example shows the configuration on PE-1:

```
*A:ALA-48>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            spoke-sdp 1:100 endpoint "y" create
                precedence primary
            exit
            spoke-sdp 2:200 endpoint "y" create
                precedence 1
            exit
            no shutdown
----------------------------------------------
*A:ALA-48>config>service>epipe#
```

2.15.8 Configuring VLL Resilience for a Switched Pseudowire Path

Figure 52 **VLL Resilience with Pseudowire Switching**

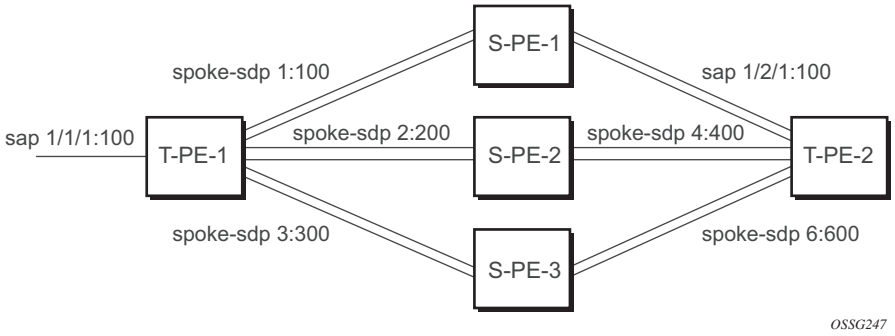

T-PE-1

The following example shows the configuration on T-PE-1.

```
*A:ALA-48>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap 1/1/1:100 endpoint "x" create
            exit
            spoke-sdp 1:100 endpoint "y" create
                precedence primary
            exit
            spoke-sdp 2:200 endpoint "y" create
                precedence 1
```

```
            exit
            spoke-sdp 3:300 endpoint "y" create
                precedence 1
            exit
            no shutdown
----------------------------------------------
*A:ALA-48>config>service>epipe#
```

T-PE-2

The following example shows the configuration on T-PE-2.

```
*A:ALA-49>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
                revert-time 100
            exit
            spoke-sdp 4:400 endpoint "y" create
                precedence primary
            exit
            spoke-sdp 5:500 endpoint "y" create
                precedence 1
            exit
            spoke-sdp 6:600 endpoint "y" create
                precedence 1
            exit
            no shutdown
----------------------------------------------
*A:ALA-49>config>service>epipe#
```

S-PE-1

The following example shows the configuration on S-PE-1.

```
*A:ALA-50>config>service>epipe# info
----------------------------------------------
...
            spoke-sdp 1:100 create
            exit
            spoke-sdp 4:400 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-49>config>service>epipe#
```

2.15.9 Configuring BGP Virtual Private Wire Service (VPWS)

2.15.9.1 Single-Homed BGP VPWS

Figure 53 shows an example topology for a BGP VPWS service used to create a virtual lease-line across an MPLS network between two sites: A and B.

Figure 53 **Single-Homed BGP VPWS Configuration Example**

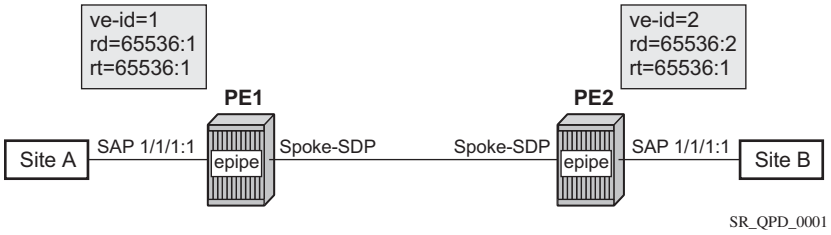

An Epipe is configured on PE1 and PE2 with BGP VPWS enabled. PE1 and PE2 are connected to site A and B, respectively, each using a SAP. The interconnection between the two PEs is achieved through a pseudowire, using Ethernet VLAN encapsulation, which is signaled using BGP VPWS over a tunnel LSP between PE1 and PE2. A MIP or MEP can be configured on a BGP VPWS SAP. However, fault propagation between a MEP and the BGP update state signaling is not supported. BGP VPWS routes are accepted only over an IBGP session.

The following example shows the BGP VPWS configuration on each PE:

```
PE1:
pw-template 1 create
    vc-type vlan
exit
epipe 1 customer 1 create
    bgp
        route-distinguisher 65536:1
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
        exit
    exit
    bgp-vpws
        ve-name PE1
            ve-id 1
        exit
        remote-ve-name PE2
            ve-id 2
        exit
        no shutdown
```

```
    exit
    sap 1/1/1:1 create
    exit
    no shutdown
exit

PE2:

pw-template 1 create
    vc-type vlan
exit
epipe 1 customer 1 create
    bgp
        route-distinguisher 65536:2
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
        exit
    exit
    bgp-vpws
        ve-name PE2
            ve-id 2
        exit
        remote-ve-name PE1
            ve-id 1
        exit
        no shutdown
    exit
    sap 1/1/1:1 create
    exit
    no shutdown
exit
```

The BGP-VPWS update can be shown using the following command:

```
A:PE1# show service l2-route-table bgp-vpws detail
===============================================================================
Services: L2 Bgp-Vpws Route Information - Summary
===============================================================================
Svc Id          : 1
VeId            : 2
PW Temp Id      : 1
RD              : *65536:2
Next Hop        : 1.1.1.2
State (D-Bit)   : up(0)
Path MTU        : 1514
Control Word    : 0
Seq Delivery    : 0
Status          : active
Tx Status       : active
CSV             : 0
Preference      : 0
Sdp Bind Id     : 17407:4294967295
===============================================================================
A:PE1#
```

2.15.9.2 Dual-Homed BGP VPWS

Single Pseudowire Example:

Figure 54 shows an example topology for a dual-homed BGP VPWS service used to create a virtual lease-line across an MPLS network between two sites: A and B. A single pseudowire is established between the designated forwarder of the dual-homed PEs and the remote PE.

Figure 54 **Example of Dual-Homed BGP VPWS with Single Pseudowire**

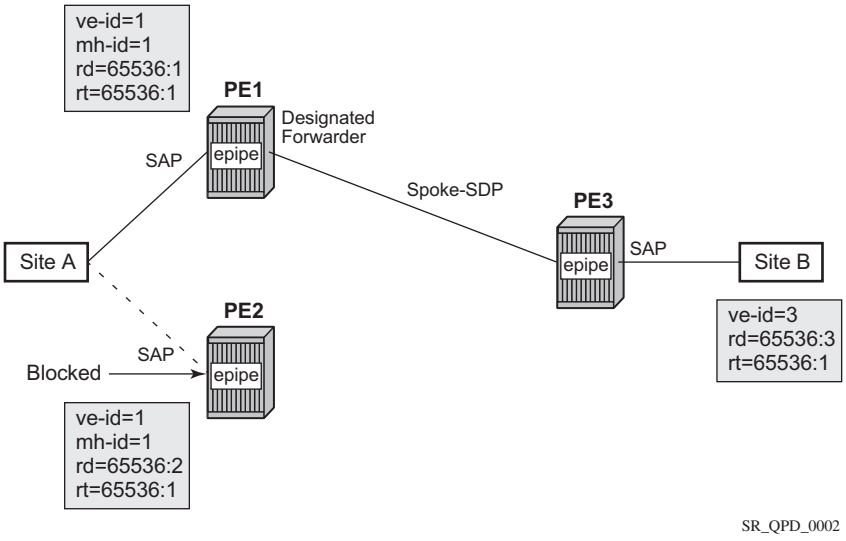

An Epipe with BGP VPWS enabled is configured on each PE. Site A is dual-homed to PE1 and PE2 with a remote PE (PE3) connected to site B; each connection uses a SAP. A single pseudowire using Ethernet Raw Mode encaps connects PE3 to PE1. The pseudowire is signaled using BGP VPWS over a tunnel LSP between the PEs.

Site A is configured on PE1 and PE2 with the BGP route selection, the site state, and the site-preference used to ensure PE1 is the designated forwarder when the network is fully operational.

The following example shows the BGP VPWS configuration on each PE.

PE1:

```
pw-template 1 create
exit
epipe 1 customer 1 create
    bgp
        route-distinguisher 65536:1
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
```

```
        exit
    exit
    bgp-vpws
        ve-name PE1
            ve-id 1
        exit
        remote-ve-name PE3
            ve-id 3
        exit
        no shutdown
    exit
    sap 1/1/1:1 create
    exit
    site "siteA" create
        site-id 1
        sap 1/1/1:1
        boot-timer 20
        site-activation-timer 5
        no shutdown
    exit
    no shutdown
exit
```

PE2:

```
pw-template 1 create
exit
epipe 1 customer 1 create
    bgp
        route-distinguisher 65536:2
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
        exit
    exit
    bgp-vpws
        ve-name PE2
            ve-id 1
        exit
        remote-ve-name PE3
            ve-id 3
        exit
        no shutdown
    exit
    sap 1/1/1:1 create
    exit
    site "siteA" create
        site-id 1
        sap 1/1/1:1
        boot-timer 20
        site-activation-timer 5
        no shutdown
    exit
    no shutdown
exit
```

PE3:

```
pw-template 1 create
exit
epipe 1 customer 1 create
    bgp
        route-distinguisher 65536:3
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
        exit
    exit
    bgp-vpws
        ve-name PE3
            ve-id 3
        exit
        remote-ve-name PE1orPE2
            ve-id 1
        exit
        no shutdown
    exit
    sap 1/1/1:1 create
    exit
    no shutdown
exit
```

Active/Standby Pseudowire Example:

Figure 55 shows an example topology for a dual-homed BGP VPWS service used to create a virtual lease-line across an MPLS network between two sites: A and B. Two pseudowires are established between the remote PE and the dual-homed PEs. The active pseudowire used for the traffic is the one connecting the remote PE to the designated forwarder of the dual-homed PEs.

Figure 55 **Example of Dual-homed BGP VPWS with Active/Standby Pseudowires**

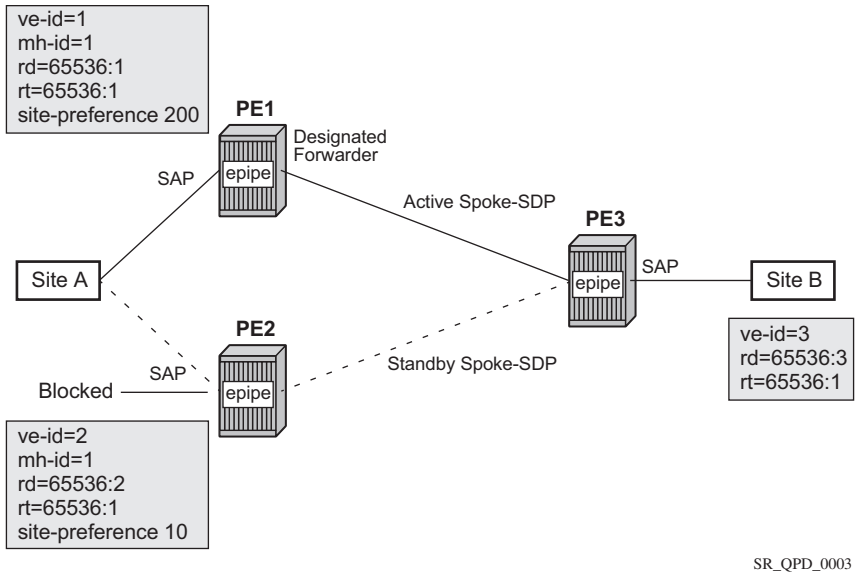

An Epipe with BGP VPWS enabled is configured on each PE. Site A is dual-homed to PE1 and PE2 with a remote PE (PE3) connected to site B; each connection uses a SAP. Active/standby pseudowires using Ethernet Raw Mode encaps connect PE3 to PE1 and PE2, respectively. The pseudowires are signaled using BGP VPWS over a tunnel LSP between the PEs.

Site A is configured on PE1 and PE2 with the site-preference set to ensure that PE1 is the designated forwarder when the network is fully operational. An endpoint is automatically created on PE3 in which the active/standby pseudowires are created.

The following example shows the BGP VPWS configuration on each PE.

PE1:

```
pw-template 1 create
exit
epipe 1 customer 1 create
    bgp
        route-distinguisher 65536:1
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
        exit
    exit
    bgp-vpws
        ve-name PE1
            ve-id 1
        exit
        remote-ve-name PE3
```

```
            ve-id 3
        exit
        no shutdown
    exit
    sap 1/1/1:1 create
    exit
    site "siteA" create
        site-id 1
        sap 1/1/1:1
        boot-timer 20
        site-activation-timer 5
        site-preference 200
        no shutdown
    exit
    no shutdown
exit
```

PE2:

```
pw-template 1 create
exit
epipe 1 customer 1 create
    bgp
        route-distinguisher 65536:2
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
        exit
    exit
    bgp-vpws
        ve-name PE2
            ve-id 2
        exit
        remote-ve-name PE3
            ve-id 3
        exit
        no shutdown
    exit
    sap 1/1/1:1 create
    exit
    site "siteA" create
        site-id 1
        sap 1/1/1:1
        boot-timer 20
        site-activation-timer 5
        site-preference 10
        no shutdown
    exit
    no shutdown
exit
```

PE3:

```
pw-template 1 create
exit
epipe 1 customer 1 create
```

```
        bgp
            route-distinguisher 65536:3
            route-target export target:65536:1 import target:65536:1
            pw-template-binding 1
            exit
        exit
        bgp-vpws
            ve-name PE3
                ve-id 3
            exit
            remote-ve-name PE1
                ve-id 1
            exit
            remote-ve-name PE2
                ve-id 2
            exit
            no shutdown
        exit
        sap 1/1/1:1 create
        exit
        no shutdown
    exit
```

2.16 Service Management Tasks

This section discusses VLL service management tasks.

2.16.1 Modifying Apipe Service Parameters

The following example shows the command usage to modify Apipe parameters, supported on the 7750 SR only:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 1/1/1:0/32 create
A:ALA-41>config>service>apipe>sap# accounting-policy 2
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe# spoke-sdp 1:4
A:ALA-41>config>service>apipe>spoke-sdp# egress
A:ALA-41>config>service>apipe>spoke-sdp>egress# vc-
 label 16
A:ALA-41>config>service>apipe>spoke-sdp>egress# exit
A:ALA-41>config>service>apipe>spoke-sdp# exit
A:ALA-41>config>service>apipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# sap 2/2/2:0/32 create
A:ALA-42>config>service>apipe>sap# accounting-policy 2
A:ALA-42>config>service>apipe>sap# exit
A:ALA-42>config>service>apipe# spoke-sdp 1:4
A:ALA-42>config>service>apipe>spoke-sdp# egress
A:ALA-42>config>service>apipe>spoke-sdp>egress# vc-
 label 16
A:ALA-42>config>service>apipe>spoke-sdp>egress# exit
A:ALA-42>config>service>apipe>spoke-sdp# exit
A:ALA-42>config>service>apipe#
```

```
PE Router 1 (ALA-41):
A:ALA-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               accounting-policy 2
```

```
                ingress
                    qos 102
                exit
                egress
                    qos 103
                exit
            exit
            spoke-sdp 1:4 create
                egress
                    vc-label 16
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#

PE Router 2 (ALA-42):
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
                accounting-policy 2
                ingress
                    qos 102
                exit
                egress
                    qos 103
                exit
            exit
            spoke-sdp 1:4 create
                egress
                    vc-label 16
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

2.16.2 Disabling an Apipe Service

An Apipe service can be shut down without deleting any service parameters.

CLI Syntax:

```
config>service#
apipe service-id
     shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# shutdown
A:ALA-41>config>service>apipe# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# shutdown
A:ALA-42>config>service>apipe# exit
```

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            shutdown
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 16
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            shutdown
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
```

```
            exit
         exit
         spoke-sdp 1:4 create
            egress
                vc-label 16
         exit
      exit
...
------------------------------------
A:ALA-42>config>service#
```

2.16.3 Re-enabling an Apipe Service

Use the following CLI syntax to re-enable an Apipe service that was shut down.

CLI Syntax:
```
config>service#
apipe service-id
     no shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# no shutdown
A:ALA-41>config>service>apipe# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# no shutdown
A:ALA-42>config>service>apipe# exit
```

2.16.4 Deleting an Apipe Service

An Apipe service cannot be deleted until the SAP is shut down. If protocols and/or a spoke-SDP are defined, they must be shut down and removed from the configuration as well.

Use the following CLI syntax to delete Apipe services.

CLI Syntax:
```
config>service#
no apipe service-id
     shutdown
     no sap sap-id
          shutdown
```

```
            no spoke-sdp [sdp-id:vc-id]
                shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 1/1/1:0/32
A:ALA-41>config>service>apipe>sap# shutdown
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe# no sap 1/1/1:0/32
A:ALA-41>config>service>apipe# spoke-sdp 1:4
A:ALA-41>config>service>apipe>spoke-sdp# shutdown
A:ALA-41>config>service>apipe>spoke-sdp# exit
A:ALA-41>config>service>apipe# no spoke-sdp 1:4
A:ALA-41>config>service>apipe# shutdown
A:ALA-41>config>service>apipe# exit
A:ALA-41>config>service# no apipe 5
```

PE router 2 (A:ALA-42):

Example:
```
Example:A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 2/2/2:0/32
A:ALA-41>config>service>apipe>sap# shutdown
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe# no sap 2/2/2:0/32
A:ALA-41>config>service>apipe# spoke-sdp 1:4
A:ALA-41>config>service>apipe>spoke-sdp# shutdown
A:ALA-41>config>service>apipe>spoke-sdp# exit
A:ALA-41>config>service>apipe# no spoke-sdp 1:4
A:ALA-41>config>service>apipe# shutdown
A:ALA-41>config>service>apipe# exit
A:ALA-41>config>service# no apipe 5
```

2.16.5 Modifying a Cpipe Service

The following example shows the Cpipe service configuration, supported on the 7750 SR only:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        cpipe 94002 customer 1 vc-type cesopsn create
            endpoint "to7705" create
            exit
            endpoint "toMC-APS" create
            exit
            sap aps-4.1.1.2.1 endpoint "toMC-APS" create
```

```
                ingress
                    qos 20
                exit
            exit
            spoke-sdp 14004:94002 endpoint "to7705" create
            exit
            spoke-sdp 100:294002 endpoint "toMC-APS" icb create
            exit
            spoke-sdp 100:194002 endpoint "to7705" icb create
            exit
            no shutdown
        exit
...
-------------------------------------------
*A:ALA-1>config>service> Cpipe#
```

2.16.6 Deleting a Cpipe Service

A Cpipe service cannot be deleted until SAPs are shut down and deleted. If a spoke-SDP is defined, it must be shut down and removed from the configuration as well.

Use the following CLI syntax to delete a Cpipe service.

CLI Syntax:

```
config>service#
[no] cpipe service-id [customer customer-id]
     [no] spoke-sdp sdp-id
          [no] shutdown
     shutdown
```

2.16.7 Modifying Epipe Service Parameters

The following example shows how to add an accounting policy to an existing SAP:

Example:

```
config>service# epipe 2
config>service>epipe# sap 2/1/3:21
config>service>epipe>sap# accounting-policy 14
config>service>epipe>sap# exit
```

The following example shows the SAP configuration:

```
ALA-1>config>service# info
-------------------------------------------
    epipe 2 customer 6 vpn 2 create
            description "Distributed Epipe service to east coast"
            sap 2/1/3:21 create
              accounting-policy 14
            exit
```

```
            spoke-sdp 2:6000 create
            exit
            no shutdown
        exit
-------------------------------------------
ALA-1>config>service#
```

2.16.8 Disabling an Epipe Service

You can shut down an Epipe service without deleting the service parameters.

CLI Syntax:
```
config>service> epipe service-id
shutdown
```

Example:
```
config>service# epipe 2
config>service>epipe# shutdown
config>service>epipe# exit
```

2.16.9 Re-enabling an Epipe Service

Use the following CLI syntax to re-enable an Epipe service that was shut down.

CLI Syntax:
```
config>service# epipe service-id
no shutdown
```

Example:
```
config>service# epipe 2
config>service>epipe# no shutdown
config>service>epipe# exit
```

2.16.10 Deleting an Epipe Service

Perform the following steps prior to deleting an Epipe service:

Step 1. Shut down the SAP and SDP.

Step 2. Delete the SAP and SDP.

Step 3. Shut down the service.

Use the following CLI syntax to delete an Epipe service.

CLI Syntax:
```
config>service
```

```
[no] epipe service-id
     shutdown
     [no] sap sap-id
          shutdown
     [no] spoke-sdp sdp-id:vc-id
          shutdown
```

Example:
```
config>service# epipe 2
config>service>epipe# sap 2/1/3:21
config>service>epipe>sap# shutdown
config>service>epipe>sap# exit
config>service>epipe# no sap 2/1/3:21
config>service>epipe# spoke-sdp 2:6000
config>service>epipe>spoke-sdp# shutdown
config>service>epipe>spoke-sdp# exit
config>service>epipe# no spoke-sdp 2:6000
config>service>epipe# epipe 2
config>service>epipe# shutdown
config>service>epipe# exit
config>service# no epipe 2
```

2.16.11 Modifying Fpipe Service Parameters

The following example shows the command usage to modify Fpipe parameters, supported on the 7750 SR only:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 1/2/1:16 create
A:ALA-41>config>service>fpipe>sap# accounting-policy 2
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe# spoke-sdp 1:4
A:ALA-41>config>service>fpipe>spoke-sdp# ingress
A:ALA-41>config>service>fpipe>spoke-sdp>filter ip 10
A:ALA-41>config>service>fpipe>spoke-sdp# exit
A:ALA-41>config>service>fpipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# sap 2/1/1.1:16 create
A:ALA-42>config>service>fpipe>sap# accounting-policy 2
A:ALA-42>config>service>fpipe>sap# exit
A:ALA-42>config>service>fpipe# spoke-sdp 1:1
A:ALA-42>config>service>fpipe>spoke-sdp# egress
```

```
A:ALA-42>config>service>fpipe>spoke-sdp>egress# filter
 ip 10
A:ALA-42>config>service>fpipe>spoke-sdp>egress# exit
A:ALA-42>config>service>fpipe>spoke-sdp# exit
A:ALA-42>config>service>fpipe#
```

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               ingress
                   filter ip 10
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               egress
                   filter ip 10
            exit
            no shutdown
        exit
```

```
...
-----------------------------------
A:ALA-42>config>service#
```

2.16.12 Disabling an Fpipe Service

An Fpipe service can be shut down without deleting any service parameters.

CLI Syntax:
```
config>service#
fpipe service-id
     shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# shutdown
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# shutdown
```

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
-----------------------------------
...
        fpipe 1 customer 1 create
            shutdown
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               ingress
                   filter ip 10
            exit
        exit
...
-----------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
----------------------------------
...
        fpipe 1 customer 1 create
            shutdown
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               egress
                   filter ip 10
            exit
        exit
...
----------------------------------
A:ALA-42>config>service#
```

2.16.13 Re-enabling an Fpipe Service

Use the following CLI syntax to re-enable an Fpipe service that was shut down.

CLI Syntax:
```
config>service#
fpipe service-id
     no shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# no shutdown
A:ALA-41>config>service>fpipe# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# no shutdown
A:ALA-42>config>service>fpipe# exit
```

2.16.14 Deleting an Fpipe Service

An Fpipe service cannot be deleted until the SAP is shut down. If protocols and/or a spoke-SDP are defined, they must be shut down and removed from the configuration as well.

Use the following CLI syntax to delete a Fpipe service.

CLI Syntax:
```
config>service#
no fpipe service-id
     shutdown
     no sap sap-id
          shutdown
     no spoke-sdp [sdp-id:vc-id]
          shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 1/1/1:0/32
A:ALA-41>config>service>fpipe>sap# shutdown
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe# no sap 1/1/1:0/32
A:ALA-41>config>service>fpipe# spoke-sdp 1:1
A:ALA-41>config>service>fpipe>spoke-sdp# shutdown
A:ALA-41>config>service>fpipe>spoke-sdp# exit
A:ALA-41>config>service>fpipe# no spoke-sdp 1:1
A:ALA-41>config>service>fpipe# shutdown
A:ALA-41>config>service>fpipe# exit
A:ALA-41>config>service# no fpipe 1
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 2/1/1.1:16
A:ALA-41>config>service>fpipe>sap# shutdown
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe# no sap 2/1/1.1:16
A:ALA-41>config>service>fpipe# spoke-sdp 1:1
A:ALA-41>config>service>fpipe>spoke-sdp# shutdown
A:ALA-41>config>service>fpipe>spoke-sdp# exit
A:ALA-41>config>service>fpipe# no spoke-sdp 1:1
A:ALA-41>config>service>fpipe# shutdown
A:ALA-41>config>service>fpipe# exit
A:ALA-41>config>service# no fpipe 1
```

2.16.15 Modifying Ipipe Service Parameters

The following example shows the command usage to modify Ipipe parameters, supported on the 7450 ESS and 7750 SR only:

Example:
```
config>service# ipipe 202
config>service>ipipe# sap 1/1/2:444
config>service>ipipe>sap# shutdown
config>service>ipipe>sap# exit
config>service>ipipe# no sap 1/1/2:444
config>service>ipipe# sap 1/1/2:555 create
config>service>ipipe>sap$ description “eth_ipipe”
config>service>ipipe>sap$ ce-address 31.31.31.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# info
```

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 202 customer 1 create
            sap 1/1/2:445 create
                description "eth_ipipe"
                ce-address 31.31.31.2
            exit
            sap 1/1/2:555 create
                description "eth_ipipe"
                ce-address 31.31.31.1
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

2.16.16 Disabling an Ipipe Service

An Ipipe service can be shut down without deleting any service parameters.

CLI Syntax:
```
config>service#
ipipe service-id
     shutdown
```

Example:
```
A:ALA-41>config>service# ipipe 202
A:ALA-41>config>service>ipipe# shutdown
```

```
A:ALA-48>config>service# info
----------------------------------------------
```

```
...
       ipipe 202 customer 1 create
           shutdown
           sap 1/1/2:445 create
               description "eth_ipipe"
               ce-address 31.31.31.2
           exit
           sap 1/1/2:555 create
               description "eth_ipipe"
               ce-address 31.31.31.1
           exit
       exit
...
-------------------------------------------
A:ALA-48>config>service#
```

2.16.17 Re-enabling an Ipipe Service

Use the following CLI syntax to re-enable an Ipipe service that was shut down.

CLI Syntax:
```
config>service#
ipipe service-id
     no shutdown
```

Example:
```
A:ALA-41>config>service# ipipe 202
A:ALA-41>config>service>ipipe# no shutdown
```

2.16.18 Deleting an Ipipe Service

An Ipipe service cannot be deleted until the SAP is shut down. If protocols and/or a spoke-SDP are defined, they must be shut down and removed from the configuration as well.

Use the following CLI syntax to delete an Ipipe service.

CLI Syntax:
```
config>service#
no ipipe service-id
     shutdown
     no sap sap-id
          shutdown
     no spoke-sdp [sdp-id:vc-id]
          shutdown
```

Example:
```
config>service# ipipe 207
config>service>ipipe# sap 1/1/2:449
```

```
config>service>ipipe>sap# shutdown
config>service>ipipe>sap# exit
config>service>ipipe# no sap 1/1/2:449
config>service>ipipe# spoke-sdp 16:516
config>service>ipipe>spoke-sdp# shutdown
config>service>ipipe>spoke-sdp# exit
config>service>ipipe# no spoke-sdp 16:516
config>service>ipipe# exit
config>service# no ipipe 207
config>service#
```

2.17 VLL Service Configuration Command Reference

This chapter describes the VLL service configuration command reference.

2.17.1 Command Hierarchies

2.17.1.1 Apipe Service Configuration Commands

```
config
    — service
        — apipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type {atm-vcc |
            atm-sdu | atm-vpc | atm-cell}] [vc-switching] [test] [name name]
        — no apipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — active-hold-delay active-hold-delay
                — no active-hold-delay
                — description description-string
                — no description
                — revert-time revert-time
                — no revert-time
            — interworking frf-5
            — no interworking
            — sap {port-id | aps-id}:[vpi/vci | vpi | vpi1.vpi2 | cp.conn-prof-id]
            — sap sap-id [no-endpoint]
            — sap sap-id [endpoint endpoint-name]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — ingress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — [no] llf
                    — oam
                        — [no] alarm-cells
                        — [no] terminate
                — [no] collect-stats
                — cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring
                    [aggregate] [car]]
```

```
— no cpu-protection
— description description-string
— no description
— dist-cpu-protection policy-name
— no dist-cpu-protection
— egress
    — [no] agg-rate
        — [no] limit-unused-bandwidth
        — [no] queue-frame-based-accounting
        — rate kilobits-per-second
        — no rate
    — policer-control-override [create]
    — no policer-control-override
        — max-rate {rate | max}
        — priority-mbs-thresholds
            — min-thresh-separation
            — [no] priority level
                — mbs-contribution size [{bytes |
                    kilobytes}]
    — policer-control-policy policy-name
    — no policer-control-policy
    — [no] policer-override
        — policer policer-id [create]
        — no policer policer-id
            — cbs size [{bytes | kilobytes}]
            — no cbs
            — mbs size [bytes | kilobytes]
            — no mbs
            — packet-byte-offset {add add-bytes | subtract
                sub-bytes}
            — percent-rate pir-percent [cir cir-percent]
            — no percent-rate
            — rate {rate | max} [cir {max | rate}]
            — stat-mode stat-mode
            — no stat-mode
    — [no] qinq-mark-top-only
    — qos policy-id [port-redirect-group queue-group-name instance
        instance-id]
    — no qos
    — [no] queue-override
        — [no] queue queue-id
            — adaptation-rule [pir adaptation-rule] [cir
                adaptation-rule]
            — no adaptation-rule
            — avg-frame-overhead percentage
            — no avg-frame-overhead
            — burst-limit {default | size [bytes |
                kilobytes]}
            — no burst-limit
            — drop-tail low
                — percent-reduction-from-mbs
                    percent
                — no percent-reduction-from-mbs
            — mbs {size [bytes | kilobytes] | default}
```

```
                                    — no mbs
                                    — monitor-depth
                                    — [no] monitor-depth
                                    — parent {[weight weight] [cir-weight cir-
                                        weight]}
                                    — no parent
                                    — percent-rate pir-percent [cir cir-percent]
                                    — no percent-rate
                                    — rate pir-rate [cir cir-rate]
                                    — no rate
                            — [no] scheduler-override
                                — [no] scheduler scheduler-name
                                    — parent [weight weight] [cir-weight cir-weight]
                                    — no parent
                                    — rate pir-rate [cir cir-rate]
                                    — no rate
                            — scheduler-policy scheduler-policy-name
                            — no scheduler-policy
                    — frame-relay
                            — scheduling-class class-id
                            — no scheduling-class
                    — ingress
                            — policer-control-override [create]
                            — no policer-control-override
                                — max-rate {rate | max}
                                — priority-mbs-thresholds
                                    — min-thresh-separation
                                    — [no] priority level
                                        — mbs-contribution size [bytes | kilobytes]
                            — [no] policer-override
                                — policer policer-id [create]
                                — no policer policer-id
                                    — cbs size [bytes | kilobytes]
                                    — no cbs
                                    — mbs {size [bytes | kilobytes] | default}
                                    — no mbs
                                    — packet-byte-offset add add-bytes | subtract
                                        sub-bytes}
                                    — percent-rate pir-percent [cir cir-percent]
                                    — no percent-rate
                                    — rate {rate | max} [cir {max | rate}]
                                    — stat-mode stat-mode
                                    — no stat-mode
                            — qos policy-id [shared-queuing] [fp-redirect-group queue-
                                group-name instance instance-id]
                            — no qos
                            — [no] queue-override
                                — [no] queue queue-id
                                    — adaptation-rule [pir adaptation-rule] [cir
                                        adaptation-rule]
                                    — no adaptation-rule
                                    — drop-tail low
                                        — percent-reduction-from-mbs
                                            percent
```

```
                                    — no percent-reduction-from-mbs
                                — mbs {size [bytes | kilobytes] | default}
                                — no mbs
                                — monitor-depth
                                — [no] monitor-depth
                                — rate pir-rate [cir cir-rate]
                                — no rate
                        — [no] scheduler-override
                            — [no] scheduler scheduler-name
                                — parent [weight weight] [cir-weight cir-weight]
                                — no parent
                                — rate pir-rate [cir cir-rate]
                                — no rate
                        — scheduler-policy scheduler-policy-name
                        — no scheduler-policy
                — multi-service-site customer-site-name
                — no multi-service-site
                — [no] shutdown
        — service-mtu octets
        — no service-mtu
        — service-name service-name
        — no service-name
        — [no] shutdown
        — signaled-vc-type-override {atm-vcc}
        — no signaled-vc-type-override
        — spoke-sdp [sdp-id[:vc-id] [no-endpoint]
        — spoke-sdp [sdp-id[:vc-id] endpoint endpoint-name [icb]
        — no spoke-sdp [sdp-id[:vc-id]
                — [no] bandwidth
                — bfd-enable
                — no bfd-enable
                — bfd-template name
                — no bfd-template
                — cell-concatenation
                    — [no] aal5-frame-aware
                    — [no] clp-change
                    — max-cells cell-count
                    — no max-cells [cell-count>]
                    — max-delay delay-time
                    — no max-delay [delay-time]
                — [no] control-channel-status
                    — [no] acknowledgment
                    — refresh-timer value
                    — no refresh-timer
                    — request-timer timer1 retry-timer timer2 [timeout-multiplier
                          multiplier]
                    — no request-timer
                — [no] control-word
                — egress
                    — qos network-policy-id port-redirect-group queue-group-name
                          [instance instance-id]
                    — no qos
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
```

```
            — ingress
                — qos network-policy-id fp-redirect-group queue-group-name
                      instance instance-id
                — no qos
                — vc-label ingress-vc-label
                — no vc-label [ingress-vc-label]
            — precedence [precedence-value | primary]
            — no precedence
            — [no] pw-path-id
                — agi agi
                — no agi
                — saii-type2 global-id:node-id:ac-id
                — no saii-type2
                — taii-type2 global-id:node-id:ac-id
                — no taii-type2
            — [no] shutdown

config
    — connection-profile conn-prof-id [create]
    — no connection-profile conn-prof-id
```

2.17.1.2 Related Apipe Commands

2.17.1.2.1 Connection Profile Commands

```
config
    — connection-profile conn-prof-id [create]
    — no connection-profile conn-prof-id
        — description description-string
        — no description
        — member encap-value [create]
        — no member encap-value
```

2.17.1.3 Cpipe Service Configuration Commands

```
config
    — service
        — cpipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type {satop-e1 |
            satop-t1 | cesopsn | cesopsn-cas}] [vc-switching] [test] [name name]
        — no cpipe service-id
            — description description-string
            — no description [description-string]
            — endpoint endpoint-name [create]
            — no endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
```

```
        — no active-hold-delay
        — description description-string
        — no description [description-string]
        — revert-time revert-time
        — no revert-time
    — sap sap-id [no-endpoint] [create]
    — sap sap-id endpoint endpoint-name [create]
    — no sap sap-id
        — accounting-policy acct-policy-id
        — no accounting-policy [acct-policy-id]
        — cem
            — packet jitter-buffer milliseconds [payload-size bytes]
            — packet payload-size bytes
            — no packet
            — [no] report-alarm [stray] [malformed] [pktloss] [overrun]
                    [underrun] [rpktloss] [rfault] [rrdi]
            — [no] rtp-header
        — [no] collect-stats
        — cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring
            [aggregate] [car]]
        — description description-string
        — no description [description-string]
        — dist-cpu-protection policy-name
        — no dist-cpu-protection
        — egress
            — [no] agg-rate
                — rate kilobits-per-second
                — no rate
                — [no] limit-unused-bandwidth
                — [no] queue-frame-based-accounting
            — [no] policer-override
                — policer policer-id [create]
                — no policer policer-id
                    — cbs size [bytes | kilobytes]
                    — no cbs
                    — mbs {size [bytes | kilobytes] | default}
                    — no mbs
                    — packet-byte-offset add add-bytes | subtract
                        sub-bytes}
                    — percent-rate pir-percent [cir cir-percent]
                    — no percent-rate
                    — rate {rate | max} [cir {max | rate}]
                    — stat-mode stat-mode
                    — no stat-mode
            — [no] qinq-mark-top-only
            — [no] qos [policy-id]
            — [no] queue-override
                — queue queue-id [create]
                — no queue queue-id
                    — adaptation-rule [pir adaptation-rule}] [cir
                        adaptation-rule}]
                    — no adaptation-rule
                    — avg-frame-overhead percent
                    — no avg-frame-overhead
```

```
                    — burst-limit {default | size [bytes | kilobytes]}
                    — no burst-limit
                    — drop-tail
                        — low
                            — percent-reduction-from-mbs
                                percent
                            — no percent-reduction-from-mbs
                    — mbs {size [bytes | kilobytes] | default}
                    — no mbs
                    — monitor-depth
                    — no monitor-depth
                    — parent {[weight weight] [cir-weight cir-weight]}
                    — no parent
                    — percent-rate pir-percent [cir cir-percent]
                    — no percent-rate
                    — rate pir-rate [cir cir-rate]
                    — no rate
                — [no] scheduler-override
                    — scheduler scheduler-name [create]
                    — no scheduler scheduler-name
                        — parent [weight weight] [cir-weight cir-weight]
                        — no parent
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — scheduler-policy scheduler-policy-name
                — no scheduler-policy
            — ingress
                — [no] policer-override
                    — policer policer-id [create]
                    — no policer policer-id
                        — cbs size [bytes | kilobytes]
                        — no cbs
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
                        — packet-byte-offset add add-bytes | subtract
                            sub-bytes}
                        — percent-rate pir-percent [cir cir-percent]
                        — no percent-rate
                        — rate {rate | max} [cir {max | rate}]
                        — stat-mode stat-mode
                        — no stat-mode
                — [no] qos [policy-id]
                — [no] queue-override
                    — queue queue-id [create]
                    — no queue queue-id
                        — adaptation-rule [pir adaptation-rule}] [cir
                            adaptation-rule}]
                        — no adaptation-rule
                        — avg-frame-overhead percent
                        — no avg-frame-overhead
                        — drop-tail
                            — low
                                — percent-reduction-from-mbs
                                    percent
```

```
                                        — no percent-reduction-from-mbs
                                    — mbs {size [bytes | kilobytes] | default}
                                    — no mbs
                                    — monitor-depth
                                    — [no] monitor-depth
                                    — rate pir-rate [cir cir-rate]
                                    — no rate
                                — [no] scheduler-override
                                    — scheduler scheduler-name [create]
                                    — no scheduler scheduler-name
                                        — parent [weight weight] [cir-weight cir-weight]
                                        — no parent
                                        — rate pir-rate [cir cir-rate]
                                        — no rate
                                — scheduler-policy scheduler-policy-name
                                — no scheduler-policy
                            — multi-service-site customer-site-name
                            — no multi-service-site
                    — service-mtu octets
                    — no service-mtu
                    — service-name service-name
                    — no service-name
                    — [no] shutdown
                    — spoke-sdp sdp-id[:vc-id] [no-endpoint] [create]
                    — spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name [icb]
                    — no spoke-sdp sdp-id[:vc-id]
                        — accounting-policy acct-policy-id
                        — no accounting-policy
                        — bandwidth bw-value
                        — bandwidth max
                        — no bandwidth
                        — bfd-enable
                        — no bfd-enable
                        — bfd-template name
                        — no bfd-template
                        — [no] collect-stats
                        — [no] control-channel-status
                            — [no] acknowledgment
                            — refresh-timer value
                            — no refresh-timer
                            — request-timer timer1 retry-timer timer2 [timeout-multiplier
                                multiplier]
                            — no request-timer
                        — [no] control-word
                        — egress
                            — qos network-policy-id port-redirect-group queue-group-name
                                [instance instance-id]
                            — no qos
                            — vc-label egress-vc-label
                            — no vc-label [egress-vc-label]
                        — ingress
                            — qos network-policy-id fp-redirect-group queue-group-name
                                instance instance-id
                            — no qos
```

```
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
                — [no] pw-path-id
                    — agi agi
                    — no agi
                    — saii-type2 global-id:node-id:ac-id
                    — no saii-type2
                    — taii-type2 global-id:node-id:ac-id
                    — no taii-type2
                — precedence [precedence-value| primary]
                — no precedence
                — [no] shutdown
```

2.17.1.4 Epipe Service Configuration Commands

- Epipe Global Commands
- Epipe SAP Configuration Commands
- Epipe Spoke SDP Configuration Commands

2.17.1.4.1 Epipe Global Commands

```
config
    — service
        — [no] epipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching]
          [test] [name name]
            — [no] bgp
                — pw-template-binding policy-id [import-rt {ext-community,.(upto 5
                  max)}]
                — no pw-template-binding policy-id
                    — [no] bfd-enable
                    — bfd-template name
                    — no bfd-template
                    — [no] shutdown
                — route-distinguisher auto-rd
                — no route-distinguisher
                — route-distinguisher rd
                — route-target {ext-community | {[export ext-community] [import ext-
                  community]}}
                — no route-target
            — [no] bgp-vpws
                — [no] remote-ve-name name
                    — ve-id value
                    — no ve-id
                — [no] shutdown
                — [no] ve-name name
                    — ve-id value
                    — no ve-id
```

- **description** *description-string*
- **no description**
- [**no**] **endpoint** *endpoint-name*
  - **active-hold-delay** *active-endpoint-delay*
  - **no active-hold-delay**
  - **description** *description-string*
  - **no description**
  - **revert-time** [*revert-time* | **infinite**]
  - **no revert-time**
  - [**no**] **standby-signaling-master**
  - [**no**] **standby-signaling-slave**
- **load-balancing**
  - [**no**] **per-service-hashing**
- **pw-port** *pw-port-id* **fpe** *fpe-id* [**create**]
- **no pw-port**
  - **egress**
    - [**no**] **shaper**
      - **int-dest-id** *name*
      - **no int-dest-id**
      - **vport** *vport*
      - **no vport**
    - **monitor-oper-group** *group-name*
    - **no monitor-oper-group**
    - [**no**] **shutdown**
- **service-mtu** *octets*
- **no service-mtu**
- **service-name** *service-name*
- **no service-name**
- **site** *name* [**create**]
- **no site**
  - **boot-timer** *seconds*
  - **no boot-timer**
  - **sap sap-***id*
  - **no sap**
  - **site-activation-timer** *seconds*
  - **no site-activation-timer**
  - **site-min-down-timer** *min-down-time*
  - **no site-min-down-timer**
  - **site-id** *value*
  - **no site-id**
  - **site-preference** *preference-value*
  - **no site-preference**
- [**no**] **shutdown**
- **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] [**no-endpoint**]
- **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **endpoint**
- **no spoke-sdp** *sdp-id*[**:***vc-id*]
  - [**no**] **bfd-enable**
  - **bfd-template** *name*
  - **no bfd-template**
  - [**no**] **control-channel-status**
    - [**no**] **acknowledgment**
    - **refresh-timer value**
    - **no refresh-timer**

```
                    — request-timer timer1 retry-timer timer2 [timeout-multiplier
                          multiplier]
                    — no request-timer
                — [no] control-word
                — hash-label
                — no hash-label
                — [no] standby-signaling-slave
                — [no] pw-path-id
                    — agi agi
                    — no agi
                    — saii-type2 global-id:node-id:ac-id
                    — no saii-type2
                    — taii-type2 global-id:node-id:ac-id
                    — no taii-type2
```

2.17.1.4.2 Epipe SAP Configuration Commands

```
config
    — service
        — epipe service-id
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
                — aarp aarpId type type
                — accounting-policy acct-policy-id
                — no accounting-policy acct-policy-id
                — app-profile app-profile-name
                — no app-profile
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — encapsulation atm-encap-type
                    — ingress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — oam
                        — [no] alarm-cells
                — cem
                    — local-ecid emulated circuit identifier
                    — no local-ecid
                    — packet jitter-buffer milliseconds [payload-size bytes]
                    — packet payload-size bytes
                    — no packet
                    — remote-ecid emulated circuit identifier
                    — no remote-ecid
                    — remote-mac ieee-address
                    — no remote-mac
                    — [no] report-alarm [stray] [malformed] [pktloss] [overrun]
                          [underrun] [rpktloss] [rfault] [rrdi]
                    — [no] rtp-header
```

```
— [no] cflowd
— [no] collect-stats
— cpu-protection policy-id {[mac-monitoring] | [eth-cfm-monitoring
    [aggregate] [car]]}
— no cpu-protection
— description description-string
— no description
— dist-cpu-protection policy-name
— no dist-cpu-protection
— egress
    — [no] agg-rate
        — [no] limit-unused-bandwidth
        — [no] queue-frame-based-accounting
        — rate kilobits-per-second
        — no rate
    — filter [ip ip-filter-id]
    — filter [ipv6 ipv6-filter-id]
    — filter [mac mac-filter-id]
    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
    — [no] hsmda-queue-override
        — packet-byte-offset {add add-bytes | subtract sub-bytes}
        — no packet-byte-offset
        — queue queue-id
        — no queue queue-id
            — mbs {size [bytes | kilobytes] | default}
            — no mbs
            — rate pir-rate
            — no rate
            — slope-policy hsmda-slope-policy-name
                allowable
            — no slope-policy
            — wrr-weight weight
            — no wrr-weight
            — secondary-shaper secondary-shaper-name
            — no secondary-shaper
            — wrr-policy hsmda-wrr-policy-name
            — no wrr-policy
        — policer-control-override [create]
        — no policer-control-override
            — max-rate {rate | max}
            — priority-mbs-thresholds
                — min-thresh-separation
                — [no] priority level
                — mbs-contribution size [bytes | kilobytes]
        — policer-control-policy policy-name
        — no policer-control-policy
        — [no] policer-override
            — policer policer-id [create]
            — no policer policer-id
                — cbs size [bytes | kilobytes]
                — no cbs
                — mbs {size [bytes | kilobytes] | default}
                — no mbs
```

```
                        — packet-byte-offset add add-bytes |
                            subtract sub-bytes}
                        — percent-rate pir-percent [cir cir-percent]
                        — no percent-rate
                        — rate {rate | max} [cir {max | rate}]
                        — stat-mode stat-mode
                        — no stat-mode
                — [no] qinq-mark-top-only
                — qos policy-id [port-redirect-group queue-group-name
                    instance instance-id]
                — no qos
                — [no] queue-override
                    — queue queue-id [create]
                    — no queue queue-id
                        — adaptation-rule [pir adaptation-rule] [cir
                            adaptation-rule]
                        — no adaptation-rule
                        — avg-frame-overhead percentage
                        — no avg-frame-overhead
                        — burst-limit {default | size [bytes |
                            kilobytes]}
                        — no burst-limit
                        — cbs size-in-kbytes
                        — no cbs
                        — drop-tail low
                            — percent-reduction-from-mbs
                                percent
                            — no percent-reduction-from-mbs
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
                        — [no] monitor-depth
                        — parent {[weight weight] [cir-weight cir-
                            weight]}
                        — percent-rate pir-percent [cir cir-percent]
                        — no percent-rate
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — [no] scheduler-override
                    — [no] scheduler scheduler-name
                        — parent [weight weight] [cir-weight cir-
                            weight]
                        — no parent
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — scheduler-policy scheduler-policy-name
                — no scheduler-policy
    — eth-cfm
        — [no] ais-enable
        — [no] collect-lmm-stats
        — collect-lmm-fc-stats
            — fc fc-name [fc-name ... (up to 8 max)]
            — no fc
            — fc-in-profile fc-name [fc-name ... (up to 8 max)]
            — no fc-in-profile
```

```
— [no] mep mep-id domain md-index association ma-index
      [direction {up | down}] primary-vlan-enable [vlan vlan-
      id]
   — [no] ais-enable
      — [no] client-meg-level [[level [level ...]]
      — low-priority-defect {allDef | macRemErrXcon}
      — [no] interface-support-enable
      — [no] interval {1 | 60}
      — [no] priority priority-value
   — [no] ccm-enable
   — [no] ccm-ltm-priority priority
   — ccm-padding-size ccm-padding
   — no ccm-padding-size ccm-padding
   — [no] csf-enable
      — multiplier multiplier-value
      — no multiplier
   — [no] description description-string
   — [no] eth-test-enable
      — [no] bit-error-threshold bit-errors
      — test-pattern {all-zeros | all-ones} [crc-enable]
      — no test-pattern
   — [no] fault-propagation-enable {use-if-tlv | suspend-
      ccm}
   — grace
      — eth-ed
         — max-rx-defect-window seconds
         — no max-rx-defect-window
         — priority priority
         — no priority
         — [no] rx-eth-ed
         — [no] tx-eth-ed
      — eth-vsm-grace
         — [no] rx-eth-vsm-grace
         — [no] tx-eth-vsm-grace
   — [no] lbm-svc-act-responder
   — low-priority-defect {allDef | macRemErrXcon |
      remErrXcon | errXcon | xcon | noXcon}
   — one-way-delay-threshold mac-address
   — no one-way-delay-threshold
   — one-way-delay-threshold seconds
   — [no] shutdown
— mip [mac mac-address] primary-vlan-enable [vlan vlan-id]
— mip default-mac
— no mip
— [no] squelch-ingress-levels [md-level [md-level…]]
— tunnel-fault [accept | ignore]
— eth-tunnel
   — path path-index tag qtag[.qtag]
   — no path path-index
— ethernet
   — [no] llf
— frame-relay
   — [no] frf-12
      — ete-fragment-threshold threshold
```

```
                    — no ete-fragment-threshold
                    — [no] interleave
                — scheduling-class class-id
                — no scheduling-class
            — [no] ignore-oper-down
            — ingress
                — filter [ip ip-filter-id]
                — filter [ipv6 ipv6-filter-id]
                — filter [mac mac-filter-id]
                — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
                — qos network-policy-id fp- redirect-group queue-group-name
                      instance instance-id
                — no qos
                — [no] hsmda-queue-override
                    — packet-byte-offset {add add-bytes | subtract sub-bytes}
                    — no packet-byte-offset
                    — [no] queue queue-id
                        — rate pir-rate
                        — no rate
                        — slope-policy hsmda-slope-policy-name
                              allowable
                        — no slope-policy
                — match-qinq-dot1p {top | bottom}
                — no match-qinq-dot1p
                — policer-control-override [create]
                — no policer-control-override
                    — max-rate {rate | max}
                    — priority-mbs-thresholds
                        — min-thresh-separation
                        — [no] priority level
                            — mbs-contribution size [bytes | kilobytes]
                — policer-control-policy policy-name
                — no policer-control-policy
                — [no] policer-override
                    — policer policer-id [create]
                    — no policer policer-id
                        — cbs size-in-kilobytes
                        — no cbs
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
                        — packet-byte-offset add add-bytes | subtract
                              sub-bytes}
                        — percent-rate pir-percent [cir cir-percent]
                        — no percent-rate
                        — rate {rate | max} [cir {max | rate}]
                        — stat-mode stat-mode
                        — no stat-mode
                — qos policy-id [shared-queuing] [fp-redirect-group queue-
                      group-name instance instance-id]
                — no qos
                — [no] queue-override
                    — [no] queue queue-id
                        — adaptation-rule [pir adaptation-rule] [cir
                              adaptation-rule]
```

```
                    — no adaptation-rule
                    — cbs size-in-kilobytes
                    — no cbs
                    — drop-tail
                        — low
                            — percent-reduction-from-mbs
                                percent
                            — no percent-reduction-from-mbs
                    — mbs {size [bytes | kilobytes] | default}
                    — no mbs
                    — [no] monitor-depth
                    — parent {[weight weight] [cir-weight cir-weight]}
                    — no parent
                    — percent-rate pir-percent [cir cir-percent]
                    — no percent-rate
                    — rate pir-rate [cir cir-rate]
                    — no rate
                — [no] scheduler-override
                    — [no] scheduler scheduler-name
                        — parent [weight weight] [cir-weight cir-weight]
                        — no parent
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — scheduler-policy scheduler-policy-name
                — no scheduler-policy
                — vlan-translation {vlan-id | copy-outer}
                — no vlan-translation
            — lag-link-map-profile link-map-profile-id
            — no lag-link-map-profile
            — lag-per-link-hash class {1 | 2 | 3} weight [1 to 1024]
            — no lag-per-link-hash
            — monitor-oper-group group-name
            — no monitor-oper-group
            — multi-service-site customer-site-name
            — no multi-service-site
            — oper-group group-name
            — no oper-group
            — ring-node ring-node-name
            — no ring-node
            — [no] shutdown
            — transit-policy {ip ip-aasub-policy-id | prefix prefix-aasub-policy-id}
            — no transit-policy
```

2.17.1.4.3 Epipe Spoke SDP Configuration Commands

```
config
    — service
        — epipe service-id
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [no-endpoint]
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] endpoint [icb]
            — no spoke-sdp sdp-id[:vc-id]
```

- **aarp** *aarpId* **type** *type*
- **accounting-policy** *acct-policy-id*
- **no accounting-policy**
- **app-profile** *app-profile-name*
- **no app-profile**
- **bandwidth** *bandwidth*
- no **bandwidth**
- [**no**] **bfd-enable**
- **bfd-template** *name*
- **no bfd-template**
- [**no**] **collect-stats**
- [**no**] **control-word**
- **cpu-protection** *policy-id* {[**mac-monitoring**] | [**eth-cfm-monitoring** [**aggregate**] [**car**]}]
- **no cpu-protection**
- [**no**] **description**
- [**no**] **egress**
  - **filter** [**ip** *ip-filter-id*]
  - **filter** [**ipv6** *ipv6-filter-id*]
  - **filter** [**mac** *mac-filter-id*]
  - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*] [**mac** *mac-filter-id*]
  - **l2tpv3**
    - **cookie** *cookie*
    - **no cookie**
  - **qos** *network-policy-id* **port-redirect-group** *queue-group-name* [**instance** *instance-id*]
  - **no qos**
  - [**no**] **vc-label** *egress-vc-label*
- [**no**] **entropy-label**
- **eth-cfm**
  - [**no**] **ais-enable**
    - [**no**] **client-meg-level** [*level* [*level* ...]]
    - [**no**] **interface-support-enable**
    - [**no**] **interval** {**1** | **60**}
    - **low-priority-defect** {**allDef** | **macRemErrXcon**}
    - [**no**] **priority** *priority-value*
  - [**no**] **ccm-enable**
  - [**no**] **ccm-ltm-priority** *priority*
  - **ccm-padding-size** *ccm-padding*
  - **no ccm-padding-size** *ccm-padding*
  - [**no**] **collect-lmm-stats**
  - **collect-lmm-fc-stats**
    - **fc** *fc-name* [*fc-name* ... (up to 8 max)]
    - **no fc**
    - **fc-in-profile** *fc-name* [*fc-name* ... (up to 8 max)]
    - **no fc-in-profile**
  - [**no**] **csf-enable**
    - **multiplier** *multiplier-value*
    - **no multiplier**
  - [**no**] **description**
  - [**no**] **eth-test-enable**
    - [**no**] **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
  - [**no**] **fault-propagation-enable** {**use-if-tlv** | **suspend-ccm**}
  - [**no**] **one-way-delay-threshold** *seconds*

```
— [no] mip [{mac mac-address | default-mac}] [primary-vlan-
      enable vlan-id]
— mep mep-id domain md-index association ma-index [direction
      {up | down}] [primary-vlan-enable]
— no mep mep-id domain md-index association ma-index
    — [no] ccm-enable
    — ccm-ltm-priority priority
    — no ccm-ltm-priority
        — [no] description
        — [no] eth-test-enable
    — ccm-padding-size ccm-padding
    — no ccm-padding-size ccm-padding
    — fault-propagation-enable {use-if-tlv | suspend-ccm}
    — no fault-propagation-enable
    — grace
        — eth-ed
            — max-rx-defect-window seconds
            — no max-rx-defect-window
            — priority priority
            — no priority
            — [no] rx-eth-ed
            — [no] tx-eth-ed
        — eth-vsm-grace
            — [no] rx-eth-vsm-grace
            — [no] tx-eth-vsm-grace
    — [no] lbm-svc-act-responder
    — low-priority-defect {allDef | macRemErrXcon |
        remErrXcon | errXcon | xcon | noXcon}
    — [no] shutdown
  — [no] squelch-ingress-levels [md-level [md-level…]]
— [no] force-qinq-vc-forwarding
— [no] force-vlan-vc-forwarding
— [no] hash-label
— [no] ingress
    — filter [ip ip-filter-id]
    — filter [ipv6 ipv6-filter-id]
    — filter [mac mac-filter-id]
    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
    — l2tpv3
        — cookie [cookie1] [cookie2]
        — no cookie
    — qos network-policy-id fp-redirect-group queue-group-name
          instance instance-id
    — no qos
    — [no] vc-label egress-vc-label
— monitor-oper-group group-name
— no monitor-oper-group
— precedence [precedence-value | primary]
— no precedence
— [no] pw-status-signaling
— [no] shutdown
— [no] standby-signaling-slave
— transit-policy {ip ip-aasub-policy-id | prefix prefix-aasub-policy-id}
— no transit-policy
```

```
            — [no] use-sdp-bmac
            — vlan-vc-tag 0 to 4094
            — no vlan-vc-tag [0 to 4094]
        — spoke-sdp-fec spoke-sdp-fec-id [fec fec-type] [aii-type aii-type] [create]
        — spoke-sdp-fec spoke-sdp-fec-id no-endpoint
        — spoke-sdp-fec spoke-sdp-fec-id [fec fec-type] [aii-type aii-type] [create]
            endpoint name [icb]
        — no spoke-sdp-fec spoke-sdp-fec-id
            — [no] auto-config
            — path name
            — no path
            — precedence prec-value
            — precedence primary
            — no precedence
            — pw-template-bind policy-id
            — no pw-template-bind
            — retry-count retry-count
            — no retry-count
            — retry-timer retry-timer
            — no retry-timer
            — saii-type2 global-id:prefix:ac-id
            — no saii-type2
            — [no] shutdown
            — signaling signaling
            — [no] standby-signaling-slave
            — taii-type2 global-id:prefix:ac-id
            — no taii-type2
```

2.17.1.4.4 Template Commands

```
configure
    — service
        — template
            — epipe-sap-template name [create]
            — no epipe-sap-template name
                — egress
                    — [no] filter
                        — ip filter-id
                        — no ip
                        — ipv6 filter-id
                        — no ipv6
                        — mac filter-id
                        — no mac
                    — qos policy-id
                    — no qos
                — ingress
                    — [no] filter
                        — ip filter-id
                        — no ip
                        — ipv6 filter-id
                        — no ipv6
```

```
                    — mac filter-id
                    — no mac
                — qos policy-id {shared-queuing | multipoint-shared}
                — qos policy-id
                — no qos
```

2.17.1.5 Fpipe Service Configuration Commands

```
config
    — service
        — fpipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type fr-dcli] [vc-
            switching] [name name]
        — no fpipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — description description-string
                — no description
                — revert-time revert-time
                — no revert-time
            — sap sap-id [no-endpoint]
            — sap sap-id endpoint endpoint-name
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — [no] collect-stats
                — cpu-protection policy-id {[mac-monitoring] | [eth-cfm-monitoring
                    [aggregate] [car]]}
                — description description-string
                — no description
                — dist-cpu-protection policy-name
                — no dist-cpu-protection
                — egress
                    — [no] agg-rate
                        — rate kilobits-per-second
                        — no rate
                        — [no] limit-unused-bandwidth
                        — [no] queue-frame-based-accounting
                    — filter [ip ip-filter-id]
                    — filter [ipv6 ipv6-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — policer-control-override [create]
                    — no policer-control-override
                        — max-rate {rate | max}
                        — priority-mbs-thresholds
                            — min-thresh-separation
                            — [no] priority level
                                — mbs-contribution size [bytes |
                                    kilobytes]
```

- **policer-control-policy** *policy-name*
- **no policer-control-policy**
- [**no**] **policer-override**
  - **policer** *policer-id* [**create**]
  - **no policer** *policer-id*
    - **cbs** ***size*** [**bytes** | **kilobytes**]
    - **no cbs**
    - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
    - **no mbs**
    - **packet-byte-offset add** ***add-bytes*** | **subtract** ***sub-bytes***}
    - **percent-rate** ***pir-percent*** [**cir** ***cir-percent***]
    - **no percent-rate**
    - **rate** {***rate*** | **max**} [**cir** {**max** | ***rate***}]
    - **stat-mode** ***stat-mode***
    - **no stat-mode**
- [**no**] **qinq-mark-top-only**
- **qos** *policy-id*
- **no qos**
- [**no**] **queue-override**
  - [**no**] **queue** *queue-id*
    - **adaptation-rule** [**pir** adaptation-rule] [**cir** adaptation-rule]
    - **no adaptation-rule**
    - **avg-frame-overhead** ***percent***
    - **no avg-frame-overhead**
    - **burst-limit** {**default** | *size* [**bytes** | **kilobytes**]}
    - **no burst-limit**
    - **drop-tail low**
      - **percent-reduction-from-mbs** ***percent***
      - **no percent-reduction-from-mbs**
    - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
    - **no mbs**
    - **monitor-depth**
    - [**no**] **monitor-depth**
    - **parent** {[**weight** ***weight***] [**cir-weight** ***cir-weight***]}
    - **no parent**
    - **percent-rate** ***pir-percent*** [**cir** ***cir-percent***]
    - **no percent-rate**
    - **rate** ***pir-rate*** [**cir** ***cir-rate***]
    - **no rate**
  - [**no**] **scheduler-override**
    - [**no**] **scheduler** *scheduler-name*
      - **parent** [**weight** ***weight***] [**cir-weight** ***cir-weight***]
      - **no parent**
      - **rate** ***pir-rate*** [**cir** ***cir-rate***]
      - **no rate**
  - **scheduler-policy** *scheduler-policy-name*
  - **no scheduler-policy**
- **frame-relay**
  - **scheduling-class** *class-id*

```
            — no scheduling-class
        — ingress
            — filter [ip ip-filter-id]
            — filter [ipv6 ipv6-filter-id]
            — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
            — [no] policer-override
                — policer policer-id [create]
                — no policer policer-id
                    — cbs size [bytes | kilobytes]
                    — no cbs
                    — mbs {size [bytes | kilobytes] | default}
                    — no mbs
                    — packet-byte-offset add add-bytes | subtract
                        sub-bytes}
                    — percent-rate pir-percent [cir cir-percent]
                    — no percent-rate
                    — rate {rate | max} [cir {max | rate}]
                    — stat-mode stat-mode
                    — no stat-mode
            — qos policy-id [shared-queuing]
            — no qos
            — [no] queue-override
                — [no] queue queue-id
                        — adaptation-rule [pir adaptation-rule] [cir
                            adaptation-rule]
                        — no adaptation-rule
                        — avg-frame-overhead percent
                        — no avg-frame-overhead
                        — drop-tail low
                            — percent-reduction-from-mbs
                                percent
                            — no percent-reduction-from-mbs
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
                        — monitor-depth
                        — [no] monitor-depth
                        — rate pir-rate [cir cir-rate]
                        — no rate
            — [no] scheduler-override
                — [no] scheduler scheduler-name
                    — parent [weight weight] [cir-weight cir-weight]
                    — no parent
                    — rate pir-rate [cir cir-rate]
                    — no rate
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
        — multi-service-site customer-site-name
        — no multi-service-site
        — [no] shutdown
    — service-mtu octets
    — no service-mtu
    — service-name service-name
```

```
— no service-name
— [no] shutdown
— spoke-sdp sdp-id[:vc-id] [no-endpoint]
— spoke-sdp sdp-id[:vc-id] endpoint endpoint-name [icb]
— no spoke-sdp sdp-id[:vc-id]
    — bandwidth bandwidth
    — no bandwidth
    — bfd-enable
    — no bfd-enable
    — bfd-template name
    — no bfd-template
    — egress
        — filter [ip ip-filter-id]
        — filter [ipv6 ipv6-filter-id]
        — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
        — qos network-policy-id port-redirect-group queue-group-name
                [instance instance-id]
        — no qos
        — vc-label ingress-vc-label
        — no vc-label [ingress-vc-label]
    — [no] entropy-label
    — [no] hash-label
    — ingress
        — filter [ip ip-filter-id]
        — filter [ipv6 ipv6-filter-id]
        — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
        — qos network-policy-id fp-redirect-group queue-group-name
                instance instance-id
        — no qos
        — vc-label ingress-vc-label
        — no vc-label [ingress-vc-label]
    — precedence [precedence-value| primary]
    — no precedence
    — [no] shutdown
```

2.17.1.6 Ipipe Service Configuration Commands

```
config
    — service
        — ipipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching] [name
            name]
        — no ipipe service-id
            — ce-address-discovery [ipv6] [keep]
            — [no] ce-address-discovery
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — description description-string
                — no description
```

```
            — revert-time revert-time
            — no revert-time
        — eth-legacy-fault-notification
            — recovery-timer timer-value
            — [no] recovery-timer
            — [no] shutdown
        — sap sap-id [no-endpoint]
        — sap sap-id endpoint endpoint-name
        — [no] sap eth-tunnel-tunnel-id [:eth-tunnel-sap-id] [create]
        — no sap sap-id
            — accounting-policy acct-policy-id
            — no accounting-policy
            — atm
                — egress
                    — traffic-desc traffic-desc-profile-id
                    — no traffic-desc
                — encapsulation atm-encap-type
                — ingress
                    — traffic-desc traffic-desc-profile-id
                    — no traffic-desc
                — oam
                    — [no] alarm-cells
            — ce-address ip-address
            — no ce-address
            — collect-stats
            — no collect-stats
            — cpu-protection policy-id {[mac-monitoring] | [eth-cfm-monitoring
              [aggregate] [car]]}
            — no cpu-protection
            — description description-string
            — no description
            — dist-cpu-protection policy-name
            — no dist-cpu-protection
            — egress
                — [no] agg-rate
                    — rate kilobits-per-second
                    — no rate
                    — [no] limit-unused-bandwidth
                    — [no] queue-frame-based-accounting
                — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
                — no filter {ip ip-filter-id | ipv6 ipv6-filter-id}
                — [no] hsmda-queue-override
                    — secondary-shaper secondary-shaper-name
                    — no secondary-shaper
                    — wrr-policy hsmda-wrr-policy-name
                    — no wrr-policy
                    — packet-byte-offset {add add-bytes | subtract sub-bytes}
                    — no packet-byte-offset
                    — queue queue-id
                    — no queue queue-id
                        — wrr-weight weight
                        — no wrr-weight
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
```

```
                    — rate pir-rate
                    — no rate
                    — slope-policy hsmda-slope-policy-name
                        allowable
                    — no slope-policy
— [no] policer-override
    — policer policer-id [create]
    — no policer policer-id
        — cbs size [bytes | kilobytes]
        — no cbs
        — mbs {size [bytes | kilobytes] | default}
        — no mbs
        — packet-byte-offset add add-bytes | subtract
            sub-bytes}
        — percent-rate pir-percent [cir cir-percent]
        — no percent-rate
        — rate {rate | max} [cir {max | rate}]
        — stat-mode stat-mode
        — no stat-mode
— qinq-mark-top-only
— qos policy-id
— no qos
— [no] queue-override
    — [no] queue queue-id
        — adaptation-rule [pir adaptation-rule] [cir
            adaptation-rule]
        — no adaptation-rule
        — avg-frame-overhead percent
        — no avg-frame-overhead
        — burst-limit {default | size [bytes |
            kilobytes]}
        — no burst-limit
        — drop-tail low
            — percent-reduction-from-mbs
                percent
            — no percent-reduction-from-mbs
        — mbs {size [bytes | kilobytes] | default}
        — no mbs
        — monitor-depth
        — [no] monitor-depth
        — parent {[weight weight] [cir-weight cir-
            weight]}
        — no parent
        — percent-rate pir-percent [cir cir-percent]
        — no percent-rate
        — rate pir-rate [cir cir-rate]
        — no rate
— [no] scheduler-override
    — [no] scheduler scheduler-name
        — parent [weight weight] [cir-weight cir-weight]
        — no parent
        — rate pir-rate [cir cir-rate]
        — no rate
— scheduler-policy scheduler-policy-name
```

```
        — no scheduler-policy
    — eth-cfm
        — [no] collect-lmm-stats]
        — collect-lmm-fc-stats
            — fc fc-name [fc-name ... (up to 8 max)]
            — no fc
            — fc-in-profile fc-name [fc-name ... (up to 8 max)]
            — no fc-in-profile
        — [no] mep mep-id domain md-index association ma-index
            [direction {up | down}]
            — [no] ccm-enable
            — [no] ccm-ltm-priority priority
            — [no] description
            — [no] eth-test-enable
                — [no] bit-error-threshold bit-errors
                — [no] test-pattern {all-zeros | all-ones} [crc-
                    enable]
            — [no] fault-propagation-enable {use-if-tlv | suspend-
                ccm}
            — grace
                — eth-ed
                    — max-rx-defect-window seconds
                    — no max-rx-defect-window
                    — priority priority
                    — no priority
                    — [no] rx-eth-ed
                    — [no] tx-eth-ed
                — eth-vsm-grace
                    — [no] rx-eth-vsm-grace
                    — [no] tx-eth-vsm-grace
            — low-priority-defect {allDef | macRemErrXcon |
                remErrXcon | errXcon | xcon | noXcon}
            — [no] one-way-delay-threshold mac-address
            — [no] one-way-delay-threshold <seconds>
            — [no] shutdown
        — [no] mip [{mac mac-address | default-mac}]
        — [no] squelch-ingress-levels [md-level [md-level…]]
        — tunnel-fault [accept | ignore]
    — eth-tunnel
        — path path-index tag qtag[.qtag]
        — no path path-index
    — ingress
        — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
        — no filter {ip ip-filter-id | ipv6 ipv6-filter-id}
        — match-qinq-dot1p {top | bottom}
        —  no match-qinq-dot1p
        — [no] policer-override
            — policer policer-id [create]
            — no policer policer-id
                — cbs size [bytes | kilobytes]
                — no cbs
                — mbs {size [bytes | kilobytes] | default}
                — no mbs
```

```
                        — packet-byte-offset add add-bytes | subtract
                            sub-bytes}
                        — percent-rate pir-percent [cir cir-percent]
                        — no percent-rate
                        — rate {rate | max} [cir {max | rate}]
                        — stat-mode stat-mode
                        — no stat-mode
                — qos policy-id [shared-queuing]
                — no qos
                — [no] queue-override
                    — [no] queue queue-id
                        — adaptation-rule [pir adaptation-rule] [cir
                            adaptation-rule]
                        — no adaptation-rule
                        — drop-tail
                            — low
                                — percent-reduction-from-mbs
                                    percent
                                — no percent-reduction-from-mbs
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
                        — monitor-depth
                        — [no] monitor-depth
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — [no] scheduler-override
                    — [no] scheduler scheduler-name
                        — parent [weight weight] [cir-weight cir-weight]
                        — no parent
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — scheduler-policy scheduler-policy-name
                — no scheduler-policy
            — lag-link-map-profile link-map-profile-id
            — no lag-link-map-profile
            — lag-per-link-hash class {1 | 2 | 3} weight [1 to 1024]
            — no lag-per-link-hash
            — mac [ieee-address]
            — no mac
            — mac-refresh [refresh interval]
            — no mac-refresh
            — multi-service-site customer-site-name
            — no multi-service-site
            — [no] shutdown
            — [no] use-broadcast-mac
        — service-mtu octets
        — no service-mtu
        — service-name service-name
        — no service-name
        — [no] shutdown
        — spoke-sdp [sdp-id[:vc-id] [no-endpoint]
        — spoke-sdp [sdp-id[:vc-id] endpoint endpoint-name [icb]
        — no spoke-sdp sap-id
            — aarp aarp-id type {subscriber-side-shunt | network-side-shunt}
```

```
— no aarp
— app-profile app-profile-name
— no app-profile
— bandwidth bandwidth
— no bandwidth
— bfd-enable
— no bfd-enable
— bfd-template name
— no bfd-template
— ce-address ip-address
— no ce-address
— [no] control-word
— egress
    — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
    — no filter {ip ip-filter-id | ipv6 ipv6-filter-id}
    — qos network-policy-id port-redirect-group queue-group-name
          [instance instance-id]
    — no qos
    — [no] vc-label vc-label
— [no] entropy-label
— [no] hash-label
— ingress
    — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
    — no filter {ip ip-filter-id | ipv6 ipv6-filter-id}
    — qos network-policy-id fp-redirect-group queue-group-name
          instance instance-id
    — no qos
    — vc-label ingress-vc-label
    — no vc-label [ingress-vc-label]
— precedence [precedence-value| primary]
— no precedence
— [no] shutdown
— [no] stack-capability-signaling
```

2.17.2 Command Descriptions

- Generic Commands
- VLL Global Commands
- VLL SAP Commands
- VLL Frame Relay Commands
- VLL SDP Commands
- Service Commands

2.17.2.1 Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context
config>service>apipe
config>service>apipe>sap
config>service>apipe>spoke-sdp
config>service>cpipe
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe
config>service>epipe>bgp-vpws
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>fpipe
config>service>fpipe>sap
config>service>fpipe>spoke-sdp
config>service>ipipe
config>service>ipipe>sap
config>service>ipipe>spoke-sdp
config>service>epipe>pw-port

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described as follows in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases **Service Admin State —** Bindings to an SDP within the service will be put into the out-of-service state when the service is shutdown. While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.

Service Operational State — A service is regarded as operational providing that at least one SAP and one SDP are operational or if two SAPs are operational.

SDP (global) — When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

description

Syntax **description** *description-string*
no description

Context config>service>apipe
config>service>apipe>sap
config>service>apipe>endpoint
config>service>cpipe
config>service>cpipe>endpoint
config>service>cpipe>sap
config>service>epipe
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>epipe>endpoint
config>service>fpipe
config>service>fpipe>sap
config>service>fpipe>endpoint
config>service>ipipe
config>service>ipipe>sap
config>service>ipipe>endpoint

Description This command creates a text description stored in the configuration file for a configuration context. The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

2.17.2.2 Service Commands

apipe

Syntax **apipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**atm-vcc** | **atm-sdu** | **atm-vpc** | **atm-cell**}] [**vc-switching**] [**test**] [**create**] [**name** *name*]
no apipe *service-id*

Context config>service

Description The Apipe service provides a point-to-point Layer 2 VPN connection to a remote SAP or to another local SAP. An Apipe can connect an ATM or Frame Relay endpoint either locally or over a PSN to a remote endpoint of the same type or of a different type and perform interworking between the two access technologies.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7450 ESS or 7750 SR on which this service is defined.

Values *service-id*:1 to 2147483648
svc-name: 64 characters maximum

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default null (0)

vc-type — Keyword that specifies a 15 bit value that defines the type of the VC signaled to the peer. Its values are defined in IETF Draft *draft-ietf-pwe3-iana-allocation* and it defines both the signaled VC type as well as the resulting datapath encapsulation over the Apipe.

Values atm-vcc, atm-sdu, atm-vpc, atm-cell

Default atm-sdu

vc-switching — Specifies if the pseudowire switching signaling is used for the spoke-SDPs configured in this service.

test — Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs. This parameter is not supported on the 7950 XRS.

name *name* — A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration data but unused by SR OS. The name is tied to the **service-name** in the service context (setting either **service-name** or **name** will cause the other to change as well).

cpipe

Syntax **cpipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**satop-e1** | **satop-t1** | [**vc-switching**] |**cesopsn** | **cesopsn-cas**}] [**vc-switching**] [**test**] [**create**] [**name** *name*]
no cpipe *service-id*

Context config>service

Description This command configures a Circuit Emulation Services instance.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. After a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

After a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

By default, no services exist until they are explicitly created with this command.

The **no** form of this command deletes the service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1to 2147483648

svc-name: Specifies an existing service name up to 64 characters in length.

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

Default null (0)

vc-type — The vc-type defines the type of unstructured or structured circuit emulation service to be configured.

Values **satop-e1**: Unstructured E1 circuit emulation service.
satop-t1: Unstructured DS1 circuit emulation service.
cesopsn: Basic structured N*64 kb/s circuit emulation service.
c**esopsn-cas:** Structured N*64 kb/s circuit emulation service with signaling.

vc-switching — Specifies if the pseudowire switching signaling is used for the spoke-SDPs configured in this service.

test — Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs. This parameter applies to the 7450 ESS and 7750 SR only.

create — Keyword used to create the service. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

name *name* — A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration data but unused by SR OS. The name is tied to the **service-name** in the service context (setting either **service-name** or **name** will cause the other to change as well).

epipe

Syntax **epipe** *service-id* **customer** *customer-id* [**vpn** *vpn-id*] [**vc-switching**] [**create**] [**name** *name*]
epipe *service-id* [**test**] [**create**] [**name** *name*]
no epipe *service-id*

Context config>service

Description This command configures an Epipe service instance. This command is used to configure a point-to-point epipe service. An Epipe connects two endpoints defined as Service Access Points (SAPs). Both SAPs may be defined in one 7450 ESS, 7750 SR, or 7950 XRS or they may be defined in separate devices connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far end SAP is generalized into a Service Distribution Point (SDP). This SDP describes a destination and the encapsulation method used to reach it.

No MAC learning or filtering is provided on an Epipe.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. After a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

After a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

By default, no epipe services exist until they are explicitly created with this command.

The **no** form of this command deletes the epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Cpipe services are enabled on the 7450 ESS in mixed mode.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7450 ESS, 7750 SR, or 7950 XRS on which this service is defined.

Values *service-id*: 1 to 2147483648
svc-name: 64 characters maximum

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

Default null (0)

vc-switching — Specifies if the pseudowire switching signaling is used for the spoke-SDPs configured in this service.

test — Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs. This parameter applies to the 7450 ESS and 7750 SR only.

create — Keyword used to create the service instance. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

name *name* — A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration data but unused by SR OS. The name is tied to the **service-name** in the service context (setting either **service-name** or **name** will cause the other to change as well).

fpipe

Syntax **fpipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-type** *fr-dlci*] [**vc-switching**] [**name** *name*]
no fpipe *service-id*

Context config>service

Description This command configures an Fpipe service. An Fpipe provides a point-to-point L2 VPN connection to a remote SAP or to another local SAP. An Fpipe connects only Frame Relay endpoints either locally or over a PSN to a remote endpoint of the same type.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values *service-id*: 1 to 2147483648
svc-name: 64 characters maximum

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default null (0)

vc-type — Specifies a 15 bit value that defines the type of the VC signaled to the peer. Its values are defined in *draft-ietf-pwe3-iana-allocation* and it defines both the signaled VC type as well as the resulting datapath encapsulation over the Apipe.

Values fr-dlci

vc-switching — Specifies if the pseudowire switching signaling is used for the spoke-SDPs configured in this service.

name *name* — A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration data but unused by SR OS. The name is tied to the **service-name** in the service context (setting either **service-name** or **name** will cause the other to change as well).

ipipe

Syntax **ipipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-switching**] [**name** *name*]
no ipipe *service-id*

Context config>service

Description This command configures an IP-Pipe service.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7450 ESS or 7750 SR on which this service is defined.

Values *service-id*: 1 to 2147483648

svc-name: 64 characters maximum

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default null (0)

vc-switching — Specifies if the pseudowire switching signaling is used for the spoke-SDPs configured in this service.

create — Keyword used to create the Ipipe service instance. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

name *name* — A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration data but unused by SR OS. The name is tied to the **service-name** in the service context (setting either **service-name** or **name** will cause the other to change as well).

2.17.2.3 VLL Global Commands

bgp

Syntax **bgp**

Context config>service>epipe

Description This command enters the context to configure the BGP related parameters BGP used for multi-homing and BGP VPWS.

The **no** form of this command removes the string from the configuration.

pw-template-binding

Syntax **pw-template-binding** *policy-id* [**import-rt** {*ext-community*,.(upto 5 max)}]
no pw-template-binding *policy-id*

Context config>service>epipe>bgp

Description This command binds the advertisements received with the route targets (RT) that match the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present, or if multiple matches are found, the numerically lowest pw-template is used.

The pw-template-binding applies to BGP-VPWS when enabled in the Epipe.

For BGP VPWS, the following additional rules govern the use of pseudowire-template:

- On transmission, the settings for the L2-Info extended community in the BGP updates are derived from the pseudowire template attributes. If multiple pseudowire template bindings (with or without import-rt) are specified for the same VPWS instance the first pw-template entry will be used.
- On reception, the values of the parameters in the L2-Info extended community of the BGP updates are compared with the settings from the corresponding pseudowire template bindings. The following steps are used to determine the local pw-template:
  - The RT values are matched to determine the pw-template.
  - If multiple pw-template-binding matches are found from the previous step, the first (numerically lowest) configured pw-template entry will be considered.
  - If the value used for Layer 2 MTU (unless the value zero is received) does not match the pseudowire is created but with the operationally down state.
  - If the value used for the S (sequenced delivery) flags is not zero the pseudowire is not created.

The **tools perform** commands can be used to control the application of changes in pw-template for BGP-VPWS.

The **no** form of the command removes the values from the configuration.

Parameters *policy-id* — Specifies an existing policy ID.

Values 1 to 2147483647

import-rt ext-comm — Specifies the communities allowed to be accepted from remote PE neighbors. An extended BGP community in the type:x:y format. The value x can be an integer or IP address. The type can be the target or origin.

Values
target:{ip-addr:comm-val | 2byte-asnumber:ext-comm-val| 4byte-snumber:comm-val}

ip-addr	a.b.c.d
comm-val	0 to 65535
2byte-asnumber	0 to 65535
ext-comm-val	0 to 4294967295
4byte-asnumber	0 to 4294967295

route-distinguisher

Syntax **route-distinguisher auto-rd**
no route-distinguisher
route-distinguisher *rd*

Context config>service>epipe>bgp

Description This command configures the Route Distinguisher (RD) component that is signaled in the MPBGP NLRI for L2VPN AFI. This value is used for BGP multi-homing and BGP-VPWS.

An RD value must be configured under BGP node.

Alternatively, the **auto-rd** option allows the system to automatically generate an RD based on the **bgp-auto-rd-range** command configured at the service level.

Format: Six bytes, other 2 bytes of type will be automatically generated.

Parameters *ip-addr:comm-val* — Specifies the IP address.

Values ip-addr: a.b.c.d
comm-val: 0 to 65535
as-number:

as-number:ext-comm-val — Specifies the AS number.

Values as-number: 1 to 65535
ext-comm-val: 0 to 4294967295

auto-rd — The system will generate an RD for the service according to the IP address and range configured in the **bgp-auto-rd-range** command.

rd — Specifies the route distinguisher.

Values

<rd>	<ip-addr:comm-val> | <2byte-asnumber:ext-comm-val>|<4byte-asnumber:comm-val>	
	ip-addr	a.b.c.d
	comm-val	0 to 65535
	2byte-asnumber	1 to 65535
	ext-comm-val	0 to 4294967295
	4byte-asnumber	0 to 4294967295

route-target

Syntax **route-target** {*ext-community* | {[**export** *ext-community*] [**import** *ext-community*]}}
no route-target

Context config>service>epipe>bgp

Description This command configures the route target (RT) component that is signaled in the related MPBGP attribute to be used for BGP multi-homing and BGP-VPWS when configured in the Epipe service. The ext-comm can have two formats:

- A two-octet AS-specific extended community, IPv4 specific extended community.
- An RT value must be configured under BGP node when BGP Epipe is configured.

Parameters *export ext-community* — Specifies communities allowed to be sent to remote PE neighbors.

import ext-community — Specifies communities allowed to be accepted from remote PE neighbors.

bgp-vpws

Syntax [**no**] **bgp-vpws**

Context config>service>epipe

Description This command enters the context to configure BGP-VPWS parameters and addressing.

Default no bgp-vpws

remote-ve-name

Syntax [**no**] **remote-ve-name** *name*

Context config>service>epipe>bgp-vpws

Description This command creates or edits a remote-ve-name. A single remote-ve-name can be created per BGP VPWS instance if the service is single-homed or uses a single pseudowire to connect to a pair of dual-homed systems. When the service requires active/standby pseudowires to be created to remote dual-homed systems then two remote-ve-names must be configured.

This context defines the remote PE to which a pseudowire will be signaled.

remote-ve-name commands can be added even if bgp-vpws is not shutdown.

The **no** form of the command removes the configured remote-ve-name from the bgp vpws node. It can be used when the BGP VPWS status is either shutdown or “no shutdown”.

Parameters *name* — Specifies a site name up to 32 characters in length.

ve-id

Syntax **ve-id** *value*

no ve-id

Context config>service>epipe>bgp-vpws>ve-name
config>service>epipe>bgp-vpws>remote-ve-name

Description This command configures a ve-id for either the local VPWS instance when configured under the ve-name, or for the remote VPWS instance when configured under the remote-ve-name.

A single ve-id can be configured per ve-name or remote-ve-name. The ve-id can be changed without shutting down the VPWS instance. When the ve-name ve-id changes, BGP withdraws the previously advertised route and sends a route-refresh to all the peers which would result in reception of all the remote routes again. The old PWs are removed and new ones are instantiated for the new ve-id value.

When the remote-ve-name ve-id changes, BGP withdraws the previously advertised route and send a new update matching the new ve-id. The old pseudowires are removed and new ones are instantiated for the new ve-id value.

NLRIs received whose advertised ve-id does not match the list of ve-ids configured under the remote ve-id will not have a spoke-SDP binding auto-created but will remain in the BGP routing table but not in the Layer 2 route table. A change in the locally configured ve-ids may result in auto-sdp-bindings either being deleted or created, based on the new matching results.

Each ve-id configured within a service must be unique.

The **no** form of the command removes the configured ve-id. It can be used just when the BGP VPWS status is shutdown. The **no shutdown** command cannot be used if there is no ve-id configured.

Default no ve-id

Parameters *value* — A two bytes identifier that represents the local or remote VPWS instance and is advertised through the BGP NLRI.

Values 1 to 65535

ve-name

Syntax [**no**] **ve-name** *name*

Context config>service>epipe>bgp-vpws

Description This command configures the name of the local VPWS instance in this service.

The **no** form of the command removes the ve-name.

Parameters *name* — Specifies a site name up to 32 characters in length.

shutdown

Syntax [**no**] **shutdown**

Context config>service>epipe>bgp-vpws

Description This command administratively enables/disables the local BGP VPWS instance. On de-activation an MP-UNREACH-NLRI is sent for the local NLRI.

The **no** form of the command enables the BGP VPWS addressing and the related BGP advertisement. The associated BGP VPWS MP-REACH-NLRI will be advertised in an update message and the corresponding received NLRIs must be considered to instantiate the data plane.

Default shutdown

site

Syntax **site** *name* [**create**]
no site name

Context config>service>epipe

Description This command configures a Epipe site.

The **no** form of the command removes the name from the configuration.

Parameters *name* — Specifies a site name up to 32 characters in length.

create — This keyword is mandatory while creating a Epipe service.

boot-timer

Syntax **boot-timer** *seconds*
no boot-timer

Context config>service>epipe>site

Description This command configures for how long the service manger waits after a node reboot before running the DF election algorithm. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged.

The **no** form of the command reverts the default.

Default 10

Parameters *seconds* — Specifies the site boot-timer in seconds.

Values 0 to 600

sap

Syntax **sap** *sap-id*
no sap

Context config>service>epipe>site

Description This command configures a SAP for the site.

The **no** form of the command removes the SAP ID from the configuration.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

site-activation-timer

Syntax **site-activation-timer** *seconds*
no site-activation-timer

Context config>service>epipe>site

Description This command configures the time-period the system keeps the local sites in standby status, waiting for BGP updates from remote PEs before running the DF (designated-forwarder) election algorithm to decide whether the site should be unblocked. This timer is terminated if an update is received for which the remote PE has transitioned from DF to non-DF.

The **no** form of the command removes the value from the configuration.

Default 2

Parameters *seconds* — Specifies the site activation timer in seconds.

Values 0 to 100

site-min-down-timer

Syntax **site-min-down-timer** *min-down-time*
no site-min-down-timer

Context config>service>epipe>site

Description This command configures the BGP multi-homing site minimum down time. When set to a non-zero value, if the site goes operationally down it will remain operationally down for at least the length of time configured for the **site-min-down-timer**, regardless of whether other state changes would have caused it to go operationally up. This timer is restarted every time that the site transitions from up to down. Setting this parameter to zero allows the minimum down timer to be disabled for this service.

The preceding operation is optimized in the following circumstances:

- If the site goes down on the designated forwarder but there are no BGP multi-homing peers with the same site in an operationally up state, then the **site-min-down-timer** is not started and is not used.
- If the site goes down on the designated forwarder but there are no active BGP multi-homing peers, then the s**ite-min-down-timer** is not started and is not used.
- If the **site-min-down-timer** is active and a BGP multi-homing update is received from the designated forwarder indicating its site has gone down, the **site-min-down-timer** is immediately terminated and this PE becomes the designated forwarder if the BGP multi-homing algorithm determines it should be the designated forwarder.

The **no** form of the command reverts to default value.

Default Taken from the value of **site-min-down-timer** configured for Multi-Chassis BGP multi-homing under the **config>redundancy>bgp-multi-homing** context.

Parameters *min-down-time* — Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down.

Values 0 to 100

site-id

Syntax **site-id** *value*
no site-id

Context config>service>epipe>site

Description This command configures the identifier for the site in this service. It must match between services but it is local to the service.

Parameters *value* — Specifies the site identifier.

Values 1 to 65535

site-preference

Syntax **site-preference** *preference-value*
no site-preference

Context config>service>epipe>site

Description This command defines the value to advertise in the VPLS preference field of the BGP VPWS and BGP Multi-homing NLRI extended community. This value can be changed without having to shutdown the site itself. The site-preference is only applicable to VPWS services.

When not configured, the default is zero, indicating that the VPLS preference is not in use.

Default no site-preference, value=0

Parameters *preference-value* — Specifies the preference value to advertise in the NLRI L2 extended community for this site.

Values 1 to 65535

primary — Sets the site-preference to 65535.

backup — Sets the site-preference to 1.

ce-address-discovery

Syntax [**no**] **ce-address-discovery** [**ipv6**] [**keep**]

Context config>service>ipipe

Description This command specifies whether the service will automatically discover the CE IP addresses.

When enabled, the addresses will be automatically discovered on SAPs that support address discovery, and on the spoke-SDPs. When enabled, addresses configuration on the Ipipe SAP and spoke-SDPs will not be allowed.

If disabled, CE IP addresses must be manually configured for the SAPs to become operationally up.

Default no ce-address-discovery

Parameters **ipv6** — The **ipv6** keyword enables IPv6 CE address discovery support on the Ipipe so that both IPv4 and IPv6 address discovery are supported. If the **ipv6** keyword is not included, then only IPv4 address discovery is supported and IPv6 packets are dropped. For the 7450 ESS platforms, it requires mixed mode support.

keep — The keep keyword is only applicable to eth-legacy-fault-notification. This option maintains the CE address discovered even when the SAP on which the address was learned fails. The ARP entry will not be maintained if the SAP is administratively shutdown, the clear service id x {arp | neighbor} is used to remove the ARP entry or the node reboots.

stack-capability-signaling

Syntax [**no**] **stack-capability-signaling**

Context config>service>ipipe

Description This command enables stack-capability signaling in the initial label mapping message of the Ipipe PW to indicate that IPv6 is supported.

When enabled, the 7750 SR includes the stack-capability TLV with the IPv6 stack bit set according to the **ce-address-discovery ipv6** keyword, and also checks the value of the stack-capability TLV received from the far end.

This command must be blocked if no **ce-address-discovery** is specified, or the **ipv6** keyword is not included with the **ce-address-discovery** command.

This command if only applicable to the Ipipe service and must be blocked for all other services.

This command has no effect if both SAPs on the Ipipe service are local to the node.

For the 7450 ESS platforms, it requires mixed mode support.

Default no stack-capability-signaling

endpoint

Syntax [**no**] **endpoint** *endpoint-name*

Context config>service>apipe
config>service>cpipe
config>service>fpipe
config>service>epipe
config>service>ipipe

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name.

load-balancing

Syntax load-balancing

Context config>service>epipe

Description This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Default not applicable

per-service-hashing

Syntax [**no**] **per-service-hashing**

Context config>service>epipe>load-balancing

Description This command enables on a per service basis, consistent per-service hashing for Ethernet services over LAG, over Ethernet tunnel (eth-tunnel) using loadsharing protection-type or over CCAG. Specifically, it enables the new hashing procedures for Epipe, VPLS, regular or PBB services.

The following algorithm describes the hash-key used for hashing when the new option is enabled:

- If the packet is PBB encapsulated (contains an I-TAG ethertype) at the ingress side, use the ISID value from the I-TAG
- If the packet is not PBB encapsulated at the ingress side
  - For regular (non-PBB) VPLS and Epipe services, use the related service ID
  - If the packet is originated from an ingress IVPLS or PBB Epipe SAP
  - If there is an ISID configured use the related ISID value
  - If there is no ISID yet configured use the related service ID
  - For BVPLS transit traffic use the related flood list id
  - Transit traffic is the traffic going between BVPLS endpoints
  - An example of non-PBB transit traffic in BVPLS is the OAM traffic
  - The preceding rules apply regardless of traffic type
  - Unicast, BUM flooded without MMRP or with MMRP, IGMP snooped

The **no** form of this command implies the use of existing hashing options.

Default no per-service-hashing

tunnel

Syntax **tunnel** *service-id* **backbone-dest-mac** *ieee-address* **isid** *ISID*
no tunnel

Context config>service>epipe>pbb

Description This command configures a Provider Backbone Bridging (PBB) tunnel with Backbone VPLS (B-VPLS) service information.

Parameters *service-id* — Specifies the B-VPLS service for the PBB tunnel associated with this service.

Values *service-id*: 1 to 2147483648
svc-name: 64 characters maximum

backbone-dest-mac *ieee-address* — Specifies the backbone destination MAC-address for PBB packets.

isid *ISID* — Specifies a 24 bit service instance identifier for the PBB tunnel associated with this service. As part of the PBB frames, it is used at the destination PE as a demultiplexor field.

Values 0 to 16777215

active-hold-delay

Syntax **active-hold-delay** *active-hold-delay*
no active-hold-delay

Context config>service>cpipe>endpoint
config>service>apipe>endpoint
config>service>fpipe>endpoint
config>service>ipipe>endpoint
config>service>epipe>endpoint

Description This command specifies that the node will delay sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby** or when any object in the endpoint. For example, SAP, ICB, or regular spoke-SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of “standby” over the spoke-SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from **standby** to **active** or when any object in the endpoint transitions to an operationally up state.

Default 0 — A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby,** the node sends immediately new T-LDP status bits indicating the new value of **standby** over the spoke-SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

Parameters *active-hold-delay* — Specifies the active hold delay in 100s of milliseconds.

Values 0 to 60

revert-time

Syntax **revert-time** [*revert-time* | **infinite**]
no revert-time

Context config>service>apipe>endpoint

config>service>fpipe>endpoint
config>service>cpipe>endpoint
config>service>epipe>endpoint
config>service>ipipe>endpoint

Description This command configures the time to wait before reverting back to the primary spoke-SDP defined on this service endpoint, after having failed over to a backup spoke-SDP.

Parameters *revert-time* — Specifies the time, in seconds, to wait before reverting to the primary SDP.

Values 0 to 600

Default 0

infinite — Causes the endpoint to be non-revertive.

eth-legacy-fault-notification

Syntax **eth-legacy-fault-notification**

Context config>service>ipipe

Description This is the top level of the hierarchy containing Ethernet to Legacy fault notification parameters. This context must activate using the no shutdown command before Ethernet to legacy fault notification can occur for iPipe services that make use of PPP, MLPPP or HDLC. This is only applicable to iPipe services with one legacy (PPP, MLPPP or HDLC) connection and an Ethernet SAP. No other services, not other combinations are supported.

recovery-timer

Syntax **recovery-timer** *timer-value*
no recovery-timer

Context config>service>ipipe>eth-legacy-fault-notification

Description This timer provides the legacy protocols PPP, MLPPP and HDLC time to establish after the Ethernet fault condition has cleared. The legacy protocol is afforded this amount of time to establish the connection before a fault is declared on the legacy side and propagated to the Ethernet segment. This timer is started as a result of a clearing Ethernet failure. Faults that may exist on the legacy side will not be detected until the expiration of this timer. Until the legacy side connection is established or the timer expires the traffic arriving on the Ethernet SAP from a peer will be discarded. The default value is unlikely to be a representative of all operator requirements and must be evaluated on a case by case basis.

Parameters *timer-value* — The value of the wait time in tenths of a second (100 ms). Granularity is in 500 ms increments, starting from 1 s and up to 30 s.

Values 10 to 300

Default 100

shutdown

Syntax [**no**] **shutdown**

Context config>service>ipipe>eth-legacy-fault-notification

Description This command enables or disables the propagation of fault from the Ethernet segment to the legacy connection using PPP, MLPPP and HDLC for an iPipe service. Issuing a “no shutdown” will activate the feature. Issuing a “shutdown” will deactivate the feature and stop fault notification from the Ethernet to PPP, MLPPP and HDLC protocols.

The **no** form of the command activates the ethernet legacy fault propagation.

Default shutdown

standby-signaling-master

Syntax [**no**] **standby-signaling-master**

Context config>service>vll>endpoint

Description When this command is enabled, the pseudowire standby bit (value 0x00000020) will be sent to T-LDP peer for each spoke-SDP of the endpoint that is selected as a standby.

This command is mutually exclusive with a VLL mate SAP created on a mc-lag/mc-aps or ICB. It is also mutually exclusive with vc-switching.

Default standby-signaling-master

standby-signaling-slave

Syntax [**no**] **standby-signaling-slave**

Context config>service>epipe>endpoint
config>service>epipe>spoke-sdp

Description When this command is enabled, the node will block the transmit forwarding direction of a spoke-SDP based on the pseudowire standby bit received from a T-LDP peer.

This command is present at the endpoint level as well as the spoke-SDP level. If the spoke-SDP is part of an explicit-endpoint, it will not be possible to change this setting at the spoke-SDP level. An existing spoke-SDP can be made part of the explicit endpoint only if the settings do not conflict. A newly created spoke-SDP, which is part of a specific explicit-endpoint, will inherit this setting from the endpoint configuration.

This command is mutually exclusive with an endpoint that is part of an mc-lag, mc-aps or an ICB.

If the command is disabled, the node assumes the existing independent mode of behavior for the forwarding on the spoke-SDP.

Default disabled

interworking

Syntax **interworking frf-5**
no interworking

Context config>service>apipe

Description This command specifies the interworking function that should be applied for packets that ingress/egress SAPs that are part of an Apipe service.

Interworking is applicable only when the two endpoints (i.e., the two SAPs or the SAP and the spoke-SDP) are of different types. Also, there are limitations on the combinations of SAP type, vc-type, and interworking values as shown in Table 12.

Table 12 **SAP types, VC-types and Interworking Values**

SAP Type	Allowed VC-Type Value	Allowed Interworking Value
ATM VC	atm-vcc, atm-sdu	None
	fr-dlci	Not Supported
FR DLCI	fr-dlci	None
	atm-sdu	frf-5

Default none (Interworking must be configured before adding a Frame-Relay SAP to an Apipe service.)

Parameters **frf-5** — Specifies Frame Relay to ATM Network Interworking (FRF.5).

service-name

Syntax **service-name** *service-name*
no service-name

Context config>service>apipe
config>service>cpipe
config>service>fpipe
config>service>ipipe
config>service>epipe

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a specific service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the SR OS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a specific service when it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0 to 9).

service-mtu

Syntax **service-mtu** *octets*
no service-mtu

Context config>service>epipe
config>service>ipipe
config>service>apipe
config>service>cpipe
config>service>fpipe

Description This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU. The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

Binding operational states are automatically re-evaluated.

For i-VPLS and Epipes bound to a b-VPLS, the service-mtu must be at least 18 bytes smaller than the b-VPLS service MTU to accommodate the PBB header.

Because this connects a Layer 2 to a Layer 3 service, adjust either the service-mtu under the Epipe service. The MTU that is advertised from the Epipe side is service-mtu minus EtherHeaderSize.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

By default if no service-mtu is configured it is (1514 - 14) = 1500.

Default apipe, fpipe: 1508

ipipe: 1500

epipe: 1514

Table 13 shows MTU values for specific VC types.

Table 13 **MTU Values**

SAP VC-Type	Example: Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (Q-in-Q with preserved bottom Qtag)	1518	1504

Parameters *octets* — The size of the MTU in octets, expressed as a decimal integer.

Values 1 to 9194

pw-port

Syntax **pw-port** *pw-port-id* **fpe** *fpe-id* [**create**]
no pw-port

Context config>service>epipe

Description This command is used to associate the PW-port with the PXC ports or PXC based LAGs referenced in the FPE. In other words, the PW-port becomes anchored by the PXC. This enables an external PW that is mapped to the anchored PW-port to be seamlessly rerouted between the I/O ports without interruption of service on the PW-port.

This mapping between the external PW (spoke-SDP) and the PXC based PXC-port is performed via an Epipe operating in vc-switching mode (creation time parameter).

Default no pw-port

Parameters *pw-port-id* — Specifies the PW-port associated with this service.

Values 1 to 10239

fpe *fpe-id* — Specifies the FPE object which contains the PXC-based ports or PXC-based LAGs.

Values 1 to 64

egress

Syntax **egress**

Context config>service>epipe>pw-port

Description This command enters the context to configure PW-port egress-side parameters

Default N/A

shaper

Syntax [**no**] **shaper**

Context config>service>epipe>pw-port>egress

Description This command enters the context to configure PW-port shaper parameters.

Default N/A

int-dest-id

Syntax **int-dest-id** *name*
no int-dest-id

Context config>service>epipe>pw-port>egress>shaper

Description This command configures an intermediate destination identifier applicable to ESM PW SAPs.

Default N/A

Parameters *name* — Specifies the default intermediate destination identifier, up to 32 characters in length, on the egress side for this PW-port entry.

vport

Syntax **vport** *vport*
no vport

Context config>service>epipe>pw-port>egress>shaper

Description This command configures specifies the virtual port name of the shaper on the egress side for this PW-port entry.

Default N/A

Parameters *vport* — Specifies a virtual port applicable to all PW SAPs.

monitor-oper-group

Syntax **monitor-oper-group** *group-name*
no monitor-oper-group

Context config>service>epipe>pw-port

Description This command configures the monitoring operational group name, up to 32 characters in length, associated with this PW-port entry.

Default N/A

Parameters *group-name* — Specifies an operational group to monitor.

signaled-vc-type-override

Syntax **signaled-vc-type-override** {**atm-vcc**}
no signaled-vc-type-override

Context <root>

Description This command overrides the pseudowire type signaled to type 0x0009 N:1 VCC cell within an Apipe VLL service of vc-type atm-cell. Normally, this service vc-type signals a pseudowire of type 0x0003 ATM Transparent Cell.

This command is not allowed in an Apipe VLL of vc-type value atm-cell if a configured ATM SAP is not using a connection profile. Conversely, if the signaling override command is enabled, only an ATM SAP with a connection profile assigned will be allowed.

The **override** command is not allowed on Apipe VLL service of vc-type value other than atm-cell. It is also not allowed on a VLL service with the vc-switching option enabled since signaling of the PW FEC in a Multi-Segment PW (MS-PW) is controlled by the T-PE nodes. Thus for this feature to be used on a MS-PW, it is required to configure an Apipe service of vc-type atm-cell at the T-PE nodes with the signaled-vc-type-override enabled, and to configure a Apipe VLL service of vc-type atm-vcc at the S-PE node with the vc-switching option enabled.

The **no** form of this command returns the Apipe VLL service to signal its default pseudowire type

Default none

Parameters **atm-vcc** — Specifies the pseudowire type to be signaled in the pseudowire establishment

connection-profile

Syntax **connection-profile** *conn-prof-id* [**create**]
no connection-profile *conn-prof-id*

Context <root>

Description This command creates a profile for the user to configure the list of discrete VPI/VCI values to be assigned to an ATM SAP of an Apipe VLL of **vc-type atm-cell.**

A connection profile can only be applied to a SAP which is part of an Apipe VLL service of **vc-type atm-cell**. The ATM SAP can be on a regular port or APS port.

A maximum of 8000 connection profiles can be created on the system.

The **no** form of this command deletes the profile from the configuration.

Default none

Parameters *conn-prof-id* — Specifies the profile number.

Values 1 to 8000

member

Syntax **member** *encap-value* [**create**]
no member *encap-value*

Context config>connection-profile

Description This command allows the adding of discrete VPI/VCI values to an ATM connection profile for assignment to an ATM SAP of an Apipe VLL of vc-type atm-cell.

Up to a maximum of 16 discrete VPI/VCI values can be configured in a connection profile. The user can modify the content of a profile which triggers a re-evaluation of all the ATM SAPs which are currently using the profile.

The **no** form of this command deletes the member from the configuration.

Default none

Parameters *encap-value* — Specifies the VPI and VCI values of this connection profile member.

Values vpi: NNI: 0 to 4095
UNI: 0 to 255
vci: 1, 2, 5 to 65535

2.17.2.4 VLL SAP Commands

sap

Syntax **sap** *sap-id* [**create**] [**no-endpoint**]
sap *sap-id* [**create**] **endpoint** *endpoint-name*
no sap *sap-id*

Context config>service>apipe
config>service>cpipe
config>service>fpipe
config>service>ipipe
config>service>epipe

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the device. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port. Channelized TDM ports are always access ports.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded.

The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The following are supported on the 7750 SR only:

- ATM VPI/VCI on an ATM port for vc-type atm-vcc and atm-sdu
- ATM VPI on an ATM port for vc-type atm-vpc
- ATM virtual trunk - a range of VPIs on an ATM port for vc-type atm-cell
- ATM port for vc-type atm-cell
- ATM connection profile for vc-type atm-cell
- Frame Relay DLCI on a port for vc-type atm-sdu
- ATM SAP carries the IPv4 packet using RFC 2684, VC-Mux or LLC/SNAP routed PDU encapsulation for an Ipipe service
- Frame Relay SAP RFC 2427, routed PDU encapsulation for an Ipipe service
- Ethernet SAP RFC 1332, PPP IPCP encapsulation of an IPv4 packet for an Ipipe service
- Ethernet SAP HDLC SAP uses the routed IPv4 encapsulation for an Ipipe service
- ATM - Frame Relay, PPP/IPCP - PPP/IPCP
- Frame Relay-Frame Relay, ATM - ATM
- Ethernet-Ethernet
- cHDLC-cHDLC
- An ATM SAP can be part of an IMA bundle.
- A PPP SAP can be part of an MLPPP bundle.
- A FR SAP can be part of a MLFR bundle.

Ethernet SAPs support null, dot1q, and qinq is supported for all routers.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Enhanced Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Default No SAPs are defined

Special Cases **Special Cases —** A SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. At most, only one sdp-id can be bound to an VLL service. Since a VLL is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. Up to 49 SDPs can be associated with a service in a single router. Each SDP must have a unique router destination or an error will be generated.

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS). This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0).

Two Frame Relay SAPs cannot be configured on an Apipe service on the 7750 SR. The limitation is for an Apipe service in local mode, which has two SAPs associated with the service, as opposed to a configuration with a SAP and a SDP in remote case, the only combination of the type of SAPs allowed is either two ATM SAPs or an ATM SAP and a Frame Relay SAP. The CLI prevents adding two Frame Relay SAPs under an Apipe service.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP.

port-id — Specifies the physical port ID.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number format. For example 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

port-id	*slot*/*mda*/*port* [.*channel*]		
	eth-sat-id	esat-*id*/*slot*/*port*	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-*id*.*sub-port*	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

endpoint — Adds a SAP endpoint association.

no endpoint — Removes the association of a SAP or a spoke-SDP with an explicit endpoint name.

create — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Output

Sample Output

```
*A:bksim2801>config>service>apipe>sap$
===============================================================
ATM PVCs, Port 1/1/1
===============================================================
VPI/VCI     Owner     Type     Ing.TD   Egr.TD  Adm  OAM     Opr
---------------------------------------------------------------
2/102       SAP       PVC      1        1       up   ETE-AIS dn
10/100      SAP       PVC      1        1       up   ETE-AIS dn
===============================================================
*A:bksim2801#
```

sap

Syntax [**no**] **sap** *eth-tunnel-tunnel-id*[:*eth-tunnel-sap-id*] [**create**]

Context config>service>epipe
config>service>ipipe
config>service>vpls

Description This command configures an Ethernet tunnel SAP.

An Ethernet tunnel control SAP has the format eth-tunnel-*tunnel-id* and is not configured with an Ethernet tunnel SAP ID. No Ethernet tunnel tags can be configured under a control SAP since the control SAP uses the control tags configured under the Ethernet tunnel port. This means that at least one member port and control tag must be configured under the Ethernet tunnel port before this command is executed. The control SAP is needed for carrying G.8031 and 802.1ag protocol traffic. This SAP can also carry user data traffic.

An Ethernet tunnel same-fate SAP has the format eth-tunnel-*tunnel-id:eth-tunnel-sap-id*. Same-fate SAPs carry only user data traffic. Multiple same-fate SAPs can be configured on one Ethernet tunnel port and share the fate of that port, provided the SAPs are properly configured with corresponding tags.

Ethernet tunnel SAPs are supported under VPLS, Epipe and Ipipe services only.

Default no sap

Parameters *tunnel-id* — Specifies the tunnel ID.

Values 1 to 1024

eth-tunnel-sap-id — Specifies a SAP ID of a same-fate SAP.

Values 0 to 4094

aarp

Syntax **aarp** *aarpId* **type** *type*
no aarp

Context config>service>epipe>sap
config>service>epipe>spoke-sdp

Description This command associates an AARP instance with a multi-homed SAP or spoke-SDP. This instance uses the same AARP ID in the same node to provide traffic flow and packet asymmetry removal for a multi-homed SAP or spoke-SDP.

The type specifies the role of this service point in the AARP: either, primary (dual-homed) or secondary (dual-homed-secondary). The AA service attributes (app-profile and transit-policy) of the primary are inherited by the secondary endpoints. All endpoints within an AARP must be of the same type (SAP or spoke), and all endpoints with an AARP must be within the same service.

The no form of the command removes the association between an AARP instance and a multi-homed SAP or spoke-SDP.

Default no aarp

Parameters *aarpid* — Specifies the AARP instance associated with this SAP. If not configured, no AARP instance is associated with this SAP.

Values 1 to 65535

type — Specifies the role of the SAP referenced by the AARP instance.

Values **dual-homed** — The primary dual-homed AA subscriber side service-point of an AARP instance; only supported for Epipe, IES, and VPRN SAP and spoke-SDP.

dual-homed-secondary — One of the secondary dual-homed AA subscriber side service-points of an AARP instance; only supported for Epipe, IES, and VPRN SAP and spoke-SDP.

lag-link-map-profile

Syntax **lag-link-map-profile** *link-map-profile-id*
no lag-link-map-profile

Context config>service>epipe>sap
config>service>ipipe>sap

Description This command assigns a pre-configured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/de-assigned, the SAP's/network interface's egress traffic will be re-hashed over LAG as required by the new configuration.

The **no** form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.

Default no lag-link-map-profile

Parameters *link-map-profile-id* — An integer from 1 to 64 that defines a unique lag link map profile on the LAG the SAP/network interface exists on.

lag-per-link-hash

Syntax **lag-per-link-hash class** {**1** | **2** | **3**} **weight** [**1 to 1024**]
no lag-per-link-hash

Context config>service>epipe>sap
config>service>ipipe>sap
config>service>vpls>sap

Description This command configures weight and class to this SAP to be used on LAG egress when the LAG uses weighted per-link-hash.

The **no** form of this command restores default configuration.

Default no lag-per-link-hash (equivalent to weight 1 class 1)

agg-rate

Syntax [**no**] **agg-rate**

Context config>service>apipe>sap>egress
config>service>cpipe>sap>egress
config>service>epipe>sap>egress
config>service>fpipe>sap>egress
config>service>ipipe>sap>egress

Description This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

rate

Syntax **rate** *kilobits-per-second*
no rate

Context config>service>apipe>sap>egress>agg-rate
config>service>cpipe>sap>egress>agg-rate
config>service>epipe>sap>egress>agg-rate
config>service>fpipe>sap>egress>agg-rate
config>service>ipipe>sap>egress>agg-rate

Description This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, Vport, and so on).

Parameters *kilobits-per-second* — The enforced aggregate rate for all queues associated with the agg-rate context, in kilobits per second.

Values 1 to 3200000000 | max

limit-unused-bandwidth

Syntax [**no**] **limit-unused-bandwidth**

Context config>service>apipe>sap>egress>agg-rate
config>service>cpipe>sap>egress>agg-rate
config>service>epipe>sap>egress>agg-rate
config>service>fpipe>sap>egress>agg-rate
config>service>ipipe>sap>egress>agg-rate

Description This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

queue-frame-based-accounting

Syntax [**no**] **queue-frame-based-accounting**

Context config>service>apipe>sap>egress>agg-rate
config>service>cpipe>sap>egress>agg-rate
config>service>fpipe>sap>egress>agg-rate
config>service>ipipe>sap>egress>agg-rate

Description This command is used to enable (or disable) frame based accounting on all policers and queues associated with the agg-rate context.

The command is supported on Ethernet ports only; it is not supported on HSMDA Ethernet ports.

Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured; however the offsets are applied to the statistics.

policer-control-override

Syntax **policer-control-override** [**create**]
no policer-control-override

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress

Description This command, within the SAP ingress or egress contexts, creates a CLI node for specific overrides to the applied policer-control-policy. A policy must be applied for a policer-control-overrides node to be created. If the policer-control-policy is removed or changed, the policer-control-overrides node is automatically deleted from the SAP.

The **no** form of the command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.

Default no policer-control-override

Parameters **create** — The create keyword is required when the policer-control-overrides node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

max-rate

Syntax **max-rate** {*rate* | **max**}

Context config>service>apipe>sap>egress>policer-control-override
config>service>apipe>sap>ingress>policer-control-override
config>service>cpipe>sap>egress>policer-control-override
config>service>cpipe>sap>ingress>policer-control-override
config>service>epipe>sap>egress>policer-control-override
config>service>epipe>sap>ingress>policer-control-override
config>service>fpipe>sap>egress>policer-control-override
config>service>fpipe>sap>ingress>policer-control-override
config>service>ipipe>sap>egress>policer-control-override
config>service>ipipe>sap>ingress>policer-control-override

Description This command, within the SAP ingress and egress contexts, overrides the root arbiter parent policer max-rate that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy max-rate parameter have no effect on the SAP's parent policer until the override is removed using the **no max-rate** command within the SAP.

Parameters *rate* — Specifies the rate override in kilobits per second.

Values 1 to 2000000000

max — Specifies the maximum rate override.

priority-mbs-thresholds

Syntax **priority-mbs-thresholds**

Context config>service>apipe>sap>egress>policer-control-override
config>service>apipe>sap>ingress>policer-control-override
config>service>cpipe>sap>egress>policer-control-override
config>service>cpipe>sap>ingress>policer-control-override
config>service>epipe>sap>egress>policer-control-override
config>service>epipe>sap>ingress>policer-control-override
config>service>fpipe>sap>egress>policer-control-override
config>service>fpipe>sap>ingress>policer-control-override
config>service>ipipe>sap>egress>policer-control-override
config>service>ipipe>sap>ingress>policer-control-override

Description This command overrides the CLI node contains the configured min-thresh-separation and the various priority level mbs-contribution override commands.

min-thresh-separation

Syntax **min-thresh-separation** *size* [**bytes** | **kilobytes**]

Context config>service>apipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>apipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>cpipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>cpipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>epipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>epipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>fpipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>fpipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>ipipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>ipipe>sap>ingress>policer-control-override>priority-mbs-threshold

Description This command, within the SAP ingress and egress contexts, is used to override the root arbiter's parent policer min-thresh-separation parameter that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy min-thresh-separation parameter have no effect on the SAP's parent policer until the override is removed using the no min-thresh-separation command within the SAP.

The no form of the command removes the override and allows the min-thresh-separation setting from the policer-control-policy to control the root arbiter's parent policer's minimum discard threshold separation size.

Default no min-thresh-separation

Parameters *size* — The minimum discard threshold separation override value.

Values 1 to 16777216 | default

bytes — Signifies that *size* is expressed in bytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

kilobytes — Signifies that *size* is expressed in kilobytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

priority

Syntax [**no**] **priority** *level*

Context config>service>apipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>apipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>cpipe>sap>egress>policer-control-override>priority-mbs-threshold

config>service>cpipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>epipe>sap>egress>policer-control-override>priority-mbs-thresholds
config>service>epipe>sap>ingress>policer-control-override>priority-mbs-thresholds
config>service>fpipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>fpipe>sap>ingress>policer-control-override>priority-mbs-threshold
config>service>ipipe>sap>egress>policer-control-override>priority-mbs-threshold
config>service>ipipe>sap>ingress>policer-control-override>priority-mbs-threshold

Description The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.

This node does not need to be created and will not be output in show or save configurations unless an mbs-contribution override exist for level.

Parameters *level* — The level parameter is required when specifying priority-level and identifies which of the parent policer instances priority level's the mbs-contribution is overriding.

Values 1 to 8

mbs-contribution

Syntax **mbs-contribution** *size* [**bytes** | **kilobytes**]

Context config>service>apipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>apipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority
config>service>cpipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>cpipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority
config>service>epipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>epipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority
config>service>fpipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>fpipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority
config>service>ipipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
config>service>ipipe>sap>ingress>policer-control-override>priority-mbs-threshold>priority

Description The mbs-contribution override command within the SAP ingress and egress contexts is used to override a parent policer's priority level's mbs-contribution parameter that is defined within the policer-control-policy applied to the SAP. This override allow the priority level's burst tolerance to be tuned based on the needs of the SAP's child policers attached to the priority level.

When the override is defined, modifications to the policer-control-policy priority level's mbs-contribution parameter have no effect on the SAP's parent policer priority level until the override is removed using the no mbs-contribution command within the SAP.

The **no** form of the command removes the override and allows the mbs-contribution setting from the policer-control-policy to control the parent policer's priority level's burst tolerance.

Default no mbs-contribution

Parameters *size* — The mbs-contribution override value.

Values 1 to 16777216 | default

bytes — Signifies that *size* is expressed in bytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

kilobytes — Signifies that *size* is expressed in kilobytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

policer-control-policy

Syntax **policer-control-policy** *policy-name* [**create**]
no policer-control-policy

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>epipe>sap>egress

Description This command, within the QoS CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied.

When applied to a sub-profile on a 7450 ESS and 7750 SR, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis.

For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and therefore the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As previously stated, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair-based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated in the Tier 1 and Tier 2 Arbiter subsection, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

Each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

To derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless.

When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of the command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP context.

Default none

Parameters *policy-name* — Each policer-control-policy must be created with a unique policy name. The name given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create — The keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

policer-override

Syntax [**no**] **policer-override**

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress

config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress

Description This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of the command is used to remove any existing policer overrides.

Default no policer-overrides

policer

Syntax **policer** *policer-id* [**create**]
no policer *policer-id*

Context config>service>apipe>sap>egress>policer-override
config>service>apipe>sap>ingress>policer-override
config>service>cpipe>sap>egress>policer-override
config>service>cpipe>sap>ingress>policer-override
config>service>fpipe>sap>egress>policer-override
config>service>fpipe>sap>ingress>policer-override
config>service>ipipe>sap>egress>policer-override
config>service>ipipe>sap>ingress>policer-override
config>service>epipe>sap>egress>policer-override

Description This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy.

The **no** form of the command is used to remove any existing overrides for the specified policer-id.

Parameters *policer-id* — The *policer-id* parameter is required when executing the policer command within the policer-overrides context. The specified *policer-id* must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the policer-id.

create — The create keyword is required when a policer policer-id override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

cbs

Syntax **cbs** *size* [{**bytes** | **kilobytes**}]

no cbs

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>cpipe>sap>egress>policer-override>policer
config>service>cpipe>sap>ingress>policer-override>policer
config>service>ipipe>sap>egress>policer-override>policer
config>service>ipipe>sap>ingress>policer-override>policer
config>service>epipe>sap>egress>policer-override>policer

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured CBS parameter for the specified *policer-id*.

The **no** form of this command returns the CBS size to the default value.

Default no cbs

Parameters *size* — The size parameter is required when specifying cbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 | default

bytes — When bytes is defined, the value given for *size* is interpreted as the policer's MBS value in bytes.

kilobytes — When kilobytes is defined, the value given for *size* is interpreted as the policer's MBS value in kilobytes.

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>cpipe>sap>egress>policer-override>policer
config>service>cpipe>sap>ingress>policer-override>policer
config>service>ipipe>sap>egress>policer-override>policer
config>service>ipipe>sap>ingress>policer-override>policer
config>service>epipe>sap>egress>policer-override>policer
config>service>epipe>sap>ingress>policer-override>policer

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured mbs parameter for the specified policer-id.

The **no** form of the command is used to restore the policer's mbs setting to the policy defined value.

Default no mbs

Parameters *size* — The size parameter is required when specifying mbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 | default

bytes — When **bytes** is defined, the value given for *size* is interpreted as the policer's MBS value in bytes.

kilobytes — When **kilobytes** is defined, the value given for *size* is interpreted as the policer's MBS value in kilobytes.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>cpipe>sap>egress>policer-override>policer
config>service>cpipe>sap>ingress>policer-override>policer
config>service>ipipe>sap>egress>policer-override>policer
config>service>ipipe>sap>ingress>policer-override>policer
config>service>epipe>sap>egress>policer-override>policer

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured packet-byte-offset parameter for the specified policer-id. Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured; however, the offsets are applied to the statistics.

The **no** packet-byte-offset command is used to restore the policer's packet-byte-offset setting to the policy defined value.

Default no packet-byte-offset

Parameters *add-bytes* — Specifies the number of bytes that are added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 1 to 31

sub-bytes — Specifies the number of bytes that are subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 to 64

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]
no percent-rate

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>cpipe>sap>egress>policer-override>policer
config>service>cpipe>sap>ingress>policer-override>policer
config>service>ipipe>sap>egress>policer-override>policer
config>service>ipipe>sap>ingress>policer-override>policer
config>service>epipe>sap>egress>policer-override>policer

Description This command configures the percent rates (CIR and PIR) override.

Parameters *pir-percent* — Specifies the policer's PIR as a percentage of the policers's parent arbiter rate.

Values 0.01 to 100.00

Default 100.00

cir-percent — Specifies the policer's CIR as a percentage of the policers's parent arbiter rate.

Values 0.00 to 100.00

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*] [
no percent-rate

Context config>service>epipe>sap>egress>queue-override>queue

Description The percent-rate command within the SAP ingress and egress QoS policy enables supports for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate or of its parent scheduler's rate.

When the rates are expressed as a port-limit, the actual rates used per instance of the queue will vary based on the port speed. For example, when the same QoS policy is used on a 1-Gb and a 10-Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb port due to the difference in port speeds. This enables the same QOS policy to be used on SAPs on different ports without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

When the rates are expressed as a local-limit, the actual rates used per instance of the queue are relative to the queue's parent scheduler rate. This enables the same QOS policy to be used on SAPs with different parent scheduler rates without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the parent scheduler rate changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

Queue rate overrides can only be specified in the form as configured in the QoS policy (a SAP override can only be specified as a percent-rate if the associated QoS policy was also defined as percent-rate). Likewise, a SAP override can only be specified as a rate (kb/s) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QOS policy.

When no percent-rate is defined within a SAP ingress or egress queue-override, the queue reverts to the defined shaping and CIR rates within the SAP ingress and egress QOS policy associated with the queue.

Parameters *percent-of-line-rate* — The percent-of-line-rate parameter is used to express the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

pir-percent — Specifies the queue's PIR as a percentage dependent on the use of the port-limit or local-limit.

Values 0.01 to 100.00

Default 100.00

cir-percent — Specifies the queue's CIR as a percentage dependent on the use of the port-limit or local-limit.

Values 0.00 to 100.00

Default 100.00

rate

Syntax **rate** {*rate* | **max**} [**cir** {*rate* | **max**}]

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>cpipe>sap>egress>policer-override>policer
config>service>cpipe>sap>ingress>policer-override>policer
config>service>epipe>sap>egress>policer-override>policer
config>service>epipe>sap>ingress>policer-override>policer
config>service>ipipe>sap>egress>policer-override>policer
config>service>ipipe>sap>ingress>policer-override>policer

Description This command within the SAP ingress and egress policer-overrides contexts is used to override the sap-ingress and sap-egress QoS policy configured rate parameters for the specified policer-id.

The **no** rate command is used to restore the policy defined metering and profiling rate to a policer.

Parameters **rate** *rate* — Specifies the policer instance metering rate for the PIR leaky bucket, in kilobits per second. The integer value is multiplied by 1000 to derive the actual rate in bits per second.

Values 1 to 2000000000

cir *rate* — Specifies the overriding value for the policy-derived profiling rate of the policer, in kilobits per second. The integer value is multiplied by 1000 to derive the actual rate in bits per second.

Values 0 to 2000000000

max — Uses the maximum policer rate, equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR or CIR used is equivalent to **max**.

stat-mode

Syntax **stat-mode** st*at-mode*
no stat-mode

Context config>service>apipe>sap>egress>policer-override>policer
config>service>apipe>sap>ingress>policer-override>policer
config>service>cpipe>sap>egress>policer-override>policer
config>service>cpipe>sap>ingress>policer-override>policer
config>service>epipe>sap>egress>policer-override>policer
config>service>epipe>sap>ingress>policer-override>policer
config>service>fpipe>sap>egress>policer-override>policer
config>service>fpipe>sap>ingress>policer-override>policer
config>service>ipipe>sap>egress>policer-override>policer
config>service>ipipe>sap>ingress>policer-override>policer

Description The SAP QoS policy's **policer stat-mode** command is used to configure the forwarding plane counters that allow offered, output, and discard accounting to occur for the policer. A policer has multiple types of offered packets (for example, soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overrides) and each of these offered types is interacting with the policers metering and profiling functions resulting in colored output packets (green, yellow, and red). Due to the potentially large number of egress policers, it is not economical to allocate counters in the forwarding plane for all

possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and indicates how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported that prevents any packet accounting, the use of the policer's parent command requires that the policer's **stat-mode** to be set at least to the minimal setting so that offered statistics are available for the policer's Fair Information Rate (FIR) to be calculated.

Each time the policer's stat mode is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources.The total/allocated/free statistics can be viewed by using the **tools dump resource-usage card fp** command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on a SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The current active stat mode setting will continue to be used by the policer.

The **no stat-mode** command attempts to return the policer's stat-mode setting to **minimal**. The command will fail if insufficient policer counter resources exist to implement **minimal** where the QoS policer is currently applied and has a forwarding class mapping.

Refer to the *7450 ESS, 7750 SR, and 7950 XRS Quality of Service Guide* for detailed information about the supported parameters for the **policer stat-mode** command.

ce-address

Syntax **ce-address** *ip-address*
no ce-address

Context config>service>ipipe>sap
config>service>ipipe>spoke-sdp

Description This command specifies the IP address of the CE device associated with an Ipipe SAP or spoke-SDP. In the case of a SAP, it is the address of the CE device directly attached to the SAP. For a spoke-SDP, it is the address of the CE device reachable through that spoke-SDP (for example, attached to the SAP on the remote node). The address must be a host address (no subnet addresses are accepted) as there must be only one CE device attached to an Ipipe SAP. The CE address specified at one end of an Ipipe will be used in processing ARP messages at the other endpoint, as the router acts as a proxy for ARP messages.

On a 7450 ESS, this command specifies the IP address of the CE device associated with an Ipipe SAP. In the case of a SAP, it is the address of the CE device directly attached to the SAP. The address must be a host address (no subnet addresses are accepted) as there must be only one CE device attached to an Ipipe SAP. The CE address specified at one end of an Ipipe will be used in processing ARP messages at the other endpoint, as the router acts as a proxy for ARP messages.

Parameters *ip-address* — Specifies the IP address of the CE device associated with an Ipipe SAP.

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>cpipe>sap>egress
config>service>apipe>sap>egress
config>service>epipe>sap>egress
config>service>fpipe>sap>egress
config>service>apipe>sap>egress

Description When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the **qinq-mark-top-only** command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When the enabled, only the P-bits/DEI bit in the top Q-tag are marked.

Default no qinq-mark-top-only

multi-service-site

Syntax **multi-service-site** *customer-site-name*
no multi-service-site

Context config>service>ipipe>sap
config>service>apipe>sap
config>service>cpipe>sap
config>service>fpipe>sap
config>service>epipe>sap

Description This command associates the SAP with a *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command will not execute. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within *customer-site-name* as parent schedulers.

The **no** form of the command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future queues to enter an orphaned state.

Default None

Parameters *customer-site-name* — The customer-site-name must exist in the context of the customer-id defined as the service owner. If customer-site-name exists and local scheduler policies have not been applied to the SAP, the current and future queues defined on the SAP will look for their parent schedulers within the scheduler hierarchies defined on customer-site-name.

Values Any valid customer-site-name created within the context of the customer-id.

ring-node

Syntax **ring-node** *ring-node-name*
no ring-node

Context config>service>epipe>sap

Description This command configures a multi-chassis ring-node for this SAP.

The **no** form of the command removes the name from the configuration.

Default none

transit-policy

Syntax **transit-policy** {**ip** *ip-aasub-policy-id* | prefix *prefix-aasub-policy-id*}
no transit-policy

Context config>service>epipe>sap
config>service>epipe>spoke-sdp

Description This command associates an AA transit policy to the service. The transit IP policy must be defined prior to associating the policy with a SAP in the **config**>**application assurance**>**group**>**policy**>**transit-ip-policy** context.

Transit AA subscribers are managed by the system through this service policy, which determines how transit subs are created and removed for that service.

The no form of the command removes the association of the policy to the service.

Default no transit-policy

Parameters *ip-aasub-policy-id* — Specifies an integer identifying an IP transit IP profile entry.

Values 1 to 65535

prefix-aasub-policy-id — Specifies an integer identifying a prefix transit profile entry.

Values 1 to 65535

use-broadcast-mac

Syntax [**no**] **use-broadcast-mac**

Context config>service>ipipe>sap

Description This command enables the user of a of broadcast MAC on SAP.

An Ipipe VLL service with the ce-address-discovery command enabled forwards unicast IP packets using the broadcast MAC address until the ARP cache is populated with a valid entry for the CE IP and MAC addresses.

The **no** form of this command enables the user of a of broadcast MAC on SAP.

Default no use-broadcast-mac

mac

Syntax [**no**] **mac** *ieee-address*

Context config>service>ipipe>sap

Description This command assigns a specific MAC address to an Ipipe SAP.

The **no** form of this command returns the MAC address of the SAP to the default value.

Default The physical MAC address associated with the Ethernet interface where the SAP is configured.

Parameters *ieee-address* — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

mac-refresh

Syntax **mac-refresh** *refresh interval*
no mac-refresh

Context config>service>ipipe>sap

Description This command specifies the interval between ARP requests sent on this Ipipe SAP. When the SAP is first enabled, an ARP request will be sent to the attached CE device and the received MAC address will be used in addressing unicast traffic to the CE. Although this MAC address will not expire while the Ipipe SAP is enabled and operational, it is verified by sending periodic ARP requests at the specified interval.

The **no** form of this command restores mac-refresh to the default value.

Default 14400

Parameters *refresh interval* — Specifies the interval, in seconds, between ARP requests sent on this Ipipe SAP.

Values 0 to 65535

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>epipe>sap
config>service>fpipe>sap
config>service>ipipe
config>service>epipe>spoke-sdp

Description This command creates the accounting policy context that can be applied to a SAP.

An accounting policy must be defined before it can be associated with a SAP. If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SAP, and the accounting policy reverts to the default.

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command configures the application profile name.

Parameters *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

bandwidth

Syntax **bandwidth** *bandwidth*
no bandwidth

Context config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>cpipe>spoke-sdp

Description This command specifies the bandwidth to be used for VLL bandwidth accounting by the VLL CAC feature.

The service manager keeps track of the available bandwidth for each SDP. The maximum value is the sum of the bandwidths of all constituent LSPs in the SDP. The SDP available bandwidth is adjusted by the user configured booking factor.

If an LSP consists of a primary and many secondary standby LSPs, then the bandwidth used in the maximum SDP available bandwidth is that of the active path. Any change to and LSP active path bandwidth will update the maximum SDP available bandwidth. Note however that a change to any constituent LSP bandwidth due to re-signaling of the primary LSP path or the activation of a secondary path which causes overbooking of the maximum SDP available bandwidth causes a warning and a trap to be issued but no further action is taken. The activation of a bypass or detour LSP in the path of the primary LSP does not change the maximum SDP available bandwidth.

When the user binds a VLL service to this SDP, an amount of bandwidth equal to bandwidth is subtracted from the SDP available bandwidth adjusted by the booking factor. When the user deletes this VLL service binding from this SDP, an amount of bandwidth equal to bandwidth is added back into the SDP available bandwidth.

If the total SDP available bandwidth when adding this VLL service is about to overbook, a warning is issued and the binding is rejected. This means that the spoke-SDP bandwidth does not update the maximum SDP available bandwidth. In this case, the spoke-SDP is put in operational down state and a status message of “pseudowire not forwarding” is sent to the remote SR-series PE node. A trap is also generated. The service manager will not put the spoke-SDP into an operationally up state until the user executes a **shutdown** command and then a **no-shutdown** command of the spoke-SDP and the bandwidth check succeeds. Thus, the service manager will not automatically audit spoke-SDPs subsequently to their creation to check if bandwidth is available.

If the VLL service contains an endpoint with multiple redundant spoke-SDPs, each spoke-SDP will have its bandwidth checked against the available bandwidth of the corresponding SDP.

If the VLL service performs a pseudowire switching (VC switching) function, each spoke-SDP is separately checked for bandwidth against the corresponding SDP.

This feature does not alter the way service packets are sprayed over multiple RSVP LSPs, which are part of the same SDP. In other words, by default load balancing of service packets occurs over the SDP LSPs based on service-id, or based on a hash of the packet header if ingress SAP shared queuing is enabled. In both cases, the VLL bandwidth is not checked against the available bandwidth of the selected LSPs but on the total SDP available bandwidth. Thus, if there is a single LSP per SDP, these two match.

If class-forwarding is enabled on the SDP, VLL service packets are forwarded to the SDP LSP which the packet forwarding class maps to, or if this is down to the default LSP. However, the VLL bandwidth is not checked against the selected LSP available bandwidth but on the total SDP available bandwidth. If there is a single LSP per SDP, these two match.

If a non-zero bandwidth is specified for a VLL service and attempts to bind the service to an LDP or a GRE SDP, a warning is issued that CAC failed but the VLL is established. A trap is also generated.

The **no** form of the command reverts to the default value.

Parameters *bandwidth* — The bandwidth to be used for VLL bandwidth accounting by the VLL CAC feature, in kilobits per second.

Values 0 to 100000000

Default 0

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>service>epipe>spoke-sdp
config>service>epipe>bgp>pw-template-binding
config>service>fpipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>cpipe>spoke-sdp

Description This command enables VCCV BFD on the PW associated with the VLL, BGP VPWS, or VPLS service. The parameters for the BFD session are derived from the named BFD template, which must have been first configured using the **bfd-template** command.

bfd-template

Syntax **bdf-template** *name*
no bfd-template

Context config>service>epipe>spoke-sdp
config>service>epipe>bgp>pw-template-binding
config>service>fpipe>spoke-sdp

config>service>apipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>cpipe>spoke-sdp

Description This command configures a named BFD template to be used by VCCV BFD on PWs belonging to the VLL, BGP VPWS, or VPLS service. The template specifies parameters, such as the minimum transmit and receive control packet timer intervals, to be used by the BFD session. Template parameters are configured under the **config>router>bfd** context.

Default no bfd-template

Parameters *name* — A text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

block-on-peer-fault

Syntax [**no**] **block-on-peer-fault**

Context config>service>epipe>spoke-sdp

Description When enabled, this command blocks the transmit direction of a PW when any of the following PW status codes is received from the far end PE:

0x00000001	Pseudowire Not Forwarding
0x00000002	Local Attachment Circuit (ingress) Receive Fault
0x00000004	Local Attachment Circuit (egress) Transmit Fault
0x00000008	Local PSN-facing PW (ingress) Receive Fault
0x00000010	Local PSN-facing PW (egress) Transmit Fault

The transmit direction is unblocked when the following PW status code is received:

0x00000000	Pseudowire forwarding (clear all failures)

This command is mutually exclusive with **no pw-status-signaling**, and **standby-signaling-slave**. It is not applicable to spoke-SDPs forming part of an MC-LAG or spoke-SDPs in an endpoint.

Default no block-on-peer-fault

cflowd

Syntax [**no**] **cflowd**

Context config>service>epipe>sap

Description This command enables cflowd to collect traffic flow samples through a service interface (SAP) for analysis. When cflowd is enabled on an Ethernet service SAP, the Ethernet traffic can be sampled and processed by the system's cflowd engine and exported to IPFIX collectors with the l2-ip template enabled.

cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When cflowd is enabled at the SAP level, all packets forwarded by the interface are subjected to analysis according to the cflowd configuration.

For L2 services, only ingress sampling is supported.

Default no cflowd

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>apipe>sap
config>service>fpipe>sap
config>service>epipe>sap

Description This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default no collect-stats

cpu-protection

Syntax **cpu-protection** *policy-id* [**mac-monitoring**] | [**eth-cfm-monitoring** [**aggregate**] [**car**]]
no cpu-protection

Context config>service>apipe>sap
config>service>epipe>spoke-sdp
config>service>epipe>sap

Description This command assigns an existing CPU protection policy to the associated service. The CPU protection policies are configured in the **config**>**sys**>**security**>**cpu-protection**>**policy** *cpu-protection-policy-id* context.

Default cpu-protection 254 (for access interfaces)

cpu-protection 255 (for network interfaces)

The configuration of **no cpu-protection** returns the interface/SAP to the default policies.

If no CPU protection policy is assigned to a service SAP then a the default policy is used to limit the overall-rate.

Parameters *policy-id* — Specifies an existing CPU protection policy.

Values *1 to 255*

mac-monitoring — This keyword enables MAC monitoring.

eth-cfm-monitoring — This keyword enables Ethernet Connectivity Fault Management monitoring.

aggregate — This keyword applies the rate limit to the sum of the per peer packet rates.

car — (Committed Access Rate) This keyword causes Eth-CFM packets to be ignored when enforcing the overall-rate.

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>epipe>sap
config>service>fpipe>sap
config>service>ipipe>sap

Description This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid existing DCP policy can be assigned to a SAP or a network interface (this rule does not apply to templates, such as an **msap-policy** template).

If no dist-cpu-protection policy is assigned to a SAP, then the default access DCP policy (_default-access-policy) is used.

If no DCP functionality is required on the SAP then an empty DCP policy can be created and explicitly assigned to the SAP

Parameters *policy-name* — Specifies the name of the DCP policy up to 32 characters in length

ethernet

Syntax **ethernet**

Context config>service>epipe>sap

Description This command enters the context to configure Ethernet properties in this SAP.

llf

Syntax [**no**] **llf**

Context config>service>apipe>sap>atm
config>service>epipe>sap>ethernet

Description This command enables Link Loss Forwarding (LLF) on an Ethernet port or an ATM port. This feature provides an end-to-end OAM fault notification for Ethernet VLL service and for ATM VLL service of vc-type atm-cell. It brings down the Ethernet port (Ethernet LLF) or sends a SONET/SDH Path AIS (ATM LLF) toward the attached CE when there is a local fault on the Pseudowire or service, or a remote fault on the SAP or pseudowire, signaled with label withdrawal or T-LDP status bits. It ceases when the fault disappears.

The Ethernet port must be configured for null encapsulation.

For the 7750 SR, the ATM port must be configured as a SAP on an Apipe service of vc-type atm-cell. The ATM port must also be configured on the following MDAs:

- 1-port OC12/STM4 ASAP MDA. At OC3/STM1 port level
- 4-port ATM MDA at OC12/STM4 or OC3/STM1 port level
- 16-port ATM MDA at OC3/STM1 port level

The ATM port must be configured as a SAP on an Apipe service of vc-type atm-cell. The ATM port must also be configured on the following MDAs:

- 1-port OC12/STM4 ASAP MDA. At OC3/STM1 port level
- 4-port ATM MDA at OC12/STM4 or OC3/STM1 port level
- 16-port ATM MDA at OC3/STM1 port level

2.17.2.5 Circuit Emulation Commands

cem

Syntax **cem**

Context config>service>cpipe>sap
config>service>epipe>sap

Description This command enters the context to specify circuit emulation (CEM) properties.

local-ecid

Syntax **local-ecid** *emulated circuit identifier*
no local-ecid

Context config>service>epipe>sap>cem

Description This command defines the Emulated Circuit Identifiers (ECID) to be used for the local (source) end of the circuit emulation service.

The **no** form of the command removes the ECID from the configuration.

Default 65535

Parameters *emulated circuit identifier* — Specifies the value to be used as the local (source) ECID for the circuit emulation service. On CES packet reception, the ECID in the packet will be compared to the configured local-ecid value. These must match for the packet payload to be used for the TDM circuit. The remote-ecid value is inserted into the MEF-8 CES packet to be transmitted.

Values 0 to 1048575

packet

Syntax **packet jitter-buffer** *milliseconds* [**payload-size** *bytes*]
packet payload-size *bytes*
no packet *bytes*

Context config>service>cpipe>sap
config>service>epipe>sap>cem

Description This command specifies the jitter buffer size, in milliseconds, and payload size, in bytes.

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots as shown in Table 14.

Table 14 **packet CEM SAP Endpoint Types**

Endpoint Type	Timeslots	Default Jitter Buffer (in ms)
unstructuredE1	n/a	5
unstructuredT1	n/a	5

Table 14 **packet CEM SAP Endpoint Types (Continued)**

Endpoint Type	Timeslots	Default Jitter Buffer (in ms)
nxDS0 (E1/T1)	—	32
	N = 1	16
	N = 2 to 4	8
	N = 5 to 15	5
nxDS0WithCas (E1)	N	8
nxDS0WithCas (T1)	N	12

Parameters *milliseconds* — Specifies the jitter buffer size in milliseconds (ms).

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffers is not allowed. Setting the jitter butter value to 0 sets it back to the default value.

Values 1 to 250

payload-size *bytes* — Specifies the payload size (in bytes) of packets transmitted to the packet service network (PSN) by the CEM SAP. This determines the size of the data that will be transmitted over the service. If the size of the data received is not consistent with the payload size then the packet is considered malformed.

Values 0 | 16 to 2048

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots as shown in Table 15.

Table 15 **CEM SAP Endpoint Types**

Endpoint Type	Timeslots	Default Payload Size (in bytes)
unstructuredE1	n/a	256
unstructuredT1	n/a	192
nxDS0 (E1/T1)	N = 1	64
	N = 2 to 4	N x 32
	N = 5 to 15	N x 16
	N >= 16	N x 8
nxDS0WithCas (E1)	N	N x 16
nxDS0WithCas (T1)	N	N x 24

For all endpoint types except for nxDS0WithCas, the valid payload size range is from the default to 2048 bytes.

For nxDS0WithCas, the payload size divide by the number of timeslots must be an integer factor of the number of frames per trunk multi-frame (for example, 16 for E1 trunk and 24 for T1 trunk).

For 1xDS0, the payload size must be a multiple of 2.

For NxDS0, where N > 1, the payload size must be a multiple of the number of timeslots.

For unstructuredE1 and unstructuredT1, the payload size must be a multiple of 32 bytes.

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffer is not allowed.

Setting the payload size to 0 sets it back to the default value.

remote-ecid

Syntax **remote-ecid** *emulated circuit identifier*
no remote-ecid

Context config>service>epipe>sap>cem

Description This command defines the Emulated Circuit Identifiers (ECID) to be used for the remote (destination) end of the circuit emulation service.

Parameters *emulated circuit identifier* — Specifies the value to be used as the remote (destination) ECID for the circuit emulation service. Upon CES packet reception, the ECID in the packet will be compared to the configured local-ecid value. These must match for the packet payload to be used for the TDM circuit. The remote-ecid value is inserted into the MEF-8 CES packet to be transmitted.

Values 0 to 1048575

remote-mac

Syntax **remote-mac** *ieee-address*
no remote-mac

Context config>service>epipe>sap>cem

Description This command defines the destination IEEE MAC address to be used to reach the remote end of the circuit emulation service.

Default 00:00:00:00:00:00

Parameters *ieee-address* — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

report-alarm

Syntax [**no**] **report-alarm** [**stray**] [**malformed**] [**pktloss**] [**overrun**] [**underrun**] [**rpktloss**] [**rfault**] [**rrdi**]

Context config>service>epipe>sap>cem

Description This command indicates the type of CEM SAP alarm.

The **no** form of the command removes the parameter from the configuration.

Default On: stray, malformed, pktloss and overrun

Off: rpktloss, rfault, rrdi

Parameters **stray** — Reports the reception of packets not destined for this CES circuit.

malformed — Reports the reception of packet not properly formatted as CES packets.

pktloss — Reports the lack of reception of CES packets.

overrun — Reports the reception of too many CES packets resulting in a overrun of the receive jitter buffer.

underrun — Reports the reception of too few CES packets resulting in a overrun of the receive jitter buffer.

rpktloss — Reports hat the remote peer is currently in packet loss status.

rfault — Reports that the remote TDM interface is currently not in service.

rrdi — Reports that the remote TDM interface is currently in RDI status.

rtp-header

Syntax [**no**] **rtp-header**

Context config>service>epipe>sap>cem
config>service>cpipe>sap>cem

Description This command specifies whether an RTP header is used when packets are transmitted to the packet service network (PSN) by the CEM SAP. This mode must be enabled for differential-timed DS1/E1s. It can optionally be enabled for other DS1/E1s for interoperability purposes.

Default no rtp-header

2.17.2.6 ETH-CFM Service Commands

eth-cfm

Syntax **eth-cfm**

Context config>service>epipe>spoke-sdp
config>service>epipe
config>service>epipe>sap
config>service>ipipe>sap

Description This command enters the context to configure ETH-CFM parameters.

ais-enable

Syntax [**no**] **ais-enable**

Context config>service>epipe>sap>eth-cfm
config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables the generation and the reception of AIS messages.

low-priority-defect

Syntax **low-priority-defect** {**allDef** | **macRemErrXcon**}

Context config>lag>eth-cfm>mep>ais
config>lag>eth-cfm>mep>ais
config>port>ethernet>eth-cfm>mep>ais
config>service>epipe>sap>eth-cfm>mep>ais
config>service>epipe>spoke-sdp>eth-cfm>mep>ais
config>service>vpls>mesh-sdp>eth-cfm>mep>ais

Description This command allows the operator to include all CCM Defect conditions or exclude the Remote Defect Indication CCM (DefRDICCM) as a trigger for generating AIS. AIS generation can only occur when the client-meg-level configuration option has been included. Changing this parameter will evaluate the MEP for AIS triggers based on the new criteria.

Parameters **allDef** — Keyword that includes any CCM defect condition to trigger AIS generation.

macRemErrXcon — Keyword that excludes RDI CCM Defect condition to trigger AIS generation.

collect-lmm-stats

Syntax **collect-lmm-stats**
no collect-lmm-stats

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm
config>service>vpls>sap>eth-cfm
config>service>vpls>spoke-sdp>eth-cfm
config>service>vpls>mesh-sdp>eth-cfm
config>service>ipipe>sap>eth-cfm

Description This command enables the collection of statistics on the SAP or MPLS SDP binding on which the ETH- LMM test is configured. The collection of LMM statistics must be enabled if a MEP is launching or responding to ETH-LMM packets. If LMM statistics collection is not enabled, the counters in the LMM and LMR PDU do not represent accurate measurements and all measurements should be ignored. The **show sap-using eth-cfm collect-lmm-stats** command and the **show sdp-using eth-cfm collect-lmm-stats command** can be used to display which entities are collecting stats.

The **no** form of the command disables and deletes the counters for this SAP or MPLS SDP binding.

Default no collect-lmm-stats

collect-lmm-fc-stats

Syntax **collect-lmm-fc-stats**

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm
config>service>ipipe>sap>eth-cfm

Description This command enters the context to configure per-forwarding class (FC) LMM information collection.

This command is mutually exclusive with the **collect-lmm-stats** command when there is entity resource contention.

fc

Syntax **fc** *fc-name* [*fc-name* ... (up to 8 max)]
no fc

Context config>service>epipe>sap>eth-cfm>collect-lmm-fc-stats
config>service>epipe>spoke-sdp>eth-cfm>collect-lmm-fc-stats
config>service>ipipe>sap>eth-cfm>collect-lmm-fc-stats

Description This command creates individual counters for the specified FCs without regard for profile. All countable packets that match a configured FC, regardless of profile, will be included in this counter.

A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.

Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-aware FC using the **fc-in-profile** command under the same context.

The **no** form of the command removes all previously defined FCs and stops counting for those FCs.

Default no fc

Parameters *fc-name* — Specifies the name of the FC for which to create an individual profile-unaware counter. In order for the counter to be used, the **config**>**oam-pm**>**session**>**ethernet**>**priority** command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the **config**>**oam-pm**>**session**>**ethernet**>**lmm**>**enable-fc-collection** command must be enabled.

Values nc, h1, ef, h2, l1, af, l2, be

fc-in-profile

Syntax **fc-in-profile** *fc-name* [*fc-name* ... (up to 8 max)]
no fc-in-profile

Context config>service>epipe>sap>eth-cfm>collect-lmm-fc-stats
config>service>epipe>spoke-sdp>eth-cfm>collect-lmm-fc-stats
config>service>ipipe>sap>eth-cfm>collect-lmm-fc-stats

Description This command creates individual counters for the specified FCs with regard for profile. All countable packets that match a configured FC and are deemed to be in-profile will be included in this counter.

A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.

Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-unaware FC using the **fc** command under the same context.

The **no** form of the command removes all previously defined FCs and stops counting for those FCs.

Default no fc-in-profile

Parameters *fc-name* — Specifies the name of the FC for which to create an individual profile-aware counter. In order for the counter to be used, the **config**>**oam-pm**>**session**>**ethernet**>**priority** command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the **config**>**oam-pm**>**session**>**ethernet**>**lmm**>**enable-fc-collection** command must be enabled.

Values nc, h1, ef, h2, l1, af, l2, be

interface-support-enable

Syntax [**no**] **interface-support-enable**

Context config>service>epipe>sap>eth-cfm>mep>ais
config>service>epipe>spoke-sdp>eth-cfm>mep>ais

Description This command enables the AIS function to consider the operational state of the entity on which it is configured. With this command, ETH-AIS on operationally down MEPs will be triggered and cleared based on the operational status of the entity on which it is configured. If CCM is also enabled then transmission of the AIS PDU will be based on either the non-operational state of the entity or on any CCM defect condition. AIS generation will cease if BOTH operationally up state and CCM has no defect conditions. If the MEP is not CCM enabled then the operational state of the entity is the only consideration assuming this command is present for the MEP.

Default no interface-support-enabled (AIS will not be generated or stopped based on the state of the entity on) which the operationally down MEP is configured.

client-meg-level

Syntax **client-meg-level** [[*level* [*level* ...]]
no client-meg-level

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>aid-enable

Description This command configures the client maintenance entity group (MEG) level or levels to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters *level* — Specifies the client MEG level.

Values 1 to 7

Default 1

interval

Syntax **interval** {**1** | **60**}
no interval

Context config>service>epipe>sap>eth-cfm>mep>ais-enable
config>service>epipe>spoke-sdp>eth-cfm>ais-enable

Description This command specifies the transmission interval of AIS messages in seconds.

Parameters **1 | 60** — Specifies the transmission interval of AIS messages in seconds.

Default 1

priority

Syntax **priority** *priority-value*
no priority

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>aid-enable

Description This command specifies the priority of AIS messages originated by the node.

Parameters *priority-value* — Specifies the priority value of the AIS messages originated by the node.

Values 0 to 7

Default 1

eth-tunnel

Syntax **eth-tunnel**

Context config>service>epipe>sap
config>service>ipipe>sap

Description The command enables the context to configure Ethernet tunnel SAP parameters.

path

Syntax **path** *path-index* **tag** *qtag*[.*qtag*]
no path *path-index*

Context config>service>epipe>sap>eth-tunnel
config>service>ipipe>sap>eth-tunnel

Description This command configures Ethernet tunnel SAP path parameters.

The **no** form of the command removes the values from the configuration.

Default none

Parameters *path-index* — Specifies the path index value.

Values 1 to 16

qtag[.*qtag*] — Specifies the qtag value.

Values 0 to 4094 | *

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}] [**primary-vlan-enable**]

no mep *mep-id* **domain** *md-index* **association** *ma-index*

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm
config>service>ipipe>sap>eth-cfm

Description This command provisions the maintenance endpoint (MEP).

The **no** form of the command reverts to the default values.

Parameters *mep-id* — Specifies the maintenance endpoint identifier.

Values 1 to 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 to 4294967295

ma-index — Specifies the maintenance association (MA) index value.

Values 1 to 4294967295

direction {up | **down}** — Indicates the direction in which the MEP faces on the bridge port. The UP direction is not supported for all Fpipe services. For example, Ipipe does not support the direction of UP for MEPs.

up — Sends ETH-CFM messages toward the MAC relay entity.

down — Sends ETH-CFM messages away from the MAC relay entity.

primary-vlan-enable — Provides a method for linking the MEP with the primary VLAN configured under the bridge-identifier for the MA. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Layer 2 Epipe and VPLS services.

ccm-enable

Syntax [**no**] **ccm-enable**

Context config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep

Description This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

Context config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep

Description This command specifies the priority value for CCMs and LTMs transmitted by the MEP.

The **no** form of the command removes the priority value from the configuration.

Default The highest priority on the bridge-port.

Parameters *priority* — Specifies the priority of CCM and LTM messages.

Values 0 to 7

ccm-padding-size

Syntax **ccm-padding-size** *ccm-padding*
no ccm-padding-size *ccm-padding*

Context config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep
config>service>epipe>sdp>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>eth-cfm>mep

config>router>if>eth-cfm>mep

Description Set the byte size of the optional Data TLV to be included in the ETH-CC PDU. This will increase the size of the ETH-CC PDU by the configured value. The base size of the ETH-CC PDU, including the Interface Status TLV and Port Status TLV, is 83 bytes not including the Layer Two encapsulation. CCM padding is not supported when the CCM-Interval is less than one second.

Default [no] ccm-padding-size

Parameters *ccm-padding* — Specifies the byte size of the Optional Data TLV.

Values 3 to 1500

csf-enable

Syntax [**no**] **csf-enable**

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables the reception and local processing of ETH-CSF frames.

multiplier

Syntax **multiplier** *multiplier-value*
no multiplier

Context config>service>epipe>sap>eth-cfm>mep>cfs-enable
config>service>epipe>spoke-sdp>eth-cfm>mep>cfs-enable

Description This command enables the multiplication factor applied to the receive time used to clear the CSF condition in increments of .5.

Default 3.5

Parameters *multiplier-value* — Specifies the multiplier used for timing out CSF.

Values 0.0 | 2.0 to 30.0

ccm-tlv-ignore

Syntax **ccm-tlv-ignore** [**interface-status**] [**port-status**]
no ccm-tlv-ignore

Context config>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>mep
config>router>if>eth-cfm>mep

Description This command allows the receiving MEP to ignore the specified TLVs in CCM PDU. Ignored TLVs will be reported as absent and will have no impact on the MEP state machine.

The **no** form of the command means the receiving MEP will process all recognized TLVs in the CCM PDU.

Default no ccm-tlv-ignore

Parameters **interface-status** — Ignores the interface status TLV on reception.

port-status — Ignores the port status TLV on reception.

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep

Description For this test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

A check is performed for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP indicates the problem.

bit-error-threshold

Syntax **bit-error-threshold** *errors*
no bit-error-threshold

Context config>service>epipe>sap>eth-cfm>mep>eth-test-enable

Description This command is used to specify the threshold value of bit errors.

Parameters *errors* — The threshold value of bit errors.

Values 0 to 11840

Default 1

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]

no test-pattern

Context config>service>epipe>spoke-sdp>eth-cfm>mep>eth-test-enable
config>service>epipe>sap>eth-cfm>mep>eth-test-enable
config>service>ipipe>sap>eth-cfm>mep>eth-test-enable

Description This command configures the test pattern for eth-test frames.

The **no** form of the command removes the values from the configuration.

Default all-zeros

Parameters **all-zeros** — Specifies to use all zeros in the test pattern.

all-ones — Specifies to use all ones in the test pattern.

crc-enable — Generates a CRC checksum.

fault-propagation-enable

Syntax **fault-propagation-enable** {**use-if-tlv** | **suspend-ccm**}
no fault-propagation-enable

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep

Description This command configures the fault propagation for the MEP.

Parameters **use-if-tlv** — Specifies to use the interface TLV.

suspend-ccm — Specifies to suspend continuity check messages.

grace

Syntax **grace**

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep

Description This command enters the context to configure Nokia ETH-CFM Grace and ITU-T Y.1731 ETH-ED expected defect functional parameters.

eth-ed

Syntax **eth-ed**

Context config>service>epipe>sap>eth-cfm>mep>grace

config>service>epipe>spoke-sdp>eth-cfm>mep>grace
config>service>ipipe>sap>eth-cfm>mep>grace

Description This command enters the context to configure ITU-T Y.1731 ETH-ED expected defect functional parameters.

max-rx-defect-window

Syntax **max-rx-defect-window** *seconds*
no max-rx-defect-window

Context config>service>epipe>sap>eth-cfm>mep>grace>eth-ed
config>service>epipe>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>ipipe>sap>eth-cfm>mep>grace>eth-ed

Description This command limits the duration of the received ETH-ED expected defect window to the lower value of either the received value from the peer or this parameter.

The **no** form of the command removes the limitation, and any valid defect window value received from a peer MEP in the ETH-ED PDU will be used.

Default no max-rx-defect-window

Parameters *seconds* — Specifies the duration, in seconds, of the maximum expected defect window.

Values 1 to 86400

priority

Syntax **priority** *priority*
no priority

Context config>service>epipe>sap>eth-cfm>mep>grace>eth-ed
config>service>epipe>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>ipipe>sap>eth-cfm>mep>grace>eth-ed

Description This command sets the priority bits and determines the forwarding class based on the mapping of priority to FC.

The **no** form of the command disables the local priority configuration and sets the priority to the **ccm-ltm-priority** associated with this MEP.

Default no priority

Parameters *priority* — Specifies the priority bit.

Values 0 to 7

rx-eth-ed

Syntax [**no**] **rx-eth-ed**

Context config>service>epipe>sap>eth-cfm>mep>grace>eth-ed
config>service>epipe>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>ipipe>sap>eth-cfm>mep>grace>eth-ed

Description This command enables the reception and processing of the ITU-T Y.1731 ETH-ED PDU on the MEP.

The **no** form of the command disables the reception of the ITU-T Y.1731 ETH-ED PDU on the MEP.

Default rx-eth-ed

tx-eth-ed

Syntax [**no**] **tx-eth-ed**

Context config>service>epipe>sap>eth-cfm>mep>grace>eth-ed
config>service>epipe>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>ipipe>sap>eth-cfm>mep>grace>eth-ed

Description This command enables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP when a system soft reset notification is received for one or more cards.

The **config**>**eth-cfm**>**system**>**grace-tx-enable** command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP.

Default no tx-eth-ed

eth-vsm-grace

Syntax **eth-vsm-grace**

Context config>service>epipe>sap>eth-cfm>mep>grace
config>service>epipe>spoke-sdp>eth-cfm>mep>grace
config>service>ipipe>sap>eth-cfm>mep>grace

Description This command enters the context to configure Nokia ETH-CFM Grace functional parameters.

rx-eth-vsm-grace

Syntax [**no**] **rx-eth-vsm-grace**

Context config>service>epipe>sap>eth-cfm>mep>grace>eth-vsm-grace
config>service>epipe>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace
config>service>ipipe>sap>eth-cfm>mep>grace>eth-vsm-grace

Description This command enables the reception and processing of the Nokia ETH-CFM Grace PDU on the MEP.

The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.

The **no** form of the command disables the reception of the Nokia ETH-CFM Grace PDU on the MEP.

Default rx-eth-vsm-grace

tx-eth-vsm-grace

Syntax [**no**] **tx-eth-vsm-grace**

Context config>service>epipe>sap>eth-cfm>mep>grace>eth-vsm-grace
config>service>epipe>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace
config>service>ipipe>sap>eth-cfm>mep>grace>eth-vsm-grace

Description This command enables the transmission of the Nokia ETH-CFM Grace PDU from the MEP when a system soft reset notification is received for one or more cards.

The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.

The **config**>**eth-cfm**>**system**>**grace-tx-enable** command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the Nokia ETH-CFM Grace PDU from the MEP.

Default tx-eth-vsm-grace

lbm-svc-act-responder

Syntax [**no**] **lbm-svc-act-responder**

Context config>service>epipe>sap>eth-cfm>mep

config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables the MEP to process service activation streams encapsulated in ETH-CFM LBM frames that are directed to the MEP. The MEP will be allocated additional resources to rapidly respond to a high-speed stream of LBM messages. A MEP created with this option will not validate any TLVs, will not validate the ETH-LBM MAC Address, and will not increment or compute any loopback statistics. Statistical computation and reporting is the responsibility of the test head-end. The ETH-CFM level of the high speed ETH-LBM stream must match the level of a MEP configured with this command. It must not target any lower ETH-CFM level the MEP will terminate. When the service activation test is complete, the MEP may be returned to standard processing by removing this command. If there is available bandwidth, the MEP will respond to other ETH-CFM PDUs, such as ETH-DMM marker packets, using standard processing.

The interaction between this command and the **tools perform service id** *service-id* **loopback eth** command must be carefully considered. It is recommended that either the **lbm-svc-act-responder** or the **tools perform service id** *service-id* **loopback eth** command be used at any given time within a service. If both commands must be configured, and the target reflection point is the MAC Swap Loopback function, the inbound stream of data must not include ETH-CFM traffic that is equal to or lower than the domain level of any configured MEP which would otherwise extract and process the ETH-CFM message. If the reflection target is a MEP configured with the **lbm-svc-act-responder** option, the mode (ingress or egress) of the SAP or SDP specified with this tools command and the MEP **direction** (up or down) must match when the functions are enabled on the same reflection point, and the domain level of the inbound ETH-LBM must be the same as that of the MEP configured with the **lbm-svc-act-responder** option. At no time should the two functions be conflicting with each other along the path of the stream. This conflict would lead to unpredictable and possibly destabilizing situations.

The **no** form of the command reverts to MEP LBM standard processing.

Default no lbm-svc-act-responder

low-priority-defect

Syntax **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}

Context config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>epipe>sap>eth-cfm>mep
config>service>ipipe>sap>eth-cfm>mep

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default macRemErrXcon

Parameters **low-priority-defect** — The low priority defect values are defined as follows:

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
errXcon	Only DefErrorCCM and DefXconCCM
xcon	Only DefXconCCM
noXcon	No defects DefXcon or lower are to be reported

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>service>vpls>sap>eth-cfm>mep

Description This command enables/disables eth-test functionality on MEP.

Parameters *seconds* — Specifies the one way-delay threshold in seconds.

Values 0 to 600

Default 3

mip

Syntax **mip** [**mac** *mac-address*] [**primary-vlan-enable** *vlan-id*]
mip default-mac [**primary-vlan-enable** *vlan-id*]
no mip

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm

Description This command allows Maintenance Intermediate Points (MIPs). The creation rules of the MIP are dependent on the mhf-creation configuration for the MA.

The **no** form of the command removes the MIP creation request.

Default no mip

Parameters **mac** — Provides a method for manually configuring the MIP MAC address.

mac-address — Specifies the MAC address of the MIP.

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MIP. The MAC address must be unicast. Using the all-zeros address is equivalent to the **no** form of this command.

default-mac — Using the **no** command deletes the MIP. This parameter should be used if the operator wants to change the MAC address back to the default MAC without having to delete and reconfigure the MIP.

primary-vlan-enable — Provides a method for linking the MIP with the primary VLAN configured under the bridge-identifier for the MA. This is only allowed if the mhf-creation method is static. MIPs cannot be changed from or to primary VLAN functions without first being deleted. This must be configured as part of the creation step and can only be changed by deleting the MIP and recreating it. Primary VLANs are only supported under Layer 2 Epipe and VPLS services.

vlan-id — Must match the vlan-id under the bridge-identifier for the MA that is appropriate for this service.

Values 0 to 4094

squelch-ingress-levels

Syntax **squelch-ingress-levels** [*md-level* [*md-level*…]]
no squelch-ingress-levels

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm
config>service>vpls>sap>eth-cfm
config>service>vpls>spoke-sdp>eth-cfm
config>service>vpls>mesh-sdp>eth-cfm
config>service>ipipe>sap>eth-cfm
config>service>template>vpls-sap-template>eth-cfm

Description This command defines the levels of the ETH-CFM PDUs that will silently be discarded on ingress into the SAP or SDP binding from the wire. All ETH-CFM PDUs inbound to the SAP or SDP binding will be dropped that match the configured levels without regard for any other ETH-CFM criteria. No statistical information or drop count will be available for any ETH-PDU that is silently discarded by this option. The operator must configure a complete contiguous list of md-levels up to the highest level that will be dropped. The command must be retyped in complete form to modify a previous configuration, if the operator does not want to delete it first.

The **no** form of the command removes the silent discarding of previously matching ETH-CFM PDUs.

Default no squelch-ingress-levels

Parameters *md-level* — Identifies the level.

Values 0 to 7

tunnel-fault

Syntax **tunnel-fault** {**accept** | **ignore**}

Context config>service>epipe>eth-cfm
config>service>epipe>sap>eth-cfm
config>service>ipipe>eth-cfm
config>service>ipipe>sap>eth-cfm

Description Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, Epipe will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an Epipe service as well as setting the operational flag. If AIS generation is the requirement for the Epipe services this command is not required. See the **ais-enable** command under **config>service>epipe>sap>eth-cfm>ais-enable** context for more details. This works in conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is “ignore” and the sap level command is “accept”. This means simply changing the service level command to “accept” will enable the feature for all SAPs. This is not required for Epipe services that only wish to generate AIS on failure.

Default ignore (Service Level)

accept (SAP Level for Epipe and VPLS)

Parameters **accept** — Shares fate with the facility tunnel MEP.

ignore — Does not share fate with the facility tunnel MEP.

2.17.2.7 Service Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>sap
config>service>ipipe>sap
config>service>epipe>sap

Description This command enters the context to configure egress SAP parameters.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing.

force-qinq-vc-forwarding

Syntax [**no**] **force-qinq-vc-forwarding**

Context config>service>epipe>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>pw-template

Description This command forces the data path to insert and remove two VLAN tags for spoke and mesh SDPS that have either **vc-type ether** or **vc-type vlan**. The use of this command is mutually exclusive with the **force-vlan-vc-forwarding** command.

The VLAN identifiers and dot 1p/DE bits used in the two VLAN tags are taken from the inner tag received on a qinq SAP or qinq mesh/spoke-SDP, or from the VLAN tag received on a dot1q SAP or mesh/spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**), or 0 if there is no service delimiting VLAN tag at the ingress SAP or mesh/spoke-SDP. Alternatively, the VLAN identifiers in both VLAN tags can be set to the value configured in the **vlan-vc-tag** parameter in the **pw-template** or under the mesh or spoke-SDP configuration.

The Ethertype used for both VLAN tags is 0x8100. A different Ethertype can be used for the outer VLAN tag by configuring the pseudowire template with the **use-provisioned-sdp** or **prefer-provisioned-sdp** options and setting the Ethertype using the **sdp vlan-vc-etype** parameter (this Ethertype value is then used for all mesh and spoke-SDPs using that SDP).

The **no** version of this command sets the default behavior.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

Context config>service>epipe>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command forces vc-vlan-type forwarding in the data path for spoke and mesh SDPs which have either vc-type. This command is not allowed on vlan-vc-type SDPs.

The **no** version of this command sets default behavior.

Default By default this feature is disabled.

ingress

Syntax **ingress**

Context config>service>apipe>sap
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>sap
config>service>ipipe>sap
config>service>epipe>sap
config>service>epipe>sap

Description This command enters the context to configure ingress SAP Quality of Service (QoS) policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing.

filter

Syntax **filter** [**ip** *ip-filter-id*]
filter [**ipv6** *ipv6-filter-id*]
filter [**mac** *mac-filter-id*]
no filter [**ip** *ip-filter-id*]
no filter [**ipv6** *ipv6-filter-id*]
no filter [**mac** *mac-filter-id*]

Context config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>epipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>spoke-sdp>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP matching criteria. Only one filter can be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *filter-id* with an ingress or egress SAP. The *filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

IPv6 filters are only supported by the 7450 ESS and 7750 SR but are not supported on a Layer 2 SAP that is configured with QoS MAC criteria. Also, MAC filters are not supported on a Layer 2 SAP that is configured with QoS IPv6 criteria.

Special Cases **Epipe —** Both MAC and IP filters are supported on an Epipe service SAP.

Parameters *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6-filter-id — Specifies the IPv6 filter policy for 7450 ESS or 7750 SR. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac-filter-id — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

l2tpv3

Syntax **l2tpv3**

Context config>service>epipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>ingress

Description This command enters the context to configure L2TPv3 spoke-SDPs for Epipe services.

cookie

Syntax **cookie** [*cookie1*] [*cookie2*]
no cookie

Context config>service>epipe>spoke-sdp>egress>l2tpv3
config>service>epipe>spoke-sdp>ingress>l2tpv3

Description This command configures the RX/TX cookie for L2TPv3 spoke-SDPs for Epipe services. The RX cookie must match the configured TX cookie on a far-end node, while the TX cookie must match the configured RX cookie on a far-end node. If a mismatch is detected between the configured (far-end binding cookie) to what is received by the local IP address of the SDP a flag is set and must be manually cleared by an operator.

The purpose of the cookie is to provide validation against misconfiguration of service endpoints, and to ensure that the right service egress is being used.

One egress cookie and up to two ingress cookies may be configured per spoke-SDP binding. One or two cookies can be configured for matching ingress packets from the far-end node, in order to support cookie rollover without dropping packets. When a cookie is not configured, SR-OS assumes a value of 00:00:00:00:00:00:00:00.

A cookie is not mandatory. An operator may delete an egress cookie or either or both ingress cookies.

Default no cookie1 cookie2

Parameters *cookie1* — Specifies the first cookie, in the form of a 64-bit colon-separated hex value.

cookie2 — Specifies the second cookie, in the form of a 64-bit colon-separated hex value.

hsmda-queue-override

Syntax [**no**] **hsmda-queue-override**

Context config>service>epipe>sap>egress
config>service>ipipe>sap>egress

Description This command configures HSMDA egress and ingress queue overrides.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
no packet-byte-offset

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4-byte CRC (everything except the preamble and inter-frame gap). For example, this command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 20 bytes may be added to the packet and up to 43 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As mentioned previously, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. When the queue group represents the last-mile bandwidth constraints for a subscriber, the offset allows the HSMDA queue group to provide an accurate accounting to prevent overrun and underrun conditions for the subscriber. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber’s packets on an Ethernet aggregation network.

The packet-byte-offset value can be overridden for the HSMDA queue at the SAP or subscriber profile level.

The **no** form of the command removes any accounting size changes to packets handled by the queue. The command does not effect overrides that may exist on SAPs or subscriber profiles associated with the queue.

Parameters *add-bytes* — Specifies a byte value to add to packets for queue and queue group-level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 0 to 31

sub-bytes — Specifies a byte value to subtract from packets for queue and queue group-level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 1 to 64

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command, within the QoS policy hsmda-queue context, is a container for the configuration parameters controlling the behavior of an HSMDA queue. Unlike the standard QoS policy queue command, this command is not used to actually create or dynamically assign the queue to the object which the policy is applied. The queue identified by queue-id always exists on the SAP or subscriber context whether the command is executed or not. In the case of HSMDA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDA queue group to the object (both ingress and egress).

Best-Effort, Expedited and Auto-Expedite Queue Behavior Based on Queue-ID

With standard service queues, the scheduling behavior relative to other queues is based on two items, the queues Best-Effort or Expedited nature and the dynamic rate of the queue relative to the defined CIR. HSMDA queues are handled differently. The create time auto-expedite and explicit expedite and best-effort qualifiers have been eliminated and instead the scheduling behavior is based solely on the queues identifier. Queues with a queue-id equal to 1 are placed in scheduling class 1. Queues with queue-id 2 are placed in scheduling class 2. And so on up to scheduling class 8. Each scheduling class is either mapped directly to a strict scheduling priority level based on the class ID, or the class may be placed into a weighted scheduling class group providing byte fair weighted round robin scheduling between the members of the group. Two weighted groups are supported and each may contain up to three consecutive scheduling classes. The weighted group assumes its highest member class inherent strict scheduling level for scheduling purposes. Strict priority level 8 has the highest priority while strict level 1 has the lowest. When grouping of scheduling classes is defined, some of the strict levels will not be in use.

Single Type of HSMDA Queues

Another difference between HSMDA queues and standard service queues is the lack of Multipoint queues. At ingress, an HSMDA SAP or subscriber does not require Multipoint queues since all forwarding types (broadcast, multicast, unicast and unknown) forward to a single destination ñ the ingress forwarding plane on the IOM. Instead of a possible eight queues per forwarding type (for a total of up to 32) within the SAP ingress QoS policy, the hsmda-queues node supports a maximum of eight queues.

Every HSMDA Queue Supports Profile Mode Implicitly

Unlike standard service queues, the HSMDA queues do not need to be placed into the special mode profile at create time in order to support ingress color aware policing. Each queue may handle in-profile, out-of-profile and profile undefined packets simultaneously. As with standard queues, the explicit profile of a packet is dependent on ingress sub-forwarding class to which the packet is mapped.

The **no** form of the command restores the defined queue-id to its default parameters. All HSMDA queues having the queue-id and associated with the QoS policy are re-initialized to default parameters.

Parameters *queue-id* — Specifies the HSMDA queue to use for packets in this forwarding class. This mapping is used when the SAP is on a HSMDA MDA.

Values 1 to 8

rate

Syntax **rate** *pir-rate*
no rate

Context config>service>epipe>sap>egress>hsmda-queue-over>queue
config>service>ipipe>sap>egress>hsmda-queue-over>queue

Description This command specifies the administrative PIR by the user.

Parameters *pir-rate* — Configures the administrative PIR specified by the user.

Values 1 to 40000000, **max**

wrr-weight

Syntax **wrr-weight** *value*
no wrr-weight

Context config>service>epipe>sap>egress>hsmda-queue-over>queue
config>service>ipipe>sap>egress>hsmda-queue-over>queue

Description This command assigns the weight value to the HSMDA queue.

The **no** form of the command returns the weight value for the queue to the default value.

Parameters *value* — Specifies the weight for the HSMDA queue.

Values 1 to 32

wrr-policy

Syntax **wrr-policy** *hsmda-wrr-policy-name*

no wrr-policy

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

Parameters *hsmda-wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue.

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command assigns an HSMDA slope policy to the SAP. The policy may be assigned to an ingress or egress HSMDA queue. The policy contains the Maximum Buffer Size (MBS) that will be applied to the queue and the high and low priority RED slope definitions. The function of the MBS and RED slopes is to provide congestion control for an HSMDA queue. The MBS parameter defines the maximum depth a queue may reach when accepting packets. The low and high priority RED slopes provides for random early detection of congestion and slope based discards based on queue depth.

An HSMDA slope policy can be applied to queues defined in the SAP ingress and SAP egress QoS policy HSMDA queues context. Once an HSMDA slope policy is applied to a SAP QoS policy queue, it cannot be deleted. Any edits to the policy are updated to all HSMDA queues indirectly associated with the policy.

Default HSMDA Slope Policy

An HSMDA slope policy named “default” always exists on the system and does not need to be created. The default policy is automatically applied to all HSMDA queues unless another HSMDA slope policy is specified for the queue. The default policy cannot be modified or deleted. Attempting to execute the **no hsmda-slope-policy** *default* command results in an error.

The **no** form of the command removes the specified HSMDA slope policy from the configuration. If the HSMDA slope policy is currently associated with an HSMDA queue, the command will fail.

Parameters *hsmda-slope-policy-name* — Specifies a HSMDA slope policy up to 32 characters in length. The HSMDA slope policy must be exist prior to applying the policy name to an HSMDA queue.

secondary-shaper

Syntax **secondary-shaper** *secondary-shaper-name*
no secondary-shaper

Context config>service>epipe>sap>egress>hsmda-queue-over
config>service>ipipe>sap>egress>hsmda-queue-over

Description This command configures an HSMDA egress secondary shaper.

Parameters *secondary-shaper-name* — Specifies a secondary shaper name up to 32 characters in length.

filter

Syntax **filter** [**ip** *ip-filter-id*]
filter [**ipv6** *ipv6-filter-id*]
no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
no filter [**ip** *ip-filter-id*]

Context config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>cpipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>spoke-sdp>ingress

Description This command associates a filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP matching criteria. Only one filter can be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress SAP. The *ip-filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6-filter-id — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

qos

Syntax **qos** *policy-id* [**shared-queuing**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>apipe>sap>ingress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>ingress
config>service>epipe>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy-id does not exist, an error will be returned.

The **qos** command, when used under the ingress context, is used to associate ingress QoS policies. The **qos** command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Default none

Parameters *policy-id* — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.

Values 1 to 65535

shared-queuing — This keyword can only be specified on SAP ingress. The shared-queuing keyword specifies the shared queue policy will be used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

fp-redirect-group — This keyword can only be used on SAP ingress and associates a SAP ingress with an instance of a named queue group template on the ingress forwarding plane of a specified IOM/IMM/XMA. The queue-group-name and **instance** *instance-id* are mandatory parameters when executing the command.

queue-group-name — Specifies the name of the queue group to be instance on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The *queue-group-name* must correspond to a valid ingress forwarding plane queue group, created under **config**>**card**>**fp**>**ingress**>**access**.

instance-id — Specifies the instance of the named queue group on the IOM/IMM/XMA ingress forwarding plane.

qos

Syntax **qos** *policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*
qos *policy-id*
no qos [*policy-id*]

Context config>service>apipe>sap>egress
config>service>cpipe>sap>egress
config>service>fpipe>sap>egress
config>service>ipipe>sap>egress
config>service>epipe>sap>egress

Description This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command, when used under the egress context, is used to associate egress QoS policies.

The **qos** command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Default none

Parameters *policy-id* — The egress policy ID to associate with SAP on egress. The policy ID must already exist.

Values 1 to 65535

queue-group-name — Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under *config>port>ethernet>access>egress*.

instance-id — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.

Values 1 to 40960

Default 1

queue-override

Syntax [**no**] **queue-override**

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress

Description This command enters the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy. If the policy was created as a template policy, this command overrides the parameter and its description and queue parameters in the policy.

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>service>apipe>sap>egress>queue-override
config>service>apipe>sap>ingress>queue-override
config>service>cpipe>sap>egress>queue-override
config>service>cpipe>sap>ingress>queue-override
config>service>fpipe>sap>egress>queue-override
config>service>fpipe>sap>ingress>queue-override
config>service>ipipe>sap>egress>queue-override
config>service>ipipe>sap>ingress>queue-override
config>service>epipe>sap>egress>queue-override
config>service>epipe>sap>ingress>queue-override

Description This command specifies the ID of the queue whose parameters are to be overridden.

Parameters *queue-id* — The queue ID whose parameters are to be overridden.

Values 1 to 32

create — This keyword is mandatory when creating a queue.

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*}] [**cir** *adaptation-rule*}]
no adaptation-rule

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default no adaptation-rule

Parameters **pir** — The **pir** parameter defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>service>apipe>sap>egress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a SONET or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.

- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 0.1 or 1000 octets.

  For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50 x 20 or 1000 octets.
- Frame based offered-load — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- Packet to frame factor — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be 1000 / 10000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- Frame based CIR — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

  As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.
- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to figure the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

On the 7450 ESS and 7750 SR, SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0.00 to 100.00

burst-limit

Syntax **burst-limit** {**default** | *size* [**bytes** | **kilobytes**]}
no burst-limit

Context config>service>apipe>sap>egress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue

Description The queue **burst-limit** command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override queue context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy.

Default no burst-limit

Parameters **default** — Reverts the queue's burst limit to the system default value.

size — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** qualifier must be added following size.

Values 1 to 13671 kilobytes
14000000 bytes

Default No default for *size*; use the **default** keyword to specify default burst limit.

bytes — Specifies that the value given for *size* must be interpreted as the burst limit in bytes.

kilobytes — Specifies that the value given for *size* must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is **kilobytes**.

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a specific access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly to drop packets.

The **no** form of this command returns the CBS size to the default value.

Default no cbs

Parameters *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is wanted, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 131072, default

drop-tail

Syntax **drop-tail**

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

Description This command enters the context to configure queue drop tail parameters.

low

Syntax **low**

Context config>service>apipe>sap>egress>queue-override>queue>drop-tail
config>service>apipe>sap>ingress>queue-override>queue>drop-tail
config>service>cpipe>sap>egress>queue-override>queue>drop-tail
config>service>cpipe>sap>ingress>queue-override>queue>drop-tail
config>service>fpipe>sap>egress>queue-override>queue>drop-tail
config>service>fpipe>sap>ingress>queue-override>queue>drop-tail
config>service>ipipe>sap>egress>queue-override>queue>drop-tail
config>service>ipipe>sap>ingress>queue-override>queue>drop-tail
config>service>epipe>sap>egress>queue-override>queue>drop-tail
config>service>epipe>sap>ingress>queue-override>queue>drop-tail

Description This command enters the context to configure the queue low drop-tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and will be discarded.

percent-reduction-from-mbs

Syntax **percent-reduction-from-mbs** *percent*
no percent-reduction-from-mbs

Context config>service>apipe>sap>egress>queue-override>queue>drop-tail>low
config>service>apipe>sap>ingress>queue-override>queue>drop-tail>low
config>service>cpipe>sap>egress>queue-override>queue>drop-tail>low
config>service>cpipe>sap>ingress>queue-override>queue>drop-tail>low
config>service>fpipe>sap>egress>queue-override>queue>drop-tail>low
config>service>fpipe>sap>ingress>queue-override>queue>drop-tail>low
config>service>ipipe>sap>egress>queue-override>queue>drop-tail>low
config>service>ipipe>sap>ingress>queue-override>queue>drop-tail>low
config>service>epipe>sap>egress>queue-override>queue>drop-tail>low
config>service>epipe>sap>ingress>queue-override>queue>drop-tail>low

Description This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes. Any out-of-profile packets will not be accepted into the queue if its depth is greater than this value, and so will be discarded.

Parameters *percent* — Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's MBS parameters. A queue uses its MBS value to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the number of buffers allowed by the MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope associated with a packet. A queue that has not exceeded its MBS is not guaranteed to have buffer available when needed or that the packet's RED slope will not force the discard of the packet. Setting correct CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS assigned to the queue to the default value.

Default default

Parameters *size* — The size parameter is an integer expression of the maximum number of kilobytes or bytes of buffering allowed for the queue. A value of 0 causes the queue to discard all packets.

Values 0 to 1073741824, default

bytes — Indicates that the *size* parameter value is expressed in bytes.

kilobytes — Indicates that the *size* parameter is expressed in kilobytes.

monitor-depth

Syntax [**no**] **monitor-depth**

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue

Description This command enables queue depth monitoring for the specified queue.

The **no** form of the command removes queue depth monitoring for the specified queue.

parent

Syntax **parent** {[**weight** *weight*] [**cir-weight** *cir-weight*]}
no parent

Context config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue

Description This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier** *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state and automatically returns to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters *weight* — These optional keywords are mutually exclusive to the **level** keyword. The weight defines the relative weight of this queue in comparison to other child schedulers, policers, and queues while vying for bandwidth on the parent *scheduler-name*. Any policers, queues, or schedulers defined as weighted receive no parental bandwidth until all policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active policers, queues, and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the policer, queue, or scheduler. A weight is considered to be active when the pertaining policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit. All child policers, queues, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted policers, queues, and schedulers at that level are operating at the maximum bandwidth or are idle.

Values 0 to 100

Default 1

cir-weight — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the policer, queue, or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]

Context config>service>epipe>sap>egress>queue-override>queue

Description The **percent-rate** command supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gb and a 10-Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

Parameters *pir-percent* — Specifies the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.01 to 100.00

Default 100.00

cir-percent — Specifies the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.00 to 100.00

Default 100.00

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>cpipe>sap>egress>queue-override>queue
config>service>cpipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile and then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default rate max cir 0

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits per second, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 3200000000, **max**

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues.

Values 0 to 3200000000, **max**, **sum**

Default 0

scheduler-override

Syntax [**no**] **scheduler-override**

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress
config>service>fpipe>sap>egress

config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress

Description This command specifies the set of attributes whose values have been overridden by management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

scheduler

Syntax [**no**] **scheduler** *scheduler-name* [**create**]

Context config>service>apipe>sap>egress>sched-override
config>service>apipe>sap>ingress>sched-override
config>service>cpipe>sap>egress>sched-override
config>service>cpipe>sap>ingress>sched-override
config>service>fpipe>sap>egress>sched-override
config>service>fpipe>sap>ingress>sched-override
config>service>ipipe>sap>egress>sched-override
config>service>ipipe>sap>ingress>sched-override
config>service>epipe>sap>egress>sched-override
config>service>epipe>sap>ingress>sched-override

Description This command can be used to override specific attributes of the specified scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers, policers, and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause policers, queues, or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the following criteria, a name syntax error will occur, the command will not execute, and the CLI context will not change.

Parameters *scheduler-name* — The name of the scheduler.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

Default None. Each scheduler must be explicitly created.

create — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

parent

Syntax **parent** [**weight** *weight*] [**cir-weight** *cir-weight*]
no parent

Context config>service>apipe>sap>ingress>sched-override>scheduler
config>service>apipe>sap>egress>sched-override>scheduler
config>service>cpipe>sap>ingress>sched-override>scheduler
config>service>cpipe>sap>egress>sched-override>scheduler
config>service>epipe>sap>ingress>sched-override>scheduler
config>service>epipe>sap>egress>sched-override>scheduler
config>service>fpipe>sap>ingress>sched-override>scheduler
config>service>fpipe>sap>egress>sched-override>scheduler
config>service>ipipe>sap>ingress>sched-override>scheduler
config>service>ipipe>sap>egress>sched-override>scheduler

Description This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers.

The **no** form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default no parent

Parameters *weight* — **Weight** defines the relative weight of this scheduler in comparison to other child schedulers, policers, and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue, or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

cir-weight — The **cir-weight** keyword defines the relative weight of this scheduler in comparison to other child schedulers, policers, and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>apipe>sap>egress>sched-override>scheduler
config>service>apipe>sap>ingress>sched-override>scheduler
config>service>cpipe>sap>egress>sched-override>scheduler
config>service>cpipe>sap>ingress>sched-override>scheduler
config>service>fpipe>sap>egress>sched-override>scheduler
config>service>fpipe>sap>ingress>sched-override>scheduler

config>service>ipipe>sap>egress>sched-override>scheduler
config>service>ipipe>sap>ingress>sched-override>scheduler
config>service>epipe>sap>egress>sched-override>scheduler
config>service>epipe>sap>ingress>sched-override>scheduler

Description This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child policers, queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child policers or queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child policers, queues, and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the values configured in the applied scheduler policy.

Parameters *pir-rate* — The **pir** parameter accepts the **max** keyword or a value of 1 to 3200000000. Any other value will result in an error without modifying the current PIR rate.

Values 1 to 3200000000, **max**

cir *cir-rate* — The **cir** parameter accepts a value of 0 to 3200000000 or the **max** keyword. Any other value will result in an error without modifying the current CIR rate.

If the **cir** parameter is set to **max**, then the CIR rate is set to infinity but bounded by the PIR rate.

The **sum** keyword specifies that the CIR will be used as the summed CIR values of the children schedulers, policers, or queues.

Values 0 to 3200000000, **max**, **sum**

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>service>apipe>sap>ingress
config>service>apipe>sap>egress

config>service>cpipe>sap>ingress
config>service>cpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>fpipe>sap>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>epipe>sap>ingress
config>service>epipe>sap>egress

Description This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created when the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues associated with the customer site. Policers or queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have policers or queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers or queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

Parameters *scheduler-policy-name* — The *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues and to egress policers managed by HQoS created on associated SAPs.

vlan-translation

Syntax **vlan-translation** {*vlan-id* | **copy-outer**}
no vlan-translation

Context config>service>epipe>sap>ingress

Description This command configures ingress VLAN translation. If enabled with an explicit VLAN value, the preserved vlan-id will be overwritten with this value. This setting is applicable to dot1q encapsulated ports. If enabled with “copy-outer” keyword, the outer vlan-id will be copied to inner position on QinQ encapsulated ports. The feature is not supported on default-dot1q saps (1/1/1:* and 1/1/1:0), nor on TopQ saps.

The **no** version of the command sets the default value and no action will be taken.

Default By default, the preserved VLAN values will not be overwritten.

Parameters *vlan-id* — Specifies that the preserved vlan-id will be overwritten with this value.

Values 0 to 4094

copy-outer — Keyword specifies to use the outer VLAN ID.

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p de

Context config>service>ipipe>sap>ingress
config>service>epipe>sap>ingress

Description This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The setting also applies to classification based on the DE indicator bit.

The **no** form of this command reverts the dot1p and de bits matching to the default tag.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 16 defines the default behavior for Dot1P evaluation.

Table 16 **Default QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Default no match-qinq-dot1p (no filtering based on p-bits)

(top or bottom must be specified to override the default QinQ dot1p behavior)

Parameters **top** — The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. Table 17 defines the dot1p evaluation behavior when the top parameter is specified.

Table 17 **Top Position QinQ dpt1p Evaluation Behavior**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	TopQ PBits

bottom — The bottom parameter is mutually exclusive to the top parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. Table 18 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 18 **Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits

Table 18 **Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Table 19 **Egress SAP Types**

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
Null	No preserved Dot1P bits	None
Null	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
Dot1Q	No preserved Dot1P bits	New PBits marked using dot1p-value
Dot1Q	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
TopQ	No preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	No preserved Dot1P bits	TopQ PBits and BottomQ PBits marked using dot1p-value
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits and BottomQ PBits marked using dot1p-value

The QinQ and TopQ SAP PBit/DEI bit marking follows the default behavior defined in the preceding table when **qinq-mark-top-only** is not specified.

The dot1p *dot1p-value* command must be configured without the qinq-mark-top-only parameter to remove the TopQ PBits only marking restriction.

A QinQ-encapsulated Ethernet port can have two different sap types:

For a TopQ SAP type, only the outer (top) tag is explicitly specified. For example, **sap 1/1/1:10.***

For QinQ SAP type, both inner (bottom) and outer (top) tags are explicitly specified. For example, **sap 1/1/1:10.100**.

2.17.2.8 VLL Frame Relay Commands

frame-relay

Syntax **frame-relay**

Context config>service>apipe>sap
config>service>fpipe>sap
config>service>ipipe>sap
config>service>epipe>sap

Description This command enters the context to configure Frame Relay parameters.

frf-12

Syntax [**no**] **frf-12**

Context config>service>fpipe>sap>frame-relay
config>service>ipipe>sap>frame-relay
config>service>epipe>sap>frame-relay

Description This command enables the use of FRF12 headers.

The **no** form of the command disables the use of FRF12 headers.

ete-fragment-threshold

Syntax **ete-fragment-threshold** *threshold*
no ete-fragment-threshold

Context config>service>fpipe>sap>frame-relay>frf-12
config>service>ipipe>sap>frame-relay>frf-12
config>service>epipe>sap>frame-relay>frf-12

Description This command specifies the maximum length of a fragment to be transmitted.

The **no** form of the command reverts to the default.

Parameters *threshold* — The maximum length of a fragment to be transmitted.

Values 128 to 512

Default 0

interleave

Syntax [**no**] **interleave**

Context config>service>epipe>sap>frame-relay>frf.12
config>service>ipipe>sap>frame-relay>frf.12

Description This command enables interleaving of high priority frames and low-priority frame fragments within a FR SAP using FRF.12 end-to-end fragmentation.

When this option is enabled, only frames of the FR SAP non-expedited forwarding class queues are subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI).

When this option is disabled, frames of all the FR SAP forwarding class queues are subject to fragmentation. The fragmentation header is however not included when the frame size is smaller than the user configured fragmentation size. In this mode, the SAP transmits all fragments of a frame before sending the next full or fragmented frame.

The receive direction of the FR SAP supports both modes of operation concurrently, with and without fragment interleaving.

The **no** form of this command restores the default mode of operation.

Default no interleave

scheduling-class

Syntax **scheduling-class** *class-id*

Context config>service>apipe>sap>frame-relay
config>service>fpipe>sap>frame-relay
config>service>ipipe>sap>frame-relay
config>service>epipe>sap>frame-relay

Description This command specifies the scheduling class to use for this SAP.

Parameters *class-id* — Specifies the scheduling class to use for this sap.

Values 0 to 3

Default 0

2.17.2.9 VLL SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**no-endpoint**]
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>cpipe
config>service>epipe

Description This command binds a service to an existing Service Distribution Point (SDP). A spoke-SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an Epipe, VPLS, VPRN, VPRN service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

This command can also be used to associate a GRE tunnel carrying Ethernet payload with an Epipe and terminate it on a PW port referenced within the same Epipe service. The spoke SDP represents a L2oGRE tunnel with SDP delivery type set to **eth-gre-bridged**. With this configuration, the **vc-id** is unused since there is no multiplexing of Ethernet payload within the same tunnel. The **vc-id** value is included only to maintain the expected spoke SDP structure within an EPIPE service. For L2oGRE tunnels, the **vc-id** can be set to any arbitrary value within its configurable range.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default No *sdp-id* is bound to a service.

Special Cases **Epipe —** At most, only one *sdp-id* can be bound to an Epipe service. Since an Epipe is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. Vc-switching VLLs are an exception. If the VLL is a “vc-switching” VLL, then the two endpoints must both be SDPs.

L2TPv3 SDP types are only supported on Epipe services and not other xpipe services.

Parameters *sdp-id* — The SDP identifier.

Values 1 to 17407

vc-id — The virtual circuit identifier. The VC-ID is not used with L2TPv3 SDPs or L2oGRE tunnels, however it must be configured.

Values 1 to 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

The VC type value for a VPLS service is defined as 0x000B.

Values ethernet

ether — Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke-SDP binding.

vlan — Defines the VC type as VLAN. The top VLAN tag, if a VLAN tag is present, is stripped from traffic received on the pseudowire, and a vlan-tag is inserted when forwarding into the pseudowire. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings.

The VLAN VC-type requires at least one dot1q tag within each encapsulated Ethernet packet transmitted to the far end.

Note: The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

no-endpoint — Removes the association of a spoke-SDP with an explicit endpoint name.

endpoint-name — Specifies the name of the service endpoint.

icb — Configures the spoke-SDP as an inter-chassis backup SDP binding.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**no-endpoint**]
spoke-sdp *sdp-id*[:*vc-id*] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>apipe
config>service>fpipe
config>service>ipipe
config>service>cpipe

Description This command binds a service to an existing Service Distribution Point (SDP). A spoke-SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end SR/ESS devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default No *sdp-id* is bound to a service.

Parameters *sdp-id* — The SDP identifier.

Values 1 to 17407

vc-id — The virtual circuit identifier.

Values 1 to 4294967295

no-endpoint — Adds or removes a spoke-SDP association.

endpoint-name — Specifies the name of the service endpoint.

icb — Configures the spoke-SDP as an inter-chassis backup SDP binding.

aarp

Syntax **aarp** *aarp-id* **type** {**subscriber-side shunt** | **network-side shunt**}
no aarp

Context config>service>ipipe>spoke-sdp

Description This command associates an AARP instance to an Ipipe spoke SDP. This instance is paired with the same *aarp-id* in a peer node as part of a configuration to provide flow and packet asymmetry removal for traffic for a multi-homed SAP or spoke SDP. The **type** parameter specifies the role of this service point in the AARP instance.

The **no** form of the command removes the association.

Default no aarp

Parameters *aarp-id* — An integer that identifies an AARP instance.

Values 1 to 65535

subscriber-side shunt — Specifies that the AARP type is an inter-chassis shunt service for subscriber-side traffic.

network-side shunt — Specifies that the AARP type is an inter-chassis shunt service for network-side traffic.

entropy-label

Syntax [**no**] **entropy-label**

Context config>service>epipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>pw-template
config>service>vpls>bgp-evpn>mpls
config>service>epipe>bgp-evpn>mpls

Description This command enables or disables the use of entropy labels for spoke-SDPs.

If **entropy-label** is configured, the entropy label and ELI are inserted in packets for which at least one LSP in the stack for the far-end of the tunnel used by the service has advertised entropy-label-capability. If the tunnel is RSVP type, **entropy-label** can also be controlled under the **config>router>mpls** or **config>router>mpls>lsp** contexts.

The entropy label and hash label features are mutually exclusive. The entropy label cannot be configured on a spoke-SDP or service where the hash label feature has already been configured.

Default no entropy-label

hash-label

Syntax **hash-label** [**signal-capability**]
no hash-label

Context config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>pw-template

Description This command enables the use of the hash label on a VLL, VPRN or VPLS service bound to any MPLS type encapsulated SDP, as well as to a VPRN service that is using the **auto-bind-tunnel** with the **resolution-filter** set to any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option. This feature is also not supported on multicast packets forwarded using RSVP P2MP LSP or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance. It is, however, supported when forwarding multicast packets using an IES/VPRN spoke-interface.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).

To allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note, however, that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL PW packets.

Packets generated in CPM and that are forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh SDP, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The 7450 ESS or 7750 SR local PE will insert the flow label interface parameters sub-TLV with F=1 in the PW ID FEC element in the label mapping message for that spoke-SDP or mesh SDP.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the PW but must not insert the hash label in the user and control packets over that spoke-SDP or mesh SDP. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-SDP or mesh SDP at the remote PE, the PW packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-SDP or mesh SDP at the remote PE, the PW received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the 7450 ESS or 7750 SR must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.

The **no** form of this command disables the use of the hash label.

Default no hash-label

Parameters **signal-capability** — Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

cell-concatenation

Syntax **cell-concatenation**

Context config>service>apipe>spoke-sdp

Description This command enters the context to provide access to the various options that control the termination of ATM cell concatenation into an MPLS frame. Several options can be configured simultaneously. The concatenation process for a specified MPLS packet ends when the first concatenation termination condition is met. The concatenation parameters apply only to ATM N:1 cell mode VLL.

aal5-frame-aware

Syntax [**no**] **aal5-frame-aware**

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of AAL5 end-of-message (EOM) to be an indication to complete the cell concatenation operation.

The **no** form of the command resets the configuration to ignore the AAL5 EOM as an indication to complete the cell concatenation.

clp-change

Syntax [**no**] **clp-change**

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of CLP change to be an indication to complete the cell concatenation operation.

The **no** form of the command resets the configuration to ignore the CLP change as an indication to complete the cell concatenation.

max-cells

Syntax **max-cells** *cell-count*
no max-cells [*cell-count*]

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of the maximum number of ATM cells to accumulate into an MPLS packet. The remote peer will also signal the maximum number of concatenated cells it is willing to accept in an MPLS packet. When the lesser of (the configured value and the signaled value) number of cells is reached, the MPLS packet is queued for transmission onto the pseudowire. It is ensured that the MPLS packet MTU conforms to the configured service MTU.

The **no** form of this command sets max-cells to the value '1' indicating that no concatenation will be performed.

Parameters *cell-count* — Specifies the maximum number of ATM cells to be accumulated into an MPLS packet before queuing the packet for transmission onto the pseudowire.

Values 1 to 128

Default 1

max-delay

Syntax **max-delay** *delay-time*

no max-delay [*delay-time*]

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of the maximum amount of time to wait while performing ATM cell concatenation into an MPLS packet before transmitting the MPLS packet. This places an upper bound on the amount of delay introduced by the concatenation process. When this amount of time is reached from when the first ATM cell for this MPLS packet was received, the MPLS packet is queued for transmission onto the pseudowire.

The **no** form of this command resets max-delay to its default value.

Parameters *delay-time* — Specifies the maximum amount of time, in hundreds of microseconds, to wait before transmitting the MPLS packet with whatever ATM cells have been received. For example, to bound the delay to 1 ms the user would configure 10 (hundreds of microseconds). The delay-time is rounded up to one of the following values 1, 5, 10, 50, 100, 200, 300 and 400.

Values 1 to 400

Default 400

control-word

Syntax [**no**] **control-word**

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe). For the 7750 SR only, ATM N:1 cell mode pseudowires (apipe vc-types atm-vcc and atm-vpc) and VT pseudowire (apipe vc-type atm-cell).

The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.

The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an “Illegal C-bit” status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes. The control word must be enabled to allow MPLS-TP OAM to be used on a static spoke-sdp in a Apipe, Epipe and Cpipe service.

pw-path-id

Syntax [**no**] **pw-path-id**

Context config>service>epipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>vpls>spoke-sdp

Description This command enters the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID.

For an IES or VPRN spoke-sdp, the pw-path-id is only valid for ethernet spoke-sdps.

The **pw-path-id** is only configurable if all of the following is true:

- SDP signaling is off
- control-word is enabled (control-word is disabled by default)
- the service type is Epipe, VPLS, Cpipe, Apipe, or IES/VPRN interface
- mate SDP signaling is off for vc-switched services

The **no** form of the command deletes the PW path ID.

Default no pw-path-id

agi

Syntax **agi** *agi*
no agi

Context config>service>epipe>spoke-sdp>pw-path-id
config>service>cpipe>spoke-sdp>pw-path-id
config>service>apipe>spoke-sdp>pw-path-id
config>service>vpls>spoke-sdp>pw-path-id

Description This command configures the attachment group identifier for an MPLS-TP PW.

Parameters *agi* — Specifies the attachment group identifier.

Values 0 to 4294967295

saii-type2

Syntax **saii-type2** *global-id:node-id:ac-id*
no saii-type2

Context config>service>epipe>spoke-sdp>pw-path-id

config>service>cpipe>spoke-sdp>pw-path-id
config>service>apipe>spoke-sdp>pw-path-id
config>service>vpls>spoke-sdp>pw-path-id

Description This command configures the Source Individual Attachment Identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the taii-type2 of the mate spoke-sdp.

Parameters *global-id* — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 to 4294967295

node-id — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 to 4294967295

ac-id — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 to 4294967295

taii-type2

Syntax **taii-type2** *global-id:node-id:ac-id*
no taii-type2

Context config>service>epipe>spoke-sdp>pw-path-id
config>service>cpipe>spoke-sdp>pw-path-id
config>service>apipe>spoke-sdp>pw-path-id
config>service>vpls>spoke-sdp>pw-path-id

Description This command configures the Target Individual Attachment Identifier (TAII) for an MPLS-TP spoke-SDP. If this is configured on a spoke-SDP for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the saii-type2 of the mate spoke-SDP.

Parameters *global-id* — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 to 4294967295

node-id — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 to 4294967295

ac-id — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 to 4294967295

control-channel-status

Syntax [**no**] **control-channel-status**

Context config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>vpls>spoke-sdp

Description This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF.

A control-channel-status no shutdown is allowed only if all of the following are true:

- SDP signaling is off.
- The control-word is enabled (the control-word is disabled by default)
- The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN
- Mate SDP signaling is off (in vc-switched services)
- The pw-path-id is configured for this spoke-SDP.

The **no** form of this command removes control channel status signaling from a spoke-SDP. It can only be removed if control channel status is shut down.

Default no control-channel-status

acknowledgment

Syntax [**no**] **acknowledgment**

Context config>service>cpipe>spoke-sdp>control-channel-status
config>service>epipe>spoke-sdp>control-channel-status
config>service>apipe>spoke-sdp>control-channel-status
config>service>vpls>spoke-sdp>control-channel-status

Description This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

refresh-timer

Syntax **refresh-timer** *value*

no refresh-timer

Context config>service>epipe>spoke-sdp>control-channel-status
config>service>cpipe>spoke-sdp>control-channel-status
config>service>apipe>spoke-sdp>control-channel-status
config>service>vpls>spoke-sdp>control-channel-status

Description This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.

Default no refresh-timer

Parameters *value* — Specifies the refresh timer value, in seconds.

Values 10 to 65535

Default 0 (off)

request-timer

Syntax **request-timer** *timer1* **retry-timer** *timer2* **timeout-multiplier** *multiplier*
no request-timer

Context config>service>cpipe>spoke-sdp>control-channel-status
config>service>epipe>spoke-sdp>control-channel-status
config>service>apipe>spoke-sdp>control-channel-status
config>service>vpls>spoke-sdp>control-channel-status

Description This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.

Parameters *timer1* — Specifies the interval, in seconds, at which pseudowire status messages, including a reliable delivery TLV with the “request” bit set, are sent.

Values 10 to 65535

timer2 — specifies the timeout interval, in seconds, if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

Values 0, 3 to 60

multiplier — If a requesting node does not receive a valid response to a pseudowire status request within a number of seconds equal to the retry timer multiplied by this multiplier, then it will assume the pseudowire is down. This parameter is optional.

Values 3 to 20

egress

Syntax **egress**

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command configures the egress SDP context.

hash-label

Syntax **hash-label** [**signal-capability**]
no hash label

Context config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>epipe>spoke-sdp

Description This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to any MPLS type encapsulated SDP, as well as to a VPRN service using the **auto-bind-tunnel** with the **resolution-filter** set to any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to 1 to indicate that.

To allow for applications whereby the egress LER infers the presence of the Hash Label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note however that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets that are generated in CPM and forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the Hash Label is set to a value of 0.

The **no** form of this command disables the use of the hash label.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-SDP or mesh SDP, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The 7450 ESS or 7750 SR local PE will insert the flow label interface parameters sub-TLV with F=1 in the PW ID FEC element in the label mapping message for that spoke-SDP or mesh SDP.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the PW but must not insert the hash label in the user and control packets over that spoke-SDP or mesh SDP. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-SDP or mesh SDP at the remote PE, the PW packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-SDP or mesh SDP at the remote PE, the PW received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the 7450 ESS or 7750 SR must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.

The **no** form of this command disables the use of the hash label.

Default no hash-label

Parameters **signal-capability** — Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

ignore-oper-down

Syntax **ignore-oper-down**
[**no**] **ignore-oper-down**

Context config>service>epipe>sap>

Description An ePipe service will not transition to Oper State: Down when a SAP fails and when this optional command is configured under that specific SAP. Only a single SAP in an ePipe may have this optional command included.

Default no ignore-oper-down

ingress

Syntax **ingress**

Context config>service>fpipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp

Description This command configures the ingress SDP context.

filter

Syntax **filter** [**ip** *ip-filter-id*]
no filter

Context config>service>fpipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>ingress

Description This command associates an IP filter policy with an ingress or egress Service Distribution Point (SDP). Filter policies control the forwarding and dropping of packets based on IP matching criteria. Only one filter can be applied to a spoke-SDP at a time.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress spoke-SDP. The *ip-filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

The **no** form of this command removes any configured filter ID association with the SDP. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **the scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters **ip** — Keyword indicating the filter policy is an IP filter.

ip-filter-id — The filter name acts as the ID for the IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

qos

Syntax **qos** *network-policy-id* **port-redirect-group** *queue-group-name* [**instance** *instance-id*]
no qos

Context config>service>apipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>egress
config>service>ipipe>spoke-sdp>egress
config>service>vpls>spoke-sdp>egress
config>service>vpls>mesh-sdp>egress
config>service>pw-template>egress

Description This command is used to redirect PW packets to an egress port queue-group for the purpose of shaping.

The egress PW shaping provisioning model allows the mapping of one or more PWs to the same instance of queues, or policers and queues, that are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only, or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface that the PW packets can be forwarded on. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different PWs to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-sdp inside a service, or to the egress context of a PW template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use queues only, or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-sdp to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on. This queue can be a queue-group queue or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a PW packet.

2. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-sdp to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on.
3. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports that have network IP interfaces. The handling of this is dealt with in the data path as follows:
   - When a PW packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.
   - When a PW packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the PW packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.
4. If a network QoS policy is applied to the egress context of a PW, any PW FC that is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the PW is redirected to exists and the redirection succeeds, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/DSCP/EXP is performed according to the relevant mappings of the {FC, profile} in the egress context of the network QoS policy applied to the PW. This is true regardless of whether an instance of the queue-group exists or not on the egress port the PW packet is forwarded to. If the packet's profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet's DEI/dot1p and the tunnel's DEI/dot1p/EXP, and the DSCP/prec will be remarked if **enable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the PW is redirected does not exist, the redirection command is failed. In this case, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface the PW packet is forwarded to.

The **no** version of this command removes the redirection of the PW to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

queue-group-name — Specifies the name of the queue group template up to 32 characters in length.

instance-id — Specifies the optional identification of a specific instance of the queue-group.

Values 1 to 40960

qos

Syntax **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
no qos

Context config>service>apipe>spoke-sdp>ingress
config>service>cpipe>spoke-sdp>ingress
config>service>epipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>ingress
config>service>vpls>spoke-sdp>ingress
config>service>vpls>mesh-sdp>ingress
config>service>pw-template>ingress

Description This command is used to redirect PW packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress PW rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more PWs to the same instance of policers that are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress PW shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally for each traffic type (unicast or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface that the PW packets can be received on. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step which means the same network QoS policy can redirect different PWs to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-sdp inside a service, or to the ingress context of a PW template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress PW rate-limiting feature:

1. When a PW FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the PW packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
2. When a PW FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the PW packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
3. When a PW FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs that have network IP interfaces. The handling of this is dealt within the data path as follows:
   - When a PW packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as “policer-output-queues”.
   - When a PW packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the PW packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
4. If a network QoS policy is applied to the ingress context of a PW, any PW FC that is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly into the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
5. If no network QoS policy is applied to the ingress context of the PW, then all packets of the PW will feed:
   - the ingress network shared queue for the packet’s FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.
   - a queue-group policer followed by the per-FP ingress shared queues, referred to as “policer-output-queues”, if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group. The only exceptions to this behavior are for packets received from an IES/VPRN spoke interface and from an R-VPLS spoke-sdp that is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet’s FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.

When a PW is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to the default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the PW. This is true regardless of whether an instance of the named policer queue-group exists on the ingress FP the pseudowire packet is received on. The user can apply a QoS filter matching the dot1-p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload’s IP header if the user enabled the ler-use-dscp option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to the default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface the pseudowire packet is received on.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification on the system.

Values 1to 65535

queue-group-name — Specifies the name of the queue group template up to 32 characters in length.

instance-id — Specifies the identification of a specific instance of the queue-group.

Values 1to 16384

vc-label

Syntax [**no**] **vc-label** *vc-label*

Context config>service>fpipe>spoke-sdp>egress
config>service>apipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>egress
config>service>ipipe>spoke-sdp>egress
config>service>apipe>spoke-sdp>ingress
config>service>cpipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>ingress

Description This command configures the egress and ingress VC label.

The actual maximum value that can be configured is limited by the **config>router>mpls-labels>static-label-range** command.

Parameters *vc-label* — A VC egress value that indicates a specific connection.

Values for egress: 16 to 1048575

Values for ingress: 32 to 18431

monitor-oper-group

Syntax **monitor-oper-group** *group-name*
no monitor-oper-group

Context config>service>epipe>spoke-sdp
config>service>epipe>sap

Description This command specifies the operational group to be monitored by the object under which it is configured. The **oper-group** *name* must be already configured under the **config>service** context before its name is referenced in this command.

The **no** form of the command removes the association.

Default none

Parameters *group-name* — Specifies an oper group name.

oper-group

Syntax **oper-group** *group-name*
no oper-group

Context config>service>epipe>sap

Description This command configures the operational group identifier.

The no form of the command removes the group name from the configuration.

Default none

Parameters *group-name* — Specifies the Operational-Group identifier up to 32 characters in length.

precedence

Syntax **precedence** [*precedence-value* | **primary**]
no precedence

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>epipe>spoke-sdp

Description This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of the command returns the precedence value to the default.

Default 4

Parameters *precedence-value* — Specifies the spoke-SDP precedence.

Values 1 to 4

primary — Assigns primary precedence to the spoke-SDP.

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>epipe>spoke-sdp

Description This command enables pseudowire status signaling for this spoke-SDP binding.

The **no** form of the command disables the status signaling.

Default pw-status-signaling

use-sdp-bmac

Syntax [**no**] **use-sdp-bmac**

Context config>service>epipe>spoke-sdp

Description This command indicates that this spoke-SDP is expected to be part of a redundant pseudowire connected to a PBB Epipe service. Enabling this parameter will cause traffic forwarded from this spoke-SDP into the B-VPLS domain to use a virtual backbone MAC as its source MAC address when both this, and the control pseudowire, are in the active state on this BEB. This virtual backbone MAC is derived from the SDP source-bmac-lsb configuration.

This command will fail when configuring it under a spoke-SDP within a PBB-Epipe that is connected to a B-VPLS with mac-notification enabled.

Default no use-sdp-bmac

vc-label

Syntax [**no**] **vc-label** *vc-label*

Context config>service>cpipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>ingress
config>service>epipe>spoke-sdp>ingress

Description This command configures the egress and ingress VC label.

The actual maximum value that can be configured is limited by the **config>router>mpls-labels>static-label-range** command.

Parameters *vc-label* — A VC egress value that indicates a specific connection.

Values for egress: 16 to 1048575

Values for ingress: 32 to 18431

vlan-vc-tag

Syntax **vlan-vc-tag** 0 to 4094
no vlan-vc-tag [*0* to 4094]

Context config>service>epipe>spoke-sdp

Description This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command.

Default no vlan-vc-tag

Parameters *0 to 4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

spoke-sdp-fec

Syntax **spoke-sdp-fec**
spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**]
spoke-sdp-fec *spoke-sdp-fec-id* **no-endpoint**
spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**] **endpoint** *name* [**icb**]

Context config>service>epipe

Description This command binds a service to an existing Service Distribution Point (SDP), using a dynamic MS-PW.

A spoke-SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

When using dynamic MS-PWs, the particular SDP to bind-to is automatically selected based on the Target Attachment Individual Identifier (TAII) and the path to use, specified under spoke-SDP FEC. The selected SDP will terminate on the first hop S-PE of the MS-PW. Therefore, an SDP must already be defined in the config>service>sdp context that reaches the first hop router of the MS-PW. The router will in order to associate an SDP with a service. If an SDP to that is not already configured, an error message is generated. If the sdp-id does exist, a binding between that sdp-id and the service is created.

It differs from the spoke-sdp command in that the spoke-sdp command creates a spoke-SDP binding that uses a pseudowire with the PW ID FEC. However, the spoke-sdp-fec command enables pseudowires with other FEC types to be used. Only the Generalized ID FEC (FEC129) may be specified using this command.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default none

Parameters *spoke-sdp-fec-id* — An unsigned integer value identifying the spoke-SDP.

Values 1 to 4294967295

fec-type — An unsigned integer value for the type of the FEC used by the MS-PW.

Values 129 to 130

aii-type — An unsigned integer value for the Attachment Individual Identifier (AII) type used to identify the MS-PW endpoints.

Values 1 to 2

endpoint-name — Specifies the name of the service endpoint.

no endpoint — Adds or removes a spoke-SDP association.

icb — Configures the spoke-SDP as an inter-chassis backup SDP binding.

auto-config

Syntax [**no**] **auto-config**

Context config>service>epipe>spoke-sdp-fec

Description This command enables single sided automatic endpoint configuration of the spoke-SDP. The router acts as the passive T-PE for signaling this MS-PW.

Automatic Endpoint Configuration allows the configuration of a spoke-SDP endpoint without specifying the TAII associated with that spoke-SDP. It allows a single-sided provisioning model where an incoming label mapping message with a TAII that matches the SAII of that spoke-SDP to be automatically bound to that endpoint. In this mode, the far end T-PE actively initiates MS-PW signaling and will send the initial label mapping message using T-LDP, while the router T-PE for which auto-config is specified will act as the passive T-PE.

The **auto-config** command is blocked in CLI if signaling active has been enabled for this spoke-SDP. It it is only applicable to spoke-SDPs configured under the Epipe, IES and VPRN interface context.

The **no** form of the command means that the router T-PE either acts as the active T-PE (if signaling active is configured) or automatically determines which router will initiate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke-SDP. If the SAII has the greater prefix value, then the router will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the router will assume that the far end T-PE will initiate MS-PW signaling and will wait for that label mapping message before responding with a T-LDP label mapping message for the MS-PW in the reverse direction.

Default no auto-config

path

Syntax **path** *name*
no path

Context config>service>epipe>spoke-sdp-fec

Description This command specifies the explicit path, containing a list of S-PE hops, that should be used for this spoke-SDP. The path-name should correspond to the name of an explicit path configured in the **config>service>pw-routing** context.

If no path is configured, then each next-hop of the MS-PW used by the spoke-SDP will be chosen locally at each T-PE and S-PE.

Default no path

Parameters *name* — The name of the explicit path to be used, as configured under **config>service>pw-routing**.

precedence

Syntax **precedence** *prec-value*
precedence primary
no precedence

Context config>service>epipe>spoke-sdp-fec

Description This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of the command returns the precedence value to the default.

Default 42

Parameters *prec-value* — Specifies the spoke-SDP precedence.

Values 1 to 4

primary — Assigns primary precedence to this spoke-SDP.

pw-template-bind

Syntax **pw-template-bind** *policy-id*
no pw-template-bind

Context config>service>epipe>spoke-sdp-fec

Description This command binds includes the parameters included in a specific PW Template to a spoke-SDP.

The **no** form of the command removes the values from the configuration.

Default none

Parameters *policy-id* — Specifies the existing policy ID.

Values 1 to 2147483647

retry-count

Syntax **retry-count** *retry-count*
no retry-count

Context config>service>epipe>spoke-sdp-fec

Description This optional command specifies the number of attempts software should make to reestablish the spoke-SDP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the spoke-sdp is put into the shutdown state.

Use the no shutdown command to bring up the path after the retry limit is exceeded.

The **no** form of this command reverts the parameter to the default value.

Default 30

Parameters *retry-count* — The maximum number of retries before putting the spoke-sdp into the shutdown state.

Values 10 to 10000

retry-timer

Syntax **retry-timer** *retry-timer*
no retry-timer

Context config>service>epipe>spoke-sdp-fec

Description This command specifies a retry-timer for the spoke-SDP. This is a configurable exponential back-off timer that determines the interval between retries to reestablish a spoke-SDP if it fails and a label withdraw message is received with the status code “AII unreachable”.

The **no** form of this command reverts the timer to its default value.

Default 30

Parameters *retry-timer* — The initial retry-timer value in seconds.

Values 10 to 480

saii-type2

Syntax **saii-type2** *global-id:prefix:ac-id*
no saii-type2

Context config>service>epipe>spoke-sdp-fec

Description This command configures the source attachment individual identifier for the spoke-sdp. This is only applicable to FEC129 AII type 2.

Parameters *global-id* — A Global ID of this router T-PE. This value must correspond to one of the global_id values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.

Values 1 to 4294967295

prefix — The prefix on this router T-PE that the spoke-sdp SDP is associated with.This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix context.**

Values an IPv4-formatted address a.b.c.d or 1 to 4294967295

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 to 4294967295

signaling

Syntax **signaling** *signaling*

Context config>service>epipe>spoke-sdp-fec

Description This command enables a user to configure this router as the active pr passive T-PE for signaling this MS-PW, or to automatically select whether this T-PE is active or passive based on the prefix. In an active role, this endpoint initiates MS-PW signaling without waiting for a T-LDP label mapping message to arrive from the far end T-PE. In a passive role, it will wait for the initial label mapping message from the far end before sending a label mapping for this end of the PW. In auto mode, if the SAII has the greater prefix value, then the router will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the router will assume that the far end T-PE will initiate MS-PW signaling and will wait for that label mapping message before responding with a T-LDP label mapping message for the MS-PW in the reverse direction.

The **no** form of the command means that the router T-PE automatically selects the which router will initiate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke-SDP, as previously described.

Default auto

Parameters *signaling* — Configures this router as the active T-PE for signaling this MS-PW.

Values auto, master

standby-signaling-slave

Syntax [**no**] **standby-signaling-slave**

Context config>service>epipe>spoke-sdp-fec

Description This command enables standby-signaling-slave for an Epipe.

taii-type2

Syntax **taii-type2** *global-id:prefix:ac-id*
no taii-type2

Context config>service>epipe>spoke-sdp-fec

Description taii-type2 configures the target attachment individual identifier for the spoke-sdp. This is only applicable to FEC129 AII type 2.

This command is blocked in CLI if this end of the spoke-SDP is configured for single-sided auto configuration (using the **auto-config** command).

Parameters *global-id* — A Global ID of this router T-PE. This value must correspond to one of the global_id values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.

Values 1 to 4294967295

prefix — The prefix on this router T-PE that the spoke-sdp SDP is associated with.This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix context.**

Values an IPv4-formatted address a.b.c.d or 1 to 4294967295

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 to 4294967295

2.17.2.10 ATM Commands

atm

Syntax **atm**

Context config>service>epipe>sap
config>service>apipe>sap
config>service>ipipe>sap
config>service>epipe>sap

Description This command enables access to the context to configure ATM-related attributes.This command can only be used when a specified context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a specified context, the command returns an error.

egress

Syntax **egress**

Context config>service>epipe>sap
config>service>epipe>sap>atm

config>service>apipe>sap>atm
config>service>fpipe>sap

Description This command configures egress ATM attributes for the SAP.

ingress

Syntax **ingress**

Context config>service>epipe>sap
config>service>epipe>sap>atm
config>service>epipe>sap
config>service>apipe>sap>atm

Description This command configures ingress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>epipe>sap>atm

Description This command specifies the data encapsulation for an ATM PVCC delimited Epipe SAP. The definition references RFC 2684, *Multiprotocol Encapsulation over ATM AAL5,* and to the ATM Forum LAN Emulation specification. Ingress traffic that does not match the configured encapsulation will be dropped.

Default aal5snap-bridged

Parameters *atm-encap-type* — Specifies the encapsulation type.

Values **aal5snap-bridged** — Bridged encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.

aal5mux-bridged-eth-nofcs — Bridged IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>ipipe>sap>atm

Description This command specifies the data encapsulation for an ATM PVCC delimited Ipipe SAP. The definition references RFC 2684, *Multiprotocol Encapsulation over ATM AAL5,* and to the ATM Forum LAN Emulation specification. Ingress traffic that does not match the configured encapsulation will be dropped.

Default aal5snap-routed

Parameters *atm-encap-type* — Specifies the encapsulation type.

Values **aal5snap-routed** — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.

aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>epipe>sap
config>service>apipe>sap>atm>egress
config>service>apipe>sap>atm>ingress
config>service>epipe>sap>atm>egress
config>service>epipe>sap>atm>ingress

Description This command assigns an ATM traffic descriptor profile to a specified context (for example, a SAP).

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specifies a defined traffic descriptor profile (see the QoS atm-td-profile command).

2.17.2.11 OAM Commands

oam

Syntax **oam**

Context config>service>epipe>sap
config>service>apipe>sap>atm

Description This command enters the context to configure OAM functionality for a PVCC delimiting a SAP.

- The ATM-capable MDAs support end-to-end and segment OAM functionality (AIS, RDI, Loopback) over both F5 (VC) and end-to-end F4 (VP) OAM:
- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM N3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>epipe>sap>oam
config>service>epipe>sap>oam
config>service>apipe>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC terminations to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.

When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes operationally down, or enters a fault state and becomes operationally up, or exits that fault state). RDI cells are generated when PVCC is operationally down. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The **no** command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, a PVCC's operational status is no longer affected by a PVCC's OAM state changes due to AIS/RDI processing (when **alarm-cells** is disabled, a PVCC will change operational status to operationally up due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as previously described.

Default Enabled for PVCCs delimiting IES SAPs

terminate

Syntax [**no**] **terminate**

Context config>service>apipe>sap>atm>oam

Description This command specifies whether this SAP will act as an OAM termination point. ATM SAPs can be configured to tunnel or terminate OAM cells.

When configured to not terminate (the default is **no terminate**), the SAP will pass OAM cells through the VLL without inspecting them. The SAP will respond to OAM loopback requests that are directed to the local node by transmitting a loopback reply. Other loopback requests are transparently tunneled through the pseudowire. In this mode, it is possible to launch a loopback request toward the directly-attached ATM equipment and see the results of the reply.

When configured to terminate, the SAP will respond to AIS by transmitting RDI and will signal the change of operational status to the other endpoint (for example, through LDP status notifications). The SAP will respond to OAM loopback requests by transmitting a loopback reply. In this mode, it is possible to launch a loopback request toward the directly-attached ATM equipment and see the results of the reply.

For Apipe services, the user has the option of enabling or disabling this option for VC types atm-vcc and atm-sdu since these service types maintain the ATM layer and/or the AAL5 layer across the VLL. It is not supported on atm-vpc and atm-cell apipe vc types since the VLL must pass the VC level (F5) OAM cells.

The **terminate** option for OAM is the only and default mode of operation supported for an ATM SAP which is part of Epipe, Ipipe, VPLS, and IES/VPRN. This is because the ATM and AAL5 layers are terminated.

For Apipe services, the user has the option of enabling or disabling this option for vc types atm-vcc and atm-sdu since these service types maintain the ATM layer and/or the AAL5 layer across the VLL. It is not supported on atm-vpc and atm-cell Apipe vc types since the VLL must pass the VC level (F5).

The **terminate** option for OAM is the only and default mode of operation supported for an ATM SAP which is part of Epipe, Ipipe, VPLS, and IES/VPRN. This is because the ATM and AAL5 layers are terminated.

Default no terminate

2.17.2.12 Cpipe Commands

endpoint

Syntax [**no**] **endpoint** *endpoint-name*

Context config>service>cpipe

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name.

active-hold-delay

Syntax **active-hold-delay** *active-hold-delay*
no active-hold-delay

Context config>service>cpipe>endpoint

Description This command specifies that the node will delay sending the change in the T-LDP status bits for the service endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from “active” to “standby” or when any object in the endpoint. For example., SAP, ICB, or regular spoke-SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from “active” to “standby”, the node sends immediately new T-LDP status bits indicating the new value of “standby” over the spoke-SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from “standby” to “active” or when any object in the endpoint transitions to an operationally up state.

Default 0 — A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from “active” to “standby”, the node sends immediately new T-LDP status bits indicating the new value of “standby” over the spoke-SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

Parameters *active-hold-delay* — Specifies the active hold delay in 100s of milliseconds.

Values 0 to 60

revert-time

Syntax **revert-time** *revert-time*
revert-time infinite
no revert-time

Context config>service>cpipe>endpoint

Description This command configures the time to wait before reverting back to the primary spoke-SDP defined on this service endpoint, after having failed over to a backup spoke-SDP.

Parameters *revert-time* — Specifies the time, in seconds, to wait before reverting to the primary SDP.

Values 0 to 600

infinite — Causes the endpoint to be non-revertive.

2.17.2.13 VLL SAP Commands

sap

Syntax **sap** *sap-id* [**no-endpoint**] [**create**]
sap *sap-id* **endpoint** *endpoint-name* [**create**]
no sap *sap-id*

Context config>service>cpipe

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the service router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config router interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded.

The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted.

Default No SAPs are defined.

Special Cases **VLL Services —** A SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. At most, only one sdp-id can be bound to an VLL service. Since a VLL is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. Up to 49 SDPs can be associated with a service in a single router. Each SDP must have a unique router destination or an error will be generated.

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services. This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0).

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

port-id — Specifies the physical port ID.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number format. For example 61/2/3 specifies port 3 on MDA 2 in slot 61.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

port-id	*slot*/*mda*/*port* [.*channel*]		
	eth-sat-id	esat-*id*/*slot*/*port*	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-*id*.*sub-port*	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

endpoint — Adds a SAP endpoint association.

no endpoint — Removes the association of a SAP or a spoke-sdp with an explicit endpoint name.

create — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

cem

Syntax **cem**

Context config>service>cpipe>sap

Description This command enters the context to specify circuit emulation (CEM) properties.

packet

Syntax **packet jitter-buffer** *milliseconds* [**payload-size** *bytes*]
packet payload-size *bytes*
no packet

Context config>service>cpipe>sap

Description This command specifies the jitter buffer size, in milliseconds, and payload size, in bytes.

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots.

Table 20 **Default CEM SAP Endpoint Types**

Endpoint Type	Timeslots	Default Jitter Buffer (in ms)
unstructuredE1	n/a	5
unstructuredT1	n/a	5
nxDS0 (E1/T1)	N = 1	32
	N = 2 to 4	16
	N = 5 to 15	8
	N >= 16	5
nxDS0WithCas (E1)	N	8
nxDS0WithCas (T1)	N	12

Parameters *milliseconds* — specifies the jitter buffer size in milliseconds (ms).

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffers is not allowed.

Setting the jitter butter value to 0 sets it back to the default value.

Values 1 to 250

payload-size *bytes* — Specifies the payload size (in bytes) of packets transmitted to the packet service network (PSN) by the CEM SAP. This determines the size of the data that will be transmitted over the service. If the size of the data received is not consistent with the payload size, then the packet is considered malformed.

Values 0, 16 to 2048

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots.

Table 21 **Payload Size CEM SAP Endpoint Types**

Endpoint Type	Timeslots	Default Payload Size (in bytes)
unstructuredE1	n/a	256
unstructuredT1	n/a	192

Table 21 **Payload Size CEM SAP Endpoint Types (Continued)**

Endpoint Type	Timeslots	Default Payload Size (in bytes)
nxDS0 (E1/T1)	N = 1	64
	N = 2 to 4	N x 32
	N = 5 to 15	N x 16
	N >= 16	N x 8
nxDS0WithCas (E1)	N	N x 16
nxDS0WithCas (T1)	N	N x 24

For all endpoint types except for nxDS0WithCas, the valid payload size range is from the default to 2048 bytes.

For nxDS0WithCas, the payload size divide by the number of timeslots must be an integer factor of the number of frames per trunk multi-frame (for example, 16 for E1 trunk and 24 for T1 trunk).

For 1xDS0, the payload size must be a multiple of 2.

For NxDS0, where N > 1, the payload size must be a multiple of the number of timeslots.

For unstructuredE1 and unstructuredT1, the payload size must be a multiple of 32 bytes.

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffer is not allowed.

Setting the payload size to 0 sets it back to the default value.

report-alarm

Syntax [**no**] **report-alarm** [**stray**] [**malformed**] [**pktloss**] [**overrun**] [**underrun**] [**rpktloss**] [**rfault**] [**rrdi**]

Context config>service>cpipe>sap>cem

Description This command indicates the type of CEM SAP alarm.

The **no** form of the command removes the parameter from the configuration.

Parameters **stray** — Reports the reception of packets not destined for this CES circuit.

malformed — Reports the reception of packet not properly formatted as CES packets.

pktloss — Reports the lack of reception of CES packets.

overrun — Reports the reception of too many CES packets resulting in a overrun of the receive jitter buffer.

underrun — Reports the reception of too few CES packets resulting in a overrun of the receive jitter buffer.

rpktloss — Reports hat the remote peer is currently in packet loss status.

rfault — Reports that the remote TDM interface is currently not in service.

rrdi — Reports that the remote TDM interface is currently in RDI status.

rtp-header

Syntax [**no**] **rtp-header**

Context config>service>cpipe>sap>cem

Description This command specifies whether an RTP header is used when packets are transmitted to the packet service network (PSN) by the CEM SAP.

Default no rtp-header

2.17.2.14 CPipe SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**no-endpoint**] [**create**]
spoke-sdp *sdp-id:vc-id* [**create**] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>cpipe

Description This command binds a service to an existing Service Distribution Point (SDP). A spoke-SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.
SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default No *sdp-id* is bound to a service.

Parameters *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 to 17407 for existing SDPs.

vc-id — The virtual circuit identifier.

Values 1 to 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

The VC type value for a VPLS service is defined as 0x000B.

Values ethernet

no endpoint — Removes the association of a spoke-SDP with an explicit endpoint name.

endpoint-name — Specifies the name of the service endpoint.

icb — Configures the spoke-SDP as an inter-chassis backup SDP binding.

egress

Syntax **egress**

Context config>service>cpipe>spoke-sdp

Description This command enters the context to configure egress spoke-SDP context.

ingress

Syntax **ingress**

Context config>service>cpipe>spoke-sdp

Description This command enters the context to configure ingress spoke-SDP context.

vc-label

Syntax **vc-label** *egress-vc-label*
no vc-label [*egress-vc-label*]

Context config>service>cpipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>ingress

Description This command configures the spoke-SDP egress and ingress VC label.

The actual maximum value that can be configured is limited by the **config>router>mpls-labels>static-label-range** command.

Parameters *egress-vc-label* — A VC egress value that indicates a specific connection.

Values for egress: 16 to 1048575

Values for ingress: 32 to 18431

precedence

Syntax **precedence** [*precedence-value* | **primary**]
no precedence

Context config>service>cpipe>spoke-sdp

Description This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of the command returns the precedence value to the default.

Default 4

Parameters *precedence-value* — Specifies the spoke-SDP precedence.

Values 1 to 4

primary — Specifies to make this the primary spoke-SDP.

2.17.2.15 Epipe SAP Template Commands

template

Syntax **template**

Context config>service

Description This is the node for service templates.

epipe-sap-template

Syntax **epipe-sap-template** *name* [**create**]
no epipe-sap-template *name*

Context config>service>template

Description This command specifies which SAP parameter template should be applied to the l2-ap SAP. This can only be changed when the l2-ap is shutdown.

The no form of the command removes the template, the SAP will use default parameters.

Default None

Parameters *name* — Specifies the SAP template name associated with this template.

egress

Syntax **egress**

Context config>service>template

Description This command enters the context to configure egress filter policies.

ingress

Syntax **ingress**

Context config>service>template>epipe-sap-template

Description This command enters the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

filter

Syntax [**no**] **filter**

Context config>service>template>epipe-sap-template>egress
config>service>template>epipe-sap-template>ingress

Description This command enters the context to configure filter parameters.

ip

Syntax **ip** *filter-id*
no ip

Context config>service>template>epipe-sap-template>egress>filter
config>service>template>epipe-sap-template>ingress>filter

Description This command associates an existing IP filter policy with the template.

Parameters *filter-id* — Specifies the IP filter policy ID. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6

Syntax **ipv6** *filter-id*
no ipv6

Context config>service>template>epipe-sap-template>egress>filter
config>service>template>epipe-sap-template>ingress>filter

Description This command associates an existing IPv6 filter policy with the template.

Parameters *ipv6-filter-id* — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac

Syntax **mac** *filter-id*
no mac

Context config>service>template>epipe-sap-template>egress>filter
config>service>template>epipe-sap-template>ingress>filter

Description This command associates an existing MAC filter policy with the template.

Parameters *mac-filter-id* — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

qos

Syntax **qos** *policy-id*
no qos

Context config>service>template>epipe-sap-template>egress

Description This command associates an existing QoS policy with the template.

Parameters *policy-id* — The egress policy ID to associate with SAP or IP interface on egress. The policy ID must already exist.

Values 1 to 65535, or a name up to 64 characters in length

qos

Syntax **qos** *policy-id* {**shared-queuing** | **multipoint-shared**}
qos *policy-id*
no qos

Context config>service>template>epipe-sap-template>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) for the Epipe SAP template.

Default none

Parameters *policy-id* — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.

Values 1 to 65535

shared-queuing — This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared — This keyword can only be specified on SAP ingress. Multipoint shared queuing is a superset of shared queuing. When multipoint shared queuing keyword is set, as well as the unicast packets, multipoint packets also used shared queues.

Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice, similar to the shared-queuing option. In addition, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

Values Multipoint or not present.

Default Present (the queue is created as non-multipoint).

2.18 VLL Show Command Reference

This section describes the VLL show command reference.

2.18.1 Command Hierarchies

2.18.1.1 Show Commands

show
- **service**
  - **egress-label** *start-label* [*end-label*]
  - **id** *service-id*
    - **all**
    - **authentication**
    - **base** [**msap**]
    - **bgp-vpws**
    - **endpoint** [*endpoint-name*]
    - **labels**
    - **retailers**
    - **sap** *sap-id* **detail**
    - **sdp** [*sdp-id*[:*vc-id*]] | **far-end** {*ip-address* | *ipv6-address*}] [**mrp**] [**detail**]]
    - **spoke-sdp-fec** [[1 to 4294967295]]
    - **stp** [**detail**]
    - **stp** **mst-instance** *mst-instance-number*
    - **vccv-bfd** [**session**]
    - **wholesalers**
  - **ingress-label** *start-label* [*end-label*]
  - **sap-using** [*msap*] [**dyn-script**] [**description**]
  - **sap-using** [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**] [**description**]
  - **sap-using** {**ingress** | **egress**} **atm-td-profile** *td-profile-id*
  - **sap-using** {**ingress** | **egress**} **filter** *any-filter-id*
  - **sap-using** {**ingress** | **egress**} **qos-policy** *qos-policy-id* [**msap**]
  - **sap-using** **etree**
  - **sdp** *sdp-id* **pw-port** [*pw-port-id*]
  - **sdp** *sdp-id* **pw-port**
  - **sdp** *sdp-id* **pw-port** *pw-port-id* [**statistics**]
  - **sdp** [**consistent** | **inconsistent** | **na**] **egressifs**
  - **sdp** *sdp-id* **keep-alive-history**
  - **sdp** **far-end** {*ip-address* | *ipv6-address*} **keep-alive-history**
  - **sdp** [*sdp-id*] **detail**
  - **sdp** **far-end** {*ip-address* | *ipv6-address*} **detail**
  - **sdp-using** **etree**
  - **sdp-using** **node-id** *node-id* [**global-id** *global-id*]
  - **sdp-using** arrp *arrpID*
  - **sdp-using** **app-profile** *app-profile-name*
  - **sdp-using** **far-end** {*ip-address* | *ipv6-address*}

```
            — sdp-using [sdp-is[:vc-id]]
            — sdp-using transit-policy ip transit-ip-policy
            — sdp-using transit-policy prefix transit-prefix-policy
            — service-using [epipe] [ies] [vpls] [vprn] [mirror] [apipe] [fpipe] [ipipe] [cpipe] [etree]
                [vsd] [b-vpls] [i-vpls] [m-vpls] [sdp sdp-id] [customer customer-id] [origin
                creation-origin]
            — spoke-sdp-fec-using [spoke-sdp-fec-id spoke-sdp-fec-id] [saii-type2 global-
                id:prefix:ac-id] [taii-type2 global-id:prefix:ac-id] [path name] [expired]
        — pw-port [pw-port-id] [detail]
        — pw-port sdp sdp-id
        — pw-port none
        — pw-port pw-port-id statistics

show
    —  connection-profile [1 to 8000]
```

2.18.1.2 Clear Commands

```
clear
    — service
        — id service-id
            — arp
            — host-tracking [statistics]
            — host-tracking sap sap-id [host ip-address] [statistics]
            — mesh-sdp sdp-id[:vc-id] ingress-vc-label
            — mesh-sdp sdp-id[:vc-id] vccv-bfd {session | statistics}
            — spoke-sdp sdp-id:vc-id ingress-vc-label
            — spoke-sdp sdp-id:vc-id l2tpv3
            — spoke-sdp sdp-id:vc-id vccv-bfd {session | statistics}
        — statistics
            — id service-id
                — counters
                — spoke-sdp sdp-id[:vc-id] {all | counters | stp | 12pt | mrp}
            — sap sap-id {all | cem | counters | stp | 12pt | mrp}
            — sap sap-id encap-group group-name [member encap-id]
            — sdp sdp-id keep-alive
            — sdp sdp-id pw-port [1 to 10239]
```

2.18.1.3 Debug Commands

```
debug
    — service
        — id service-id
            — [no] sap sap-id
                — [no] event-type {arp | config-change | oper-status-change |
                    neighbor -discovery}
            — [no] sdp sdp-id:vc-id
```

— [no] event-type {arp | config-change | oper-status-change | neighbor -discovery}

2.18.1.4 Tools Commands

```
tools
    — dump
        — epipe-map-access-to-egress-port service service-id [end-service service-id]
        — epipe-map-access-to-egress-port lag lag-id summary
```

2.18.2 Command Descriptions

2.18.2.1 VLL Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

egress-label

Syntax **egress-label** *egress-label1* [*egress-label2*]

Context show>service

Description This command displays services using the range of egress labels. If only the mandatory *egress-label1* parameter is specified, only services using the specified label are displayed.

If both *egress-label1* and *egress-label2* parameters are specified, the services using the range of labels X where *egress-label1* <= X <= *egress-label2* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *egress-label1* — The starting egress label value for which to display services using the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0, 2049 to 131071

egress-label2 — The ending egress label value for which to display services using the label range.

Default The *egress-label1* value.

Values 2049 to 131071

Output The following output is an example of egress label information, and Table 22 describes the output fields.

Sample Output

```
*A:ALA-12# show service egress-label 0 10000
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          100:1              Mesh 0            0
...
1          107:1              Mesh 0            0
1          108:1              Mesh 0            0
1          300:1              Mesh 0            0
1          301:1              Mesh 0            0
1          302:1              Mesh 0            0
1          400:1              Mesh 0            0
1          500:2              Spok 131070       2001
1          501:1              Mesh 131069       2000
100        300:100            Spok 0            0
200        301:200            Spok 0            0
300        302:300            Spok 0            0
400        400:400            Spok 0            0
-------------------------------------------------------------------------------
Number of Bindings Found : 23
===============================================================================
*A:ALA-12#
```

Table 22 **Show Service Egress Label Output Fields**

Label	Description
Svc Id	The ID that identifies a service.
Sdp Id	The ID that identifies an SDP.
Type	Indicates whether the SDP binding is spoke or mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.
Number of bindings found	The total number of SDP bindings that exist within the specified egress label range.

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description This command displays services using the range of ingress labels. If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router vprn-service-id ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting ingress label value for which to display services using the label range. If only *start-label* is specified, services only using *start-label* are displayed.

Values 0, 18432 to 524287

end-label — The ending ingress label value for which to display services using the label range.

Values 18432 to 524287

Default The *start-label* value.

Output The following output is an example of ingress label information, and Table 23 describes the output fields.

Sample Output

```
*A:ALA-12# show service ingress-label 0
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          50:1               Mesh 0            0
1          100:1              Mesh 0            0
1          101:1              Mesh 0            0
1          102:1              Mesh 0            0
1          103:1              Mesh 0            0
1          104:1              Mesh 0            0
1          105:1              Mesh 0            0
1          106:1              Mesh 0            0
1          107:1              Mesh 0            0
1          108:1              Mesh 0            0
1          300:1              Mesh 0            0
1          301:1              Mesh 0            0
1          302:1              Mesh 0            0
1          400:1              Mesh 0            0
100        300:100            Spok 0            0
200        301:200            Spok 0            0
```

```
300         302:300             Spok 0             0
400         400:400             Spok 0             0
-------------------------------------------------------------------------------
Number of Bindings Found : 21
-------------------------------------------------------------------------------
*A:ALA-12#
```

Table 23 **Show Service Ingress-Label Output Fields**

Label	Description
Svc ID	The service identifier.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of Bindings Found	The number of SDP bindings within the label range specified.

sap-using

Syntax **sap-using** [**msap**] [**dyn-script**] [**description**]
sap-using [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**] [**description**]
sap-using {**ingress** | **egress**} **atm-td-profile** *td-profile-id*
sap-using {**ingress** | **egress**} **filter** *any-filter-id*
sap-using {**ingress** | **egress**} **qos-policy** *qos-policy-id* [**msap**]
sap-using etree

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress** — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

qos-policy-id — The ingress or egress QoS Policy ID for which to display matching SAPs.

Values 1 to 65535

td-profile-id — Displays SAPs using this traffic description for the 7750 SR only.

filter-id — The ingress or egress filter policy ID for which to display matching SAPs.

Values 1 to 65535

sap-id — Specifies the physical port identifier portion of the SAP definition.

dyn-script — Displays dynamic service SAPs information.

msap — Displays MSAPs.

vlan-translation — Displays VLAN translation information.

anti-spoof — Displays anti-spoof information.

Output The following output is an example if SAP using information, and Table 24 describes the output fields.

Sample Output

```
*A:Dut-A# show service sap-using

===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.   Ing.    Egr.   Egr.    Adm  Opr
                                           QoS    Fltr    QoS    Fltr
-------------------------------------------------------------------------------
1/1/1:1                         1          1      none    1      none    Up   Up
2/1/2:10/11                     1          1      none    1      none    Up   Up
2/1/2:10/12                     1          1      none    1      none    Up   Up
2/1/2:20/11                     1          1      none    1      none    Up   Up
2/1/2:20/12                     1          1      none    1      none    Up   Up
2/1/4:cp.10                     10         1      none    1      none    Up   Up
2/1/4:cp.20                     20         1      none    1      none    Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 7
-------------------------------------------------------------------------------
===============================================================================
```

The following is a sample of SAP information for a specific SAP for the 7450 ESS or 7750 SR:

```
A:ALA-42#
*A:ALA-48# show service sap-using sap 1/1/21:0
===============================================================================
Service Access Points Using Port 1/1/21:0
===============================================================================
PortId                  SvcId      Ing.   Ing.    Egr.   Egr.    Anti    Adm  Opr
                                   QoS    Fltr    QoS    Fltr    Spoof
-------------------------------------------------------------------------------
1/1/21:0                1          1      none    1      none    none    Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
===============================================================================
*A:ALA-48#
```

Following is a sample of SAP information for the egress traffic policy for the 7750 SR.

```
*A:ALA-12#  show service sap-using egress atm-td-profile 2
===============================================================================
Service Access Point Using ATM Traffic Profile 2
===============================================================================
PortId SvcId I.QoS I.Fltr E.QoS E.Fltr A.Pol Adm Opr
-------------------------------------------------------------------------------
5/1/1:0/11 511111 2 none 2 none none Up Up
5/1/1:0/12 511112 2 none 2 none none Up Up
5/1/1:0/13 511113 2 none 2 none none Up Up
5/1/1:0/14 511114 2none 2 none none Up Up
5/1/1:0/15 511115 2 none 2 none none Up Up
5/1/1:0/16 511116 2 none 2 none none Up Up
5/1/1:0/17 511117 2 none 2 none none Up Up
5/1/1:0/18 511118 2 none 2 none none Up Up
5/1/1:0/19 511119 2 none 2 none none Up Up
5/1/1:0/20 511120 2 none 2 none none Up Up
5/1/1:0/21 511121 2 none 2 none none Up Up
5/1/1:0/22 511122 2 none 2 none none Up Up
5/1/1:0/23 511123 2 none 2 none none Up Up
5/1/1:0/24 511124 2 none 2 none none Up Up
5/1/1:0/25 511125 2 none 2 none none Up Up
...
===============================================================================
*A:ALA-12#
```

Table 24 **Show Service SAP Output Fields**

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
Sap MTU	The SAP MTU value.
Ing. QoS	The SAP ingress QoS policy number specified on the ingress SAP.
Ing Fltr	The MAC or IP filter policy ID applied to the ingress SAP.
Egr. QoS	The SAP egress QoS policy number specified on the egress SAP for the 7450 ESS and 7750 SR only.
Egr. Fltr	The MAC or IP filter policy ID applied to the egress SAP.
Adm	The administrative state of the SAP.
Opr	The operational state of the SAP.

sdp-using

Syntax **sdp-using etree**

sdp-using node-id *node-id* [**global-id** *global-id*]
sdp-using aarp *aarpID*
sdp-using app-profile *app-profile-name*
sdp-using far-end {*ip-address* | *ipv6-address*}
sdp-using [*sdp-id*[:*vc-id*]]
sdp-using transit-policy ip *transit-ip-policy*
sdp-using transit-policy prefix *transit-prefix-policy*

Context show>service

Description Display services using SDP or far-end address options.

Parameters *node-id* — Specifies the node ID.

Values a.b.c.d, 1 to 4294967295

global-id — Specifies the global ID.

Values 1 to 4294967295

aarpID — Specifies the AARP ID.

Values 1 to 65535

app-profile-name — 32 characters max.

sdp-id — Displays only services bound to the specified SDP ID.

Values 1 to 17407

vc-id — The virtual circuit identifier.

Values 1 to 4294967295

ip-address — Displays only services matching with the specified far-end IP address. 64 characters maximum.

Default Services with any far-end IP address.

ipv6-address — Displays only services matching with the specified far-end IPv6 address. 64 characters maximum.

transit-ip-policy — Specifies a transit IP policy ID.

Values 1 to 65535

transit-prefix-policy — Specifies a transit prefix policy ID.

Values 1 to 65535

Output The following output is an example of SDP using information, and Table 25 describes the output fields.

Sample Output

```
*A:ALA-1# show service sdp-using 300
===============================================================================
Service Destination Point (Sdp Id : 300)
```

```
===============================================================================
SvcId       SdpId               Type Far End        Opr State I.Label  E.Label
-------------------------------------------------------------------------------
1           300:1               Mesh 10.0.0.13      Up        131071   131071
2           300:2               Spok 10.0.0.13      Up        131070   131070
100         300:100             Mesh 10.0.0.13      Up        131069   131069
101         300:101             Mesh 10.0.0.13      Up        131068   131068
102         300:102             Mesh 10.0.0.13      Up        131067   131067
-------------------------------------------------------------------------------
Number of SDPs : 5
-------------------------------------------------------------------------------
*A:ALA-1#
```

Table 25 **Show Service SDP Using Output Fields**

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.
Type	Type of SDP: spoke or mesh.
Far End	The far end address of the SDP.
Oper State	The operational state of the service.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] [**apipe**] [**fpipe**] [**ipipe**] [**cpipe**] [**etree**] [**b-vpls**] [**i-vpls**] [**m-vpls**] [**sdp** *sdp-id*] [**customer** *customer-id*] [**creation** *creation-origin*]

Context show>service

Description This command displays the services matching certain usage properties. Not all syntax options are available for each router type.

If no optional parameters are specified, all services defined on the system are displayed.

Parameters **epipe** — Displays epipe services.

ies — Displays IES services.

vpls — Displays VPLS services.

vprn — Displays VPRN services.

mirror — Displays mirror services.

apipe — Displays Apipe services.

fpipe — Displays Fpipe services.

ipipe — Displays Ipipe services.

cpipe — Displays Cpipe services.

b-vpls — Specifies the B-component instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature. It represents the multi-point tunneling component that multiplexes multiple customer VPNs (ISIDs) together. It is similar to a regular VPLS instance that operates on the backbone MAC addresses.

i-vpls — Specifies the I-component instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature. It identifies the specific VPN entity associated to a customer multipoint (E-LAN) service. It is similar to a regular VPLS instance that operates on the customer MAC addresses.

m-vpls — Specifies the M-component (managed VPLS) instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature.

sdp-id — Displays only services bound to the specified SDP ID.

Values 1 to 17407

Default Services bound to any SDP ID.

customer-id — Displays services only associated with the specified customer ID.

Values 1 to 2147483647

Default Services associated with any customer.

creation-origin — Specifies the method by which the service was created.

Values manual, multi-segment-p-w, dyn-script, vsd

Output The following output is an example of service using information, and Table 26 describes the output fields.

Sample Output

```
*A:ALA-12# show service service-using customer 10
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            VPLS      Up     Up         10                09/05/2006 13:24:15
100          IES       Up     Up         10                09/05/2006 13:24:15
300          Epipe     Up     Up         10                09/05/2006 13:24:15
-------------------------------------------------------------------------------
Matching Services : 3
===============================================================================
*A:ALA-12#


*A:ALA-12# show service service-using
===============================================================================
```

```
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            uVPLS     Up     Up         1                 10/26/2006 15:44:57
2            Epipe     Up     Down       1                 10/26/2006 15:44:57
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
999          uVPLS     Down   Down       1                 10/26/2006 16:14:33
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
*A:ALA-12#
```

The following is a sample of epipe service information for the 7450 ESS or 7750 SR.

```
*A:ALA-12# show service service-using epipe
===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
6            Epipe     Up     Up         6                 06/22/2006 23:05:58
7            Epipe     Up     Up         6                 06/22/2006 23:05:58
8            Epipe     Up     Up         3                 06/22/2006 23:05:58
103          Epipe     Up     Up         6                 06/22/2006 23:05:58
-------------------------------------------------------------------------------
Matching Services : 4
===============================================================================
*A:ALA-12#
```

Table 26 **Show Service-Using Output Fields**

Label	Description
Service Id	The service identifier.
Type	Specifies the service type configured for the service ID.
Adm	The desired state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

spoke-sdp-fec-using

Syntax **spoke-sdp-fec-using** [**spoke-sdp-fec-id** *spoke-sdp-fec-id*] [**saii-type2** *global-id:prefix:ac-id*] [**taii-type2** *global-id:prefix:ac-id*] [**path** *name*] [**expired**]

Context show>service

Description Displays the SDPs used by spoke-sdp-fecs at this node.

Parameters *spoke-sdp-fec-id* — Specifies the spoke-SDP FEC ID for which to show SDP information.

Values 1 to 4294967295

global-id — Specifies the SAII or TAII global ID.

Values 1 to 4294967295

prefix — Specifies the SAII or TAII prefix.

Values a.b.c.d, 1 to 4294967295

ac-id — Specifies the SAII or TAII AC ID.

Values 1 to 4294967295

name — Specifies the path name. 32 characters maximum.

expired — Displays information for expired SDPs.

Output The following output is an example of spoke-SDP information.

Sample Output

```
*A:Dut-C# show service spoke-sdp-fec-using
===============================================================================
Service Spoke-SDP-Fec Information
===============================================================================
SvcId      SpokeSdpFec  Oper-SdpBind     SAII-Type2
 Path                                     TAII-Type2
-------------------------------------------------------------------------------
1          1            17407:4294967245 3:10.20.1.3:1
 n/a                                      6:10.20.1.6:1
2          2            17407:4294967247 3:10.20.1.3:2
 n/a                                      6:10.20.1.6:2
3          3            17407:4294967248 3:10.20.1.3:3
 n/a                                      6:10.20.1.6:3
4          4            17407:4294967249 3:10.20.1.3:4
 n/a                                      6:10.20.1.6:4
5          5            17407:4294967250 3:10.20.1.3:5
 n/a                                      6:10.20.1.6:5
6          6            17407:4294967251 3:10.20.1.3:6
 n/a                                      6:10.20.1.6:6
7          7            17407:4294967252 3:10.20.1.3:7
 n/a                                      6:10.20.1.6:7
8          8            17407:4294967253 3:10.20.1.3:8
 n/a                                      6:10.20.1.6:8
9          9            17407:4294967254 3:10.20.1.3:9
 n/a                                      6:10.20.1.6:9
10         10           17407:4294967255 3:10.20.1.3:10
```

```
 n/a                                        6:10.20.1.6:10
-------------------------------------------------------------------------------
Entries found: 10
```

vccv-bfd

Syntax **vccv-bfd** [**session**]

Context show>service>id

Description This command shows whether VCCV BFD is configured for a particular service and information about the VCCV session state.

The **show>service>id>vccv-bfd session** command gives a summary of all the VCCV sessions. Using both the sdp-id and the vc-id parameters gives VCCV BFD session information about a specific spoke-SDP.

For services where auto-discovery and signaling are used (for example, BGP VPWS, VPLS, and BGP-AD VPLS), use the **show>service>id>detail** command to determine the sdp-id and vc-id parameters allocated by the system.

Parameters **session** — Displays a summary of all VCCV sessions.

Output The following output is an example of VCCV BFD information.

Sample Output

```
*A:Dut-C# show service id 1000 vccv-bfd session
===============================================================================
BFD Session
===============================================================================
Interface/Lsp Name              State               Tx Intvl  Rx Intvl  Multipl
Remote Address/Info         Protocols             Tx Pkts   Rx Pkts   Type
LAG port/sdp-id:vc-id      LAG ID/SvcId
-------------------------------------------------------------------------------
N/A                           Up (3)                1000      1000      3
N/A                           vccv                  152       151       central
100:100                       1000
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **endpoint** | **fdb** | **interface** | **labels** | **sap** | **sdp** | **split-horizon-group** | **stp**}

Context show>service

Description This command displays information for a particular service-id.

Parameters *service-id* — The service identification number that identifies the service in the domain.

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters in length.

all — Displays detailed information about the service.

arp — Displays ARP entries for the service.

base — Displays basic service information.

endpoint — Displays service endpoint information.

interface — Displays service interfaces.

labels — Displays labels being used by this service.

sap — Displays SAPs associated to the service.

sdp — Displays SDPs associated with the service.

split-horizon-group — Displays split horizon group information.

stp — Displays STP information.

Output The following output is an example of service ID information.

Sample Output

```
A:bksim1611>config>service>ipipe# show service id 1009 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1009               Vpn Id            : 0
Service Type      : Ipipe
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 09/15/2010 13:06:46 Last Mgmt Change  : 09/15/2010 13:06:02
Admin State       : Up                 Oper State        : Up
MTU               : 1500
Vc Switching      : False
SAP Count         : 1                  SDP Bind Count    : 1
CE IPv4 Discovery : Enabled            CE IPv6 Discovery : Enabled
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 5:1009  -(5.5.5.5)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 5:1009                  Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp    : (Not Specified)
VC Type            : Ipipe                   VC Tag            : 0
Admin Path MTU     : 0                       Oper Path MTU     : 1568
Far End            : 5.5.5.5                 Delivery          : MPLS
Tunnel Far End     : n/a
Hash Label         : Disabled
```

```
Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 131048                   Egress Label      : 131053
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)     : 0
Last Status Change : 09/15/2010 13:06:46      Signaling         : TLDP
Last Mgmt Change   : 09/15/2010 13:06:02
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel

KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
Max Drop Count     : 3                        Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 15                       I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 1460                     I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 17                       E. Fwd. Octets    : 1604

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP LIST :
Lsp Name           : to-bksim180-1
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 16h07m44s

Lsp Name           : to-bksim180-2
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 16h07m45s

-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Enabled                  EnforceDSTELspFc  : Disabled
Default LSP        : to-bksim180-1            Multicast LSP     : None

===============================================================================
FC Mapping Table
===============================================================================
FC Name              LSP Name
-------------------------------------------------------------------------------
ef                   to-bksim180-2
===============================================================================

-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IP Addr  : n/a                  Peer CE IP Addr   : 0.0.0.0
```

```
Peer IPv6 Capability   : No
Peer IPv6 LL Addr      :  FE80::2009:2009:2
Peer IPv6 Global Addr  :  3FFE:1200:2009:2009:9:9:9:8

------------------------------------------------------------------------------
Number of SDPs : 1
------------------------------------------------------------------------------
------------------------------------------------------------------------------
Service Access Points
------------------------------------------------------------------------------

------------------------------------------------------------------------------
SAP 1/7/3:1009
------------------------------------------------------------------------------
Service Id         : 1009
SAP                : 1/7/3:1009              Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 09/15/2010 13:06:21
Last Mgmt Change   : 09/15/2010 13:06:02
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1518                    Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
Ing Agg Rate Limit : max                     Egr Agg Rate Limit: max
Endpoint           : N/A
Q Frame-Based Acct : Disabled

Acct. Pol          : None                    Collect Stats     : Disabled

Oper Group         : (none)                  Monitor Oper Grp  : (none)

------------------------------------------------------------------------------
ETH-CFM SAP specifics
------------------------------------------------------------------------------
Tunnel Faults      : n/a                     CFM Hold-Timer    : n/a
------------------------------------------------------------------------------
Ipipe SAP Configuration Information
------------------------------------------------------------------------------
Configured CE IP   : n/a                     Discovered CE IP  : 209.1.1.1
SAP MAC Address    : ac:55:01:07:00:03       Mac Refresh Inter*: 14400

------------------------------------------------------------------------------
Ipipe SAP IPv4 ARP Entry Info
------------------------------------------------------------------------------
209.1.1.1                           00:11:22:33:44:55 dynamic

------------------------------------------------------------------------------
Ipipe SAP IPv6 Neighbor Entry Info
------------------------------------------------------------------------------
FE80::2009:2009:1                   00:11:22:33:44:55 dynamic
```

```
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 2                        172
Off. HiPrio           : 0                        0
Off. LowPrio          : 17                       1978
Off. Uncolor          : 0                        0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 17                       1978

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 15                       1790
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                  Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 17                       1978
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 17                       1978

Egress Queue 1
For. InProf           : 0                        0
For. OutProf          : 15                       1790
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.

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

```
VPLS Sites
===============================================================================
Site                Site-Id   Dest                Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
show service id x all
-------------------------------------------------------------------------------
SAP 1/1/4:500
-------------------------------------------------------------------------------
Service Id         : 500
SAP                : 1/1/4:500               Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Down
Flags              : PortOperDown
Multi Svc Site     : None
Last Status Change : 09/19/2013 11:43:04
Last Mgmt Change   : 09/19/2013 11:43:05
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1518                    Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
                                             Egr Agg Rate Limit: max
Endpoint           : N/A
Q Frame-Based Acct : Disabled
Vlan-translation   : None

Acct. Pol          : None                    Collect Stats     : Disabled

Application Profile: None
Transit Policy     : None

Oper Group         : (none)                  Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Ignore Oper Down   : Disabled
Lag Link Map Prof  : (none)
Cflowd             : Disabled

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : n/a                     AIS               : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : 0 1 2 3 4 5 6 7

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
.
.
.
-------------------------------------------------------------------------------
```

```
Service Destination Points(SDPs)
------------------------------------------------------------------------------
------------------------------------------------------------------------------
 Sdp Id 1:2  -(1.1.1.1)
------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id          : 1:2                       Type             : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp    : (Not Specified)
VC Type         : Ether                     VC Tag           : n/a
Admin Path MTU  : 0                         Oper Path MTU    : 0
Delivery        : GRE
Far End         : 1.1.1.1
Tunnel Far End  : n/a                       LSP Types        : n/a
Hash Label      : Disabled                  Hash Lbl Sig Cap : Disabled
Oper Hash Label    : Disabled

Admin State     : Up                        Oper State       : Down
Acct. Pol       : None                      Collect Stats    : Disabled
Ingress Label   : 0                         Egress Label     : 0
Ingr Mac Fltr-Id   : n/a                    Egr Mac Fltr-Id  : n/a
Ingr IP Fltr-Id    : n/a                    Egr IP Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                    Egr IPv6 Fltr-Id : n/a
Admin ControlWord  : Not Preferred          Oper ControlWord : False
Last Status Change : 09/11/2013 20:02:40    Signaling        : TLDP
Last Mgmt Change   : 09/15/2013 13:56:56    Force Vlan-Vc    : Disabled
Endpoint        : N/A                       Precedence       : 4
PW Status Sig   : Enabled
Class Fwding State : Down
Flags           : SdpOperDown
                  NoIngVCLabel NoEgrVCLabel
                  PathMTUTooSmall
Time to RetryReset : never                  Retries Left     : 3
Mac Move        : Blockable                 Blockable Level  : Tertiary
Local Pw Bits   : None
Peer Pw Bits    : None
Peer Fault Ip   : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None
Transit Policy     : None
Max Nbr of MAC Addr: No Limit               Total MAC Addr   : 0
Learned MAC Addr   : 0                      Static MAC Addr  : 0
OAM MAC Addr       : 0                      DHCP MAC Addr    : 0
Host MAC Addr      : 0                      Intf MAC Addr    : 0
SPB MAC Addr       : 0                      Cond MAC Addr    : 0

MAC Learning       : Enabled                Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                  Block On Mesh Fail: False
Oper Group         : (none)                 Monitor Oper Grp  : (none)
Rest Prot Src Mac  : Disabled
Auto Learn Mac Prot: Disabled               RestProtSrcMacAct : Disable

Ingress Qos Policy : (none)                 Egress Qos Policy : (none)
```

```
Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)

-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
V-MEP Filtering    : Disabled

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Squelch Levels     : 0 1 2 3 4 5 6 7
```

authentication

Syntax **authentication**

Context show>service>id

Description This command enables the context to display subscriber authentication information.

statistics

Syntax **statistics** [**policy** *name*] [**sap** *sap-id*]

Context show>service>id>authentication

Description This command displays session authentication statistics for this service.

Parameters *name* — Specifies the subscriber authentication policy statistics to display.

sap-id — Specifies the SAP ID statistics to display.

Output The following output is an example of statistics information.

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
===============================================================
Authentication statistics
===============================================================
Interface / SAP                Authentication  Authentication
                               Successful      Failed
---------------------------------------------------------------
vpls-11-90.1.0.254              1582            3
---------------------------------------------------------------
Number of entries: 1
```

```
===============================================================
*A:ALA-1#
```

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output The following output is an example of all service ID information, and Table 27 describes the output fields.

Sample Output

```
A:SR12# show service id 10 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 10                      Vpn Id            : 0
Service Type      : Apipe                VLL Type          : ATMCell
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 2
Last Status Change: 10/07/2010 05:03:47 Last Mgmt Change  : 10/07/2010 05:03:51
Admin State       : Up                   Oper State        : Down
MTU               : 1508                 Signaling Override: ATMVCC
Vc Switching      : False
SAP Count         : 1                    SDP Bind Count    : 1

......... (No change to SDP description)

-------------------------------------------------------------------------------
SAP 2/1/4:cp.10
-------------------------------------------------------------------------------
Service Id         : 10
SAP                : 2/1/4:cp.10              Encap             : atm
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 10/16/2010 06:58:41
Last Mgmt Change   : 10/16/2010 06:58:41
Sub Type           : regular
Split Horizon Group: (Not Specified)

Admin MTU          : 1524                     Oper MTU          : 1524
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
Endpoint           : N/A
```

```
Acct. Pol          : None                      Collect Stats     : Disabled

Oper Group         : (none)                    Monitor Oper Grp  : (none)

*B:ALA-Dut-H# show service id 100 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 100                 Vpn Id            : 100
Service Type      : Epipe
Description       : Default epipe description for service id 100
Customer Id       : 100
Last Status Change: 02/06/2007 10:03:11
Last Mgmt Change  : 02/06/2007 09:43:27
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 4
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
 Sdp Id 1:10100  -(10.20.1.7)
-------------------------------------------------------------------------------
SDP Id             : 1:10100                 Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 1560                    Oper Path MTU     : 1560
Far End            : 10.20.1.7               Delivery          : MPLS

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 130065                  Egress Label      : 130368
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 02/06/2007 10:03:24     Signaling         : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint           : y                       Precedence        : 4
Flags              : SapOperDown
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Enabled                 Oper State        : Alive
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name           : lsp1_G
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 01h40m15s
-------------------------------------------------------------------------------
```

```
 Sdp Id 2:100  -(10.20.1.4)
------------------------------------------------------------------------------
SDP Id             : 2:100                     Type             : Spoke
VC Type            : Ether                     VC Tag           : n/a
Admin Path MTU     : 1560                      Oper Path MTU    : 1560
Far End            : 10.20.1.4                 Delivery         : MPLS
Admin State        : Up                        Oper State       : Up
Acct. Pol          : None                      Collect Stats    : Disabled
Ingress Label      : 130671                    Egress Label     : 130367
Ing mac Fltr       : n/a                       Egr mac Fltr     : n/a
Ing ip Fltr        : n/a                       Egr ip Fltr      : n/a
Ing ipv6 Fltr      : n/a                       Egr ipv6 Fltr    : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord : False
Last Status Change : 02/06/2007 10:03:11       Signaling        : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint           : y                         Precedence       : 0
Flags              : None
Peer Pw Bits       : pwFwdingStandby
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Enabled                   Oper State       : Alive
Hello Time         : 10                        Hello Msg Len    : 0
Max Drop Count     : 3                         Hold Down Time   : 10

Statistics           :
I. Fwd. Pkts.      : 0                         I. Dro. Pkts.    : 0
E. Fwd. Pkts.      : 0                         E. Fwd. Octets   : 0

Associated LSP LIST :
Lsp Name           : lsp2_D
Admin State        : Up                        Oper State       : Up
Time Since Last Tr*: 01h40m16s
------------------------------------------------------------------------------
 Sdp Id 3:100  -(10.20.1.5)
------------------------------------------------------------------------------
SDP Id             : 3:100                     Type             : Spoke
VC Type            : Ether                     VC Tag           : n/a
Admin Path MTU     : 1560                      Oper Path MTU    : 1560
Far End            : 10.20.1.5                 Delivery         : MPLS

Admin State        : Up                        Oper State       : Up
Acct. Pol          : None                      Collect Stats    : Disabled
Ingress Label      : 130971                    Egress Label     : 130368
Ing mac Fltr       : n/a                       Egr mac Fltr     : n/a
Ing ip Fltr        : n/a                       Egr ip Fltr      : n/a
Ing ipv6 Fltr      : n/a                       Egr ipv6 Fltr    : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord : False
Last Status Change : 02/06/2007 10:03:17       Signaling        : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint           : y                         Precedence       : 4
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel
```

```
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Enabled                 Oper State        : Alive
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0
Associated LSP LIST :
Lsp Name           : lsp3_E
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 01h40m16s
-------------------------------------------------------------------------------
...
===============================================================================
*B:ALA-Dut-H#


A:SR12# show service id 20 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 20                  Vpn Id            : 0
Service Type      : Apipe               VLL Type          : ATMCell
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 2
Last Status Change: 10/07/2010 05:03:47 Last Mgmt Change  : 10/07/2010 05:03:51
Admin State       : Up                  Oper State        : Down
MTU               : 1508                Signaling Override: ATMVCC
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
APIPE SDU-mode specifics
-------------------------------------------------------------------------------
Interworking      : None
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
Sdp Id 3:1  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 3:1                    Type              : Spoke
VC Type            : AAL5SDU                VC Tag            : 0
Admin Path MTU     : 1600                   Oper Path MTU     : 1600
Far End            : 1.1.1.1                Delivery          : GRE

Admin State        : Up                     Oper State        : Up
Acct. Pol          : None                   Collect Stats     : Disabled
Ingress Label      : 119665                 Egress Label      : 103665
Ing mac Fltr       : n/a                    Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                    Egr ip Fltr       : n/a
Admin ControlWord  : Preferred              Oper ControlWord  : True
Last Status Change : 04/04/2007 20:52:24    Signaling         : TLDP
Last Mgmt Change   : 04/04/2007 20:48:24
Endpoint           : N/A                    Precedence        : 4
Flags              : None
```

```
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : pwe3ControlWord
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
 Sdp Id 6:2  -(4.4.4.4)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 6:2                     Type              : Spoke
VC Type            : AAL5SDU                 VC Tag            : 0
Admin Path MTU     : 1600                    Oper Path MTU     : 1600
Far End            : 4.4.4.4                 Delivery          : GRE

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 103664                  Egress Label      : 119665
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Admin ControlWord  : Preferred               Oper ControlWord  : True
Last Status Change : 04/04/2007 20:53:13     Signaling         : TLDP
Last Mgmt Change   : 04/04/2007 20:48:24
Endpoint           : N/A                     Precedence        : 4
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : pwe3ControlWord mplsRouterAlertLabel
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
No Sap Associations
-------------------------------------------------------------------------------
Service Endpoints
```

```
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:ALA-DutC>config>service#

*A:ALU-76# show service id 200 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 200                 Vpn Id            : 0
Service Type      : Cpipe               VLL Type          : CESoPSN
Customer Id       : 1
Last Status Change: 09/11/2008 19:05:29
Last Mgmt Change  : 09/10/2008 19:51:06
Admin State       : Up                  Oper State        : Up
MTU               : 1400
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
 Sdp Id 5:200  -(5.5.5.5)
-------------------------------------------------------------------------------
SDP Id             : 5:200                   Type              : Spoke
VC Type            : CESoPSN                 VC Tag            : 0
Admin Path MTU     : 0                       Oper Path MTU     : 1568
Far End            : 5.5.5.5                 Delivery          : MPLS
 Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 131061                  Egress Label      : 131066
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Admin ControlWord  : Preferred               Oper ControlWord  : True
Admin BW(Kbps)     : 0                       Oper BW(Kbps)     : 0
Last Status Change : 09/11/2008 19:05:29     Signaling         : TLDP
Last Mgmt Change   : 09/10/2008 19:51:06
Endpoint           : N/A                     Precedence        : 4
Class Fwding State : Down
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : pwe3ControlWord mplsRouterAlertLabel

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                       I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name           : to-ALU-80-1
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 03d21h52m
-------------------------------------------------------------------------------
```

```
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : disabled                EnforceDSTELspFc  : disabled
Default LSP        : Uknwn                   Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
CPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Local Bit-rate     : 12                      Peer Bit-rate     : 12
Local Payload Size : 192                     Peer Payload Size : 192
Local Sig Pkts     : No Sig.                 Peer Sig Pkts     : No Sig.
Local CAS Framing  : No CAS                  Peer CAS Framing  : No CAS
Local RTP Header   : No                      Peer RTP Header   : No
Local Differential : No                      Peer Differential : No
Local Timestamp    : 0                       Peer Timestamp    : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
SAP 1/5/1.1.1.1
-------------------------------------------------------------------------------
Service Id         : 200
SAP                : 1/5/1.1.1.1             Encap             : cem
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 09/10/2008 19:51:27
Last Mgmt Change   : 09/10/2008 19:51:06
Sub Type           : regular

Admin MTU          : 1578                    Oper MTU          : 1578
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
Ing Agg Rate Limit : max                     Egr Agg Rate Limit: max
Endpoint           : N/A

Acct. Pol          : None                    Collect Stats     : Disabled
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                          Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
```

```
Off. Uncolor          : 0                       0
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
CEM SAP Configuration Information
-------------------------------------------------------------------------------
Endpoint Type          : NxDS0            Bit-rate                    : 12
Payload Size           : 192              Jitter Buffer (ms)          : 8
Jitter Buffer (packets): 4                Playout Threshold (packets): 3
Use RTP Header         : No               Differential                : No
Timestamp Freq         : 0                CAS Framing                 : No CAS

Cfg Alarm      : stray malformed pktloss overrun underrun
Alarm Status   :
-------------------------------------------------------------------------------
CEM SAP Statistics
-------------------------------------------------------------------------------
                        Packets           Seconds           Events
Egress Stats
Forwarded             :  0
Dropped               :  0
Missing               :  0
Reordered Forwarded   :  0
Underrun              :  0                                  0
Overrun               :  0                                  0
```

```
Misordered Dropped    :  0
Malformed Dropped     :  0
LBit Dropped          :  0
Multiple Dropped      :  0
Error                 :                    0
Severely Error        :                    0
Unavailable           :                    0
Failure Count         :                                  0
Jitter Buffer Depth   :  0
Ingress Stats
Forwarded             :  0
Dropped               :  0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:ALU-76#


*A:bksim180# show service id 1000 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1000                Vpn Id            : 0
Service Type      : Ipipe
Customer Id       : 1
Last Status Change: 03/11/1973 10:20:24
Last Mgmt Change  : 03/11/1973 10:20:23
Admin State       : Up                  Oper State        : Up
MTU               : 1400
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
CE Addr Discovery : enabled
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
 Sdp Id 22:1000  -(2.2.2.2)
-------------------------------------------------------------------------------
SDP Id             : 22:1000                  Type              : Spoke
VC Type            : Ipipe                    VC Tag            : 0
Admin Path MTU     : 0                        Oper Path MTU     : 1568
Far End            : 2.2.2.2                  Delivery          : MPLS

Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 131070                   Egress Label      : 131062
Ing mac Fltr       : n/a                      Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                      Egr ip Fltr       : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)     : 0
Last Status Change : 03/11/1973 10:20:24      Signaling         : TLDP
Last Mgmt Change   : 03/11/1973 10:19:21
Endpoint           : N/A                      Precedence        : 4
Class Fwding State : Down
Flags              : None
Time to RetryReset : 1999616832 seconds       Retries Left      : 2984947
Mac Move           : Ukwn                     Blockable Level   : Unknown
Peer Pw Bits       : None
```

```
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                       I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name           : to-bksim176-1
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h01m28s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : disabled
Default LSP        : Uknwn                   Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IP Addr  : n/a
                  Peer CE IP Addr   : 1.1.1.2
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
SAP 1/7/1
-------------------------------------------------------------------------------
Service Id         : 1000
SAP                : 1/7/1                   Encap             : null
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 03/11/1973 10:20:23
Last Mgmt Change   : 03/11/1973 10:19:21
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Admin MTU          : 1514                    Oper MTU          : 1514
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
Ing Agg Rate Limit : max                     Egr Agg Rate Limit: max
Endpoint           : N/A
```

```
Q Frame-Based Acct : Disabled                                                  \

Acct. Pol          : None                      Collect Stats     : Disabled
-------------------------------------------------------------------------------
Ipipe SAP Info
-------------------------------------------------------------------------------
Configured CE IP   : n/a
                     Discovered CE IP  : 1.1.1.1
SAP MAC Address    : 8c:c7:01:07:00:01         Mac Refresh Inter*: 14400
-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
1.1.1.1        8c:c7:01:07:00:03  dynamic   04h00m00s
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                         Egress qos-policy : 1
Shared Q plcy      : n/a                       Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
Off. Uncolor          : 0                        0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                  Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 0                        0
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 0                        0
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
```

```
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:bksim180#

*A:ces-A# show service id 1 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Cpipe               VLL Type          : SAToPT1
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 07/06/2010 19:21:14
Last Mgmt Change  : 07/06/2010 19:21:14
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 12:1  -(2.2.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:1                  Type              : Spoke
VC Type           : SAToPT1               VC Tag            : 0
Admin Path MTU    : 0                     Oper Path MTU     : 9190
Far End           : 2.2.2.2               Delivery          : MPLS

Admin State       : Up                    Oper State        : Up
Acct. Pol         : None                  Collect Stats     : Disabled
Ingress Label     : 131064                Egress Label      : 131064
Admin ControlWord : Preferred             Oper ControlWord  : True
Admin BW(Kbps)    : 0                     Oper BW(Kbps)     : 0
Last Status Change : 07/06/2010 19:21:14  Signaling          : TLDP
Last Mgmt Change  : 07/06/2010 19:21:14
Endpoint          : N/A                   Precedence        : 4
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : pwe3ControlWord mplsRouterAlertLabel

KeepAlive Information :
Admin State       : Enabled               Oper State        : Alive
Hello Time        : 10                    Hello Msg Len     : 0
Max Drop Count    : 3                     Hold Down Time    : 10
```

```
Statistics            :
I. Fwd. Pkts.      : 141578                  I. Fwd. Octs.     : 31430316
E. Fwd. Pkts.      : 141583                  E. Fwd. Octets    : 31431426

Associated LSP LIST :
Lsp Name           : to_b_1_2
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 04h08m22s


*A:Dut-B# show service id 1 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1                    Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                    Creation Origin   : manual
Last Status Change: 01/28/2015 22:05:35
Last Mgmt Change  : 01/28/2015 22:05:22
Test Service      : No
Admin State       : Up                   Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                    SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

-------------------------------------------------------------------------------
BGP Information
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore


-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 230:1  -(10.20.1.3)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 230:1                 Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type           : Ether                 VC Tag            : n/a
Admin Path MTU    : 0                     Oper Path MTU     : 1582
Delivery          : MPLS
Far End           : 10.20.1.3
Tunnel Far End    : n/a                   LSP Types         : SR-ISIS
Hash Label        : Disabled              Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled

Admin State       : Up                    Oper State        : Up
Acct. Pol         : None                  Collect Stats     : Disabled
Ingress Label     : 262135                Egress Label      : 262135
```

```
Ingr Mac Fltr-Id   : n/a                       Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                       Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                       Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord  : False
Admin BW(Kbps)     : 0                         Oper BW(Kbps)     : 0
BFD Template       : None
BFD-Enabled        : no                        BFD-Encap         : ipv4
Last Status Change : 01/28/2015 22:05:35       Signaling         : TLDP
Last Mgmt Change   : 01/28/2015 22:05:22
Endpoint           : N/A                       Precedence        : 4
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                  Force Qinq-Vc     : Disabled
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel

Application Profile: None
Transit Policy     : None
Standby Sig Slave  : False
Block On Peer Fault: False
Use SDP B-MAC      : False

Ingress Qos Policy : (none)                    Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                    Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                    Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State        : Disabled                  Oper State        : Disabled
Hello Time         : 10                        Hello Msg Len     : 0
Max Drop Count     : 3                         Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                         I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                         I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                         E. Fwd. Octets    : 0

-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                  Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false

-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
Squelch Levels     : None

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
```

```
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                   Multicast LSP     : None

===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/8:1.1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/8:1.1               Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 01/28/2015 22:05:22
Last Mgmt Change   : 01/28/2015 22:05:22
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
                                             Egr Agg Rate Limit: max
Endpoint           : N/A
Q Frame-Based Acct : Disabled                Limit Unused BW   : Disabled
Vlan-translation   : None

Acct. Pol          : None                    Collect Stats     : Disabled

Application Profile: None
Transit Policy     : None

Oper Group         : (none)                  Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Ignore Oper Down   : Disabled
Lag Link Map Prof  : (none)
Cflowd             : Disabled

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
```

```
-------------------------------------------------------------------------------
Tunnel Faults      : accept                   AIS              : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : None

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 2                        Egress qos-policy : 2
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                      Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
CPM Ingress           : 0                       0

Forwarding Engine Stats
Dropped               : 0                       0
Received Valid        : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats(Ingress QoS Policy 2)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 2)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
```

```
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
-------------------------------------------------------------------------------

===============================================================================
VLL Sites
===============================================================================
Site                Site-Id   Dest                Admin       Oper  Fwdr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
================================================================================
*A:Dut-B#


*A:Dut-B>config>service>sdp# show service id 1 all

================================================================================
Service Detailed Information
================================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 05/27/2015 03:08:37
Last Mgmt Change  : 05/27/2015 02:56:37
Test Service      : No
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

-------------------------------------------------------------------------------
BGP Information
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore


-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 230:1  -(10.20.1.3)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 230:1                 Type              : Spoke
Spoke Descr     : (Not Specified)
```

```
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 1582
Delivery           : MPLS
Far End            : 10.20.1.3
Tunnel Far End     : n/a                      LSP Types         : SR-OSPF
Hash Label         : Disabled                 Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled

Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 262142                   Egress Label      : 262141
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)     : 0
BFD Template       : None
BFD-Enabled        : no                       BFD-Encap         : ipv4
Last Status Change : 05/27/2015 03:08:37      Signaling         : TLDP
Last Mgmt Change   : 05/27/2015 02:56:37
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                 Force Qinq-Vc     : Disabled
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel

Application Profile: None
Transit Policy     : None
Eth Seg Name       : <none>
Standby Sig Slave  : False
Block On Peer Fault: False
Use SDP B-MAC      : False

Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
Max Drop Count     : 3                        Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0

------------------------------------------------------------------------------
Control Channel Status
------------------------------------------------------------------------------
PW Status          : disabled                 Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false
```

```
-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
Squelch Levels     : None

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated

-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                   Multicast LSP     : None

===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings

-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
OSPF               : enabled                 LSP Id            : 524289
Oper Instance Id   : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/8:1.1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/8:1.1               Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 05/27/2015 02:56:37
Last Mgmt Change   : 05/27/2015 02:56:37
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Eth Seg Name       : <none>
Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
```

```
                                                Egr Agg Rate Limit: max
Endpoint           : N/A
Q Frame-Based Acct : Disabled                   Limit Unused BW   : Disabled
Vlan-translation   : None

Acct. Pol          : None                       Collect Stats     : Disabled

Application Profile: None
Transit Policy     : None

Oper Group         : (none)                     Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Ignore Oper Down   : Disabled
Lag Link Map Prof  : (none)
Cflowd             : Disabled

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : accept                     AIS               : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : None

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 2                          Egress qos-policy : 2
Ingress FP QGrp    : (none)                     Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                     Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                        Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                   Octets
CPM Ingress           : 0                         0

Forwarding Engine Stats
Dropped               : 0                         0
Received Valid        : 0                         0
Off. HiPrio           : 0                         0
Off. LowPrio          : 0                         0
Off. Uncolor          : 0                         0
Off. Managed          : 0                         0

Queueing Stats(Ingress QoS Policy 2)
Dro. HiPrio           : 0                         0
Dro. LowPrio          : 0                         0
For. InProf           : 0                         0
For. OutProf          : 0                         0

Queueing Stats(Egress QoS Policy 2)
Dro. InProf           : 0                         0
Dro. OutProf          : 0                         0
```

```
For. InProf           : 0                      0
For. OutProf          : 0                      0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                Octets
-------------------------------------------------------------------------------
Sap per Policer stats
-------------------------------------------------------------------------------
                        Packets                Octets

Ingress Policer 1 (Stats mode: minimal)
Off. All              : 0                      0
Dro. All              : 0                      0
For. All              : 0                      0

Egress Policer 1 (Stats mode: minimal)
Off. All              : 0                      0
Dro. All              : 0                      0
For. All              : 0                      0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
-------------------------------------------------------------------------------

===============================================================================
VLL Sites
===============================================================================
Site                Site-Id   Dest                Admin       Oper   Fwdr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
*A:Dut-B>config>service>sdp#
```

Table 27 describes the Show service ID output fields when the **all** option is specified.

Table 27 **Show Service ID Output Fields**

Label	Description
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	Specifies the type of service.
VLL Type	Specifies the VLL type.
SDP Id	The SDP identifier for the 7450 ESS or 7750 SR.
Description	Generic information about the service.
Customer Id	The customer identifier.

Table 27 **Show Service ID Output Fields (Continued)**

Label	Description (Continued)
Last Mgmt Change	The date and time of the most recent management-initiated change.
Endpoint	Specifies the name of the service endpoint for the 7450 ESS or 7750 SR.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch,RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service for the 7450 ESS or 7750 SR.
Split Horizon Group Specifics	
Split Horizon Group	Name of the split horizon group for this VPLS for the 7450 ESS or 7750 SR.
Description	Description of the split horizon group for the 7450 ESS or 7750 SR.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group for the 7450 ESS or 7750 SR.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier for the 7450 ESS or 7750 SR.
Type	Indicates whether this Service SDP binding is a spoke or a mesh for the 7450 ESS or 7750 SR.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented for the 7450 ESS or 7750 SR.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented for the 7450 ESS or 7750 SR.

Table 27 **Show Service ID Output Fields (Continued)**

Label	Description (Continued)
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS for the 7450 ESS or 7750 SR.
Admin State	The administrative state of this SD for the 7450 ESS or 7750 SR.
Oper State	The operational state of this SDP for the 7450 ESS or 7750 SR.
Jitter Buffer (packets)	Indicates the jitter buffer length in number of packet buffers for the 7450 ESS or 7750 SR.
Playout Threshold (packets)	Indicates the playout buffer packets threshold in number of packet buffers for the 7450 ESS or 7750 SR.
Playout Threshold (packets)	Indicates the current packet depth of the jitter buffer for the 7450 ESS or 7750 SR.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire for the 7450 ESS or 7750 SR. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the above failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signaling method to indicate faults. pwNotForwarding — Pseudowire not forwarding lacIngressFault Local — Attachment circuit RX fault lacEgressFault Local — Attachment circuit TX fault psnIngressFault Local — PSN-facing PW RX fault psnEgressFault Local — PSN-facing PW TX fault pwFwdingStandby — Pseudowire in standby mode
Signaling Override	Indicates the overriding signaled pseudowire type, as configured under the **signaled-vc-type-override** option for Apipes. This field is only displayed if **signaled-vc-type-override** is configured for the 7750 SR.

base

Syntax **base** [**msap**]

Context show>service>id

Description Displays basic information about the service ID including service type, description, SAPs and SDPs.

Parameters **msap** — Displays MSAPs.

Output The following output is an example of base service ID information, and Table 28 describes the output fields.

Sample Output

```
*A:ALA-48>config>service>epipe>sap# show service id 6 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 6                   Vpn Id            : 6
Service Type      : Epipe
Description       : Distributed Epipe service to east coast
Customer Id       : 6
Last Status Change: 02/02/2009 09:27:55
Last Mgmt Change  : 02/02/2009 09:27:57
Admin State       : Up                  Oper State        : Down
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access and Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm   Opr
-------------------------------------------------------------------------------
sap:1/2/9:0                             q-tag        1518    1518    Up    Down
sdp:2:6 S(10.10.10.104)                 n/a          0       0       Up    Down
===============================================================================
*A:ALA-48>config>service>epipe>sap#
```

Table 28 **Show Service-ID Base Output Fields**

Label	Description
Service Id	The service identifier.
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	The type of service: Epipe, Apipe, Fpipe, Ipipe, VPLS, IES, VPRN.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The desired state of the service.
Oper	The operating state of the service.

Table 28 **Show Service-ID Base Output Fields (Continued)**

Label	Description (Continued)
Mtu	The largest frame size (in octets) that the service can handle.
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service.
SDP Bind Count	The number of SDPs bound to the service.
Identifier	Specifies the service access (SAP) and destination (SDP) points.
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
AdmMTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented.
PBB Tunnel Point	Specifies the endpoint in the B-VPLS environment where the Epipe terminates.
Admin MTU	Specifies the B-VPLS admin MTU.
Backbone-Flooding	Specifies whether or not the traffic is flooded in the B-VPLS for the destination instead of unicast. If the backbone destination MAC is in the B-VPLS FDB, then it will be unicast.
ISID	The 24 bit field carrying the service instance identifier associated with the frame. It is used at the destination PE as a demultiplexor field.

bgp-vpws

Syntax **bgp-vpws**

Context show>service>id

Description This command displays BGP VPWS related information for the service.

Output The following output is an example of BGP VPWS information.

Sample Output

```
*A:cses-E11>config>service>epipe>bgp-vpws# show service id 2 bgp-vpws
===============================================================================
BGP VPWS Information
```

```
===============================================================================
Admin State          : Enabled
VE Name              : PE1                  VE Id            : 1
PW Template          : 2
Route Dist           : 65536:3
Rte-Target Import    : 65536:2              Rte-Target Export: 65536:2

PW-Template Id       : 2
Import Rte-Tgt       : None
-------------------------------------------------------------------------------
Remote-Ve Information
-------------------------------------------------------------------------------
Remote VE Name       : PE2                  Remote VE Id     : 2
===============================================================================
*A:cses-E11>config>service>epipe>bgp-vpws#
```

endpoint

Syntax **endpoint** [*endpoint-name*]

Context show>service>id

Description This command displays service endpoint information.

Parameters *endpoint-name* — Specifies the name of an existing endpoint for the service.

Output The following output is an example of service endpoint information.

Sample Output

```
*A:ALA-48>config>service>epipe# show service id 6 endpoint
===============================================================================
Service 6 endpoints
===============================================================================
Endpoint name      : x
Revert time        : 0
Act Hold Delay     : 0
Tx Active          : none
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
No members found.
===============================================================================
Endpoint name      : y
Revert time        : 0
Act Hold Delay     : 0
Tx Active          : none
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
No members found.
===============================================================================
*A:ALA-48>config>service>epipe#
```

labels

Syntax **labels**

Context show>service>id

Description Displays the labels being used by the service.

Output The following output is an example of service label information, and Table 29 describes the output fields.

Sample Output

```
*A:ALA-12# show service id 1 labels
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          40:1               Mesh 130081       131061
1          60:1               Mesh 131019       131016
1          100:1              Mesh 0            0
-------------------------------------------------------------------------------
Number of Bound SDPs : 6
-------------------------------------------------------------------------------
*A:ALA-12#
```

Table 29 **Show Service-ID Labels Output Fields**

Label	Description
Svc Id	The service identifier.
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.

retailers

Syntax **retailers**

Context show>service>id

Description This command displays the service ID of the retailer subscriber service to which this DHCP lease belongs.

wholesalers

Syntax **wholesalers**

Context show>service>id

Description This command displays service wholesaler information.

connection-profile

Syntax **connection-profile** [*1 to 8000*]

Context show

Description This command displays connection profile information.

Parameters *conn-prof-id* — Specifies the connection profile ID.

Values 1 to 8000

Output The following output is an example of connection profile information.

Sample Output

```
*A:Dut-A# show connection-profile
===============================================================================
Connection Profile Summary Information
===============================================================================
CP Index  Number of
          Members
-------------------------------------------------------------------------------
10        2
20        2
===============================================================================
*A:Dut-A#


*A:Dut-A# show connection-profile 10
===============================================================================
Connection Profile 10 Information
===============================================================================
Description : (Not Specified)
Last Change : 10/16/2010 06:53:30
-------------------------------------------------------------------------------
VPI/VCI
-------------------------------------------------------------------------------
10/11
10/12
===============================================================================
```

```
*A:Dut-A#


*A:Dut-A#
*A:bksim2801# show connection-profile
===============================================================================
Connection Profile Summary Information
===============================================================================
CP Index  Number of
          Members
-------------------------------------------------------------------------------
5000      0
===============================================================================
*A:bksim2801#


*A:bksim2801# show connection-profile 2
===============================================================================
Connection Profile 2 Information
===============================================================================
Description : (Not Specified)
Last Change : 09/09/2010 07:55:28
-------------------------------------------------------------------------------
VPI/VCI
-------------------------------------------------------------------------------
2/102
10/100
===============================================================================
*A:bksim2801#
```

sap

Syntax **sap** *sap-id* [**detail**]

Context show>service>id

Description This command displays information for the SAPs associated with the service. If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters *sap-id* — The ID that displays SAPs for the service in the form *slot*/*mda*/*port*[**.***channel*].

detail — Displays detailed information for the SAP.

Output The following output is an example of SAP information, and Table 30 describes the output fields.

Sample Output

```
*B:Dut-A# show service id 10 sap 2/1/4:cp.10
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 10
SAP                : 2/1/4:cp.10              Encap             : atm
Description        : Default sap description for service id 10
```

```
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 11/01/2010 11:33:16
Last Mgmt Change   : 11/01/2010 13:46:15
===============================================================================


A:SR12# configure service apipe 1 sap
  - no sap <sap-id>
  - sap <sap-id> [create] [no-endpoint]
  - sap <sap-id> [create] endpoint <endpoint-name>

 <sap-id>             : null           - <port-id|bundle-id|bpgrp-id|lag-id|
                                         aps-id>
...
                        atm             - <port-id|aps-id>[:vpi/vci|vpi|
                                         vpi1.vpi2|
...
                        ima-grp         - <bundle-id>[:vpi/vci|vpi|vpi1.vpi2]

A:ALA-48>config>service>epipe# show service id 8 sap 881/1/2:4094
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 8
SAP                : 8/1/2:4094               Encap             : bcpDot1q
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Last Status Change : 02/06/2007 12:01:14
Last Mgmt Change   : 02/06/2007 12:01:17
Admin MTU          : 1522                     Oper MTU          : 1522
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
Ingress Filter-Id  : n/a                      Egress Filter-Id  : n/a
tod-suite          : None

Multi Svc Site     : None
Acct. Pol          : None                     Collect Stats     : Disabled
===============================================================================
A:ALA-48>config>service>epipe#
*A:bksim2801# config>service>apipe>sap$ show service id 1 sap 1/1/1:cp.2
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id                      : 1
SAP : 1/1/1:cp.2 Encap          : atm
Description                     : (Not Specified)
Admin State                     : Up               Oper State : Down
Flags                           : ServiceAdminDown
                                  PortOperDown
Multi Svc Site                  : None
Last Status Change              : 09/09/2010 07:55:28
Last Mgmt Change                : 09/09/2010 08:02:44
===============================================================================
*A:bksim2801#

A:ALA-48>config>service>epipe# show service id 8 sap 881/1/2:4094 detail
```

```
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 8
SAP                : 8/1/2:4094              Encap             : bcpDot1q
Admin State        : Up                      Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Last Status Change : 02/06/2007 12:01:14
Last Mgmt Change   : 02/06/2007 12:01:17
Admin MTU          : 1522                    Oper MTU          : 1522
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
Ingress Filter-Id  : n/a                     Egress Filter-Id  : n/a
tod-suite          : None

Multi Svc Site     : None
Acct. Pol          : None                    Collect Stats     : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                          Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
Off. Uncolor          : 0                        0
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                          Packets                  Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 0                        0
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
Egress Queue 1
For. InProf           : 0                        0
For. OutProf          : 0                        0
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
===============================================================================
A:ALA-48>config>service>epipe#
```

Table 30 **Show Service-ID SAP Output Fields**

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch,RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	The time of the most recent operating status change to this SAP.
Last Mgmt Change	The time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.

Table 30 **Show Service-ID SAP Output Fields (Continued)**

Label	Description (Continued)
LLF Admin State	Displays the Link Loss Forwarding administrative state.
LLF Oper State	Displays the Link Loss Forwarding operational state.
pw-port	**pw**-*id[:qtag1[.qtag2]]* **pw**-*id[:qtag1[.qtag2]]* pw-2:1.1

Sample Output

```
A:ALA-48# show service id 1 sap 1/1/1:2
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:5                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Up
Flags              : None
Last Status Change : 10/05/2006 17:06:03
Last Mgmt Change   : 10/05/2006 22:30:03
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingress qos-policy : 1190                     Egress qos-policy : 1190
Shared Q plcy      : n/a                      Multipoint shared : Disabled
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : suite_sixteen            qinq-pbit-marking : both
Egr Agg Rate Limit : max
ARP Reply Agent    : Unknown                  Host Conn Verify  : Disabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled
Multi Svc Site     : None
I. Sched Pol       : SchedPolCust1_Night
E. Sched Pol       : SchedPolCust1Egress_Night
Acct. Pol          : None                     Collect Stats     : Disabled
Anti Spoofing      : None                     Nbr Static Hosts  : 0
===============================================================================
A:ALA-48#

A:kerckhot_4# show service id 1 sap 1/1/1:6
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:6                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Admin State        : Up                       Oper State        : Down
Flags              : TodResourceUnavail
Last Status Change : 12/01/2006 09:59:42
Last Mgmt Change   : 12/01/2006 09:59:45
```

```
...
A:kerckhot_4#
```

sdp

Syntax **sdp** [[*sdp-id*[**:***vc-id*]] [**detail**]
sdp far-end {*ip-address* | *ipv6-address*} [**detail**]
sdp *sdp-id*[**:***vc-id*] **mrp**

Context show>service>id

Description This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Displays only information for the specified SDP ID.

Values 1 to 17407

vc-id — Displays only information for the specified virtual circuit ID.

Values 1 to 4294967295

ip-address — Displays only SDPs matching the specified far-end IPv4 address.

ipv6-address — Displays only SDPs matching the specified far-end IPv6 address.

detail — Displays detailed SDP information.

mrp — Displays detailed MRP information.

Output The following output is an example of SDP information, and Table 31 describes the output fields.

Sample Output

```
A:Dut-A# show service id 1 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
 Sdp Id 1:1  -(10.20.1.2)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id              : 1:1                       Type              : Spoke
VC Type             : Ether                     VC Tag            : n/a
Admin Path MTU      : 0                         Oper Path MTU     : 9186
Far End             : 10.20.1.2                 Delivery          : MPLS

Admin State         : Up                        Oper State        : Up
Acct. Pol           : None                      Collect Stats     : Disabled
Ingress Label       : 2048                      Egress Label      : 2048
Ing mac Fltr        : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr         : n/a                       Egr ip Fltr       : n/a
Ing ipv6 Fltr       : n/a                       Egr ipv6 Fltr     : n/a
Admin ControlWord   : Not Preferred             Oper ControlWord  : False
```

```
Last Status Change : 05/31/2007 00:45:43      Signaling         : None
Last Mgmt Change   : 05/31/2007 00:45:43
Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
Max Drop Count     : 3                        Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
MCAC Policy Name   :
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Associated LSP LIST :
Lsp Name           : A_B_1
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_2
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_3
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_4
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_5
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_6
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_7
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_8
```

```
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_9
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_10
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m34s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : enabled
Default LSP        : A_B_10                    Multicast LSP     : A_B_9
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
af                  A_B_3
be                  A_B_1
ef                  A_B_6
h1                  A_B_7
h2                  A_B_5
l1                  A_B_4
l2                  A_B_2
nc                  A_B_8
===============================================================================
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Forwarding
Port Number        : 2049                      Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                       Designated Port Id: 0
Fwd Transitions    : 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
A:Dut-A#
```

Table 31 **Show Service-ID SDP Output Fields**

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SDP belongs to.
VC Type	The VC type, Ether, VLAN, or VPLS.
VC Tag	The explicit dot1q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case).
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	Transmission frequency of the SDP echo request messages.

Table 31 **Show Service-ID SDP Output Fields (Continued)**

Label	Description (Continued)
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS.
Ingress Cookie1 Ingress Cookie2	Specifies the ingress cookies configured for an L2TPv3 spoke-SDP binding for an Epipe service. One or two L2TPv3 ingress cookies may be configured.
Egress Cookie	Specifies the egress cookies configured for an L2TPv3 spoke-SDPs for an Epipe service.
Session Mismatch	Specifies a mismatch detected between the configured (far-end binding) cookie to what is received by the local IP address of the L2TPv3 SDP. The flag is set when a mismatch is detected and must be manually cleared by an operator.

The following examples show both sides (PE nodes) when control word is enabled:

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
-------------------------------------------------------------------------------
 Sdp Id 1:2001  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1:2001                    Type             : Spoke
VC Type            : Ether                     VC Tag           : n/a
Admin Path MTU     : 1600                      Oper Path MTU    : 1600
Far End            : 1.1.1.1                   Delivery         : GRE

Admin State        : Up                        Oper State       : Up
Acct. Pol          : None                      Collect Stats    : Disabled
```

```
Ingress Label      : 115066                  Egress Label      : 119068
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Preferred               Oper ControlWord  : True
Last Status Change : 02/05/2007 16:39:22     Signaling         : TLDP
Last Mgmt Change   : 02/05/2007 16:39:22
Class Fwding State : Up
Endpoint           : N/A                     Precedence        : 4
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0

MAC Learning       : Enabled                 Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                BPDU Translation  : Disabled
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
*A:ALA-Dut-B>config>service>epipe#
```

The following is an example when one side (PE) has the control word enabled (the pipe will be down).

This is the side with control word disabled:

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
Sdp Id 1:2001  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1:2001                  Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 1600                    Oper Path MTU     : 1600
Far End            : 1.1.1.1                 Delivery          : GRE

Admin State        : Up                      Oper State        : Down
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 115066                  Egress Label      : 119068
```

```
Ing mac Fltr       : n/a                        Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                        Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                        Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred              Oper ControlWord  : False
Last Status Change : 02/05/2007 16:47:54        Signaling         : TLDP
Last Mgmt Change   : 02/05/2007 16:47:54
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                   Total MAC Addr    : 0
Learned MAC Addr   : 0                          Static MAC Addr   : 0
MAC Learning       : Enabled                    Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                   BPDU Translation  : Disabled
MAC Pinning        : Disabled
KeepAlive Information :
Admin State        : Disabled                   Oper State        : Disabled
Hello Time         : 10                         Hello Msg Len     : 0
Max Drop Count     : 3                          Hold Down Time    : 10
Statistics           :
I. Fwd. Pkts.      : 0                          I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                          E. Fwd. Octets    : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
*A:ALA-Dut-B>config>service>epipe#
```

This is the side with control word enabled:

```
*A:ALA-B# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
Sdp Id 1:12000  -(3.3.3.3)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1:12000                    Type              : Spoke
VC Type            : Ether                      VC Tag            : n/a
Admin Path MTU     : 1600                       Oper Path MTU     : 1600
Far End            : 3.3.3.3                    Delivery          : GRE
Admin State        : Up                         Oper State        : Down
Acct. Pol          : None                       Collect Stats     : Disabled
Ingress Label      : 119066                     Egress Label      : 0
Ing mac Fltr       : n/a                        Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                        Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                        Egr ipv6 Fltr     : n/a
Admin ControlWord  : Preferred                  Oper ControlWord  : True
Last Status Change : 02/04/2007 22:52:43        Signaling         : TLDP
Last Mgmt Change   : 02/04/2007 02:06:08
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
```

```
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
MAC Learning       : Enabled                 Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                BPDU Translation  : Disabled
MAC Pinning        : Disabled
KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
*A:ALA-B#
```

The following is an example when both sides have control word disabled:

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
Sdp Id 1:2001  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1:2001                  Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 1600                    Oper Path MTU     : 1600
Far End            : 1.1.1.1                 Delivery          : GRE
Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 115066                  Egress Label      : 119068
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 02/05/2007 16:49:05     Signaling         : TLDP
Last Mgmt Change   : 02/05/2007 16:47:54
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
MAC Learning       : Enabled                 Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                BPDU Translation  : Disabled
MAC Pinning        : Disabled
KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10
```

```
Statistics            :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
*A:ALA-Dut-B>config>service>epipe#


*A:SetupCLI>config>service>epipe>spoke-sdp# show service id 2 sdp 2000:1 detail
===============================================================================
Service Destination Point (Sdp Id : 2000:1) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 2000:1  -(101.101.101.101)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 2000:1                   Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 1500                     Oper Path MTU     : 1500
Far End            : 101.101.101.101          Delivery          : MPLS
Hash Label         : Enabled
Admin State        : Up                       Oper State        : Down
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 0                        Egress Label      : 0
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)     : 0
Last Status Change : 10/08/2009 06:55:54      Signaling         : TLDP
Last Mgmt Change   : 10/08/2009 07:04:27      Force Vlan-Vc     : Disabled
Endpoint           : N/A                      Precedence        : 4
Class Fwding State : Down
Flags              : SvcAdminDown SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Application Profile: None

KeepAlive Information :
Admin State        : Enabled                  Oper State        : No response
Hello Time         : 600                      Hello Msg Len     : 1500
Max Drop Count     : 3                        Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP LIST :
No LSPs Associated
```

```
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                   Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:SetupCLI>config>service>epipe>spoke-sdp#
```

Sample Output for L2TPv3 SDP binding

The following is sample output for L2TPv3 SDP binding, (not an MPLS or GRE SDP binding):

```
*A:cses-V36# show service id 999 sdp detail

===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 999:999  -(2001:db8::1)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 999:999                 Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type           : Ether                   VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 8890
Delivery          : L2TPv3
Far End           : 2001:db8::1
Local End         : 2001:db8:aaab::36
Tunnel Far End    : n/a                     LSP Types         : n/a
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled
Entropy Label     : Enabled

Admin State       : Up                      Oper State        : Up
Acct. Pol         : None                    Collect Stats     : Disabled
Ingress Label     : 0                       Egress Label      : 0
Ingr Mac Fltr-Id  : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id   : n/a                     Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id : n/a                     Egr IPv6 Fltr-Id  : n/a
Admin ControlWord : Not Preferred           Oper ControlWord  : False
Admin BW(Kbps)    : 0                       Oper BW(Kbps)     : 0
BFD Template      : None
BFD-Enabled       : no                      BFD-Encap         : ipv4
Last Status Change : 06/19/2014 17:31:16    Signaling         : None
Last Mgmt Change  : 06/19/2014 17:23:47     Force Vlan-Vc     : Disabled
Endpoint          : N/A                     Precedence        : 4
PW Status Sig     : Disabled
Force Qinq-Vc     : Disabled
Class Fwding State : Down
```

```
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None
Transit Policy     : None
Standby Sig Slave  : False
Block On Peer Fault: False
Use SDP B-MAC      : False

Ingress Qos Policy : (none)                  Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                       I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

-------------------------------------------------------------------------------
L2TPv3 Information
-------------------------------------------------------------------------------
Ingress Cookie     : AB:BA:BA:BB:A0:00:00:00
Ingress Cookie2    : BA:BA:BA:BA:BA:BA:BA:BA
Egress Cookie      : AB:BA:BA:BB:A0:00:00:00
Session Mismatch   : false
Sess Mismatch Clrd : 06/19/2014 17:23:21

-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false

-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
Squelch Levels     : None

-------------------------------------------------------------------------------
MPLS-TP LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated

-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
```

```
Default LSP        : Uknwn                    Multicast LSP     : None

===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

spoke-sdp-fec

Syntax **spoke-sdp-fec** [1 to 4294967295]

Context show>service>id

Description This command displays spoke-SDP FEC information.

Parameters **detail** — Displays detailed information.

Output The following output is an example of spoke-SDP FEC information.

Sample Output

```
===============================================================================
Service Spoke-SDP FEC Information
===============================================================================
Spoke-Sdp-Fec-Id     : 1                    Admin State      : enabled
FEC Type             : 129                  AII Type         : 2
Standby Sig Slave    : disabled             ICB              : disabled
Signaling            : auto                 Auto Config      : disabled
PW Template Id       : (none)               Precedence       : 4
Retry Timer          : 10 secs              Retry Count      : 10
Retry Timer Remaining: 0 secs               Retries Remaining: 0
SAII Type2           : 3:10.20.1.3:1
TAII Type2           : 6:10.20.1.6:1
Path                 : n/a
Endpoint             : n/a
Oper SDP-Bind        : 17407:4294967246
Last Error           : <none>
===============================================================================
Entries found: 1
===============================================================================
```

stp

Syntax **stp** [**detail**]
stp mst-instance *mst-inst-number*

Context show>service>id

Description This command displays information for the spanning tree protocol instance for the service.

Parameters **detail** — Displays detailed information.

mst-inst-number — Specifies an existing Multiple Spanning Tree Instance number.

Values 1 to 4094

spoke-sdp-fec

Syntax **spoke-sdp-fec** [[*1 to 4294967295*]]

Context show>service>id

Description This command displays the details of a spoke-sdp-fec spoke-sdp.

Output The following output is an example of spoke-SDP FEC information.

Sample Output

```
===============================================================================
Service Spoke-SDP FEC Information
===============================================================================
Spoke-Sdp-Fec-Id     : 1                    Admin State      : enabled
FEC Type             : 129                  AII Type         : 2
Standby Sig Slave    : disabled             ICB              : disabled
Signaling            : auto                 Auto Config      : disabled
PW Template Id       : (none)               Precedence       : 4
Retry Timer          : 10 secs              Retry Count      : 10
Retry Timer Remaining: 0 secs               Retries Remaining: 0
SAII Type2           : 3:10.20.1.3:1
TAII Type2           : 6:10.20.1.6:1
Path                 : n/a
Endpoint             : n/a
Oper SDP-Bind        : 17407:4294967246
Last Error           : <none>
===============================================================================
Entries found: 1
===============================================================================
```

sdp

Syntax **sdp** *sdp-id* **pw-port** [*pw-port-id*]
sdp *sdp-id* **pw-port**
sdp *sdp-id* **pw-port** *pw-port-id* [**statistics**]
sdp [**consistent** | **inconsistent** | **na**] **egressifs**
sdp *sdp-id* **keep-alive-history**
sdp far-end {*ip-address* | *ipv6-address*} **keep-alive-history**

sdp [*sdp-id*] **detail**
sdp far-end {*ip-address* | *ipv6-address*} **detail**

Context show>service

Description This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Specifies the SDP ID for which to display information.

Values 1 to 17407

pw-port-id — Specifies the pseudowire port identifier.

Values 1 to 10239

ip-address — Displays only SDPs with the specified far-end IPv4 address. 64 characters maximum.

ipv6-address — Displays only SDPs with the specified far-end IPv6 address. 64 characters maximum.

detail — Displays detailed SDP information.

Default SDP summary output

keep-alive-history — Displays the last fifty SDP keepalive events for the SDP.

Default SDP summary output

Output The following outputs are examples of SDP information.

Sample Output

```
*A:ALA-12>config>service# show service sdp 1 pw-port
===============================================================================
Service Destination Point (sdp Id 1 Pw-Port)
===============================================================================
Pw-port   VC-Id     Adm     Encap       Opr     VC Type     Egr       Monitor
                                                           Shaper    Oper
                                                           VPort     Group
-------------------------------------------------------------------------------
1         1         up      dot1q       up      ether
2         2         up      qinq        up      ether
3         3         up      dot1q       up      ether
4         4         up      qinq        up      ether
-------------------------------------------------------------------------------
Entries found : 4
===============================================================================

*A:ALA-12>config>service# show service sdp 1 pw-port 3

===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 3)
===============================================================================
SDP Binding port     : lag-1
VC-Id                : 3                    Admin Status       : up
Encap                : dot1q                Oper Status        : up
```

```
VC Type              : ether
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)
===============================================================================

*A:ALA-12>config>service# show service sdp 1 pw-port 3 statistics

===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 3)
===============================================================================
SDP Binding port     : lag-1
VC-Id                : 3                    Admin Status       : up
Encap                : dot1q                Oper Status        : up
VC Type              : ether
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)

Statistics            :
I. Fwd. Pkts.        : 0                    I. Dro. Pkts.      : 0
I. Fwd. Octs.        : 0                    I. Dro. Octs.      : 0
E. Fwd. Pkts.        : 0                    E. Fwd. Octets     : 0
===============================================================================


*A:Dut-B# show service sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 1  -10.20.1.3
-------------------------------------------------------------------------------
Description           : Default sdp description
SDP Id               : 1                    SDP Source         : manual
Admin Path MTU       : 1514                 Oper Path MTU      : 1514
Delivery             : MPLS
Far End              : 10.20.1.3
Tunnel Far End       : 10.20.1.3            LSP Types          : LDP
Admin State          : Up                   Oper State         : Up
Signaling            : TLDP                 Metric             : 0
Acct. Pol            : None                 Collect Stats      : Disabled
Last Status Change   : 06/13/2017 17:14:05  Adv. MTU Over.     : No
Last Mgmt Change     : 06/13/2017 17:17:19  VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                  PBB Etype          : 0x88e7
Oper Max BW(Kbps)    : 0                    Avail BW(Kbps)     : 0
Net-Domain           : default              Egr Interfaces     : Consistent
FPE LSP Id           : 0
Weighted ECMP        : Enabled
Flags                : None
Mixed LSP Mode Information :
Mixed LSP Mode       : Disabled             Active LSP Type    : LDP
KeepAlive Information :
Admin State          : Disabled             Oper State         : Disabled
Hello Time           : 10                   Hello Msg Len      : 0
Hello Timeout        : 5                    Unmatched Replies  : 0
Max Drop Count       : 3                    Hold Down Time     : 10
Tx Hello Msgs        : 0                    Rx Hello Msgs      : 0
Src B-MAC LSB        : <none>               Ctrl PW VC ID      : <none>
Ctrl PW Active       : n/a

-------------------------------------------------------------------------------
```

```
LDP Information :
-------------------------------------------------------------------------------
LDP LSP Id           : 65662

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding     : Disabled              EnforceDSTELspFc   : Disabled
Default LSP          : Uknwn                 Multicast LSP      : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
ISIS                 : disabled
OSPF                 : disabled
TE-LSP               : disabled
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:Dut-B#
```

pw-port

Syntax **pw-port** [*pw-port-id*] [**detail**]
pw-port sdp *sdp-id*
pw-port sdp none
pw-port *pw-port-id* **statistics**

Context show>pw-port

Description Displays pseudowire port information.

If no optional parameters are specified, the command displays a summary of all defined PW ports. The optional parameters restrict output to only ports matching the specified properties.

Parameters *pw-port-id* — Specifies the pseudowire port identifier.

Values 1 to 10239

detail — Displays detailed port information that includes all the **pw-port** output fields.

sdp-id — The SDP ID for which to display matching PW port information.

Values 1 to 17407

statistics — Displays statistics information.

Output The following output is an example of PW port information, and Table 32 described the output fields.

Sample Output

```
*A:ALA-48>config>service# show pw-port

===============================================================================
PW Port Information
===============================================================================
PW Port    Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
1          dot1q          1          1526726657          1
2          qinq           1          1526726658          2
3          dot1q          1          1526726659          3
4          qinq           1          1526726660          4
===============================================================================

*A:ALA-48>config>service# show pw-port 3
===============================================================================
PW Port Information
===============================================================================
PW Port    Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
3          dot1q          1          1526726659          3
===============================================================================
*A:ALA-48>config>service# show pw-port 3 detail

===============================================================================
PW Port Information
===============================================================================
PW Port             : 3
Encap               : dot1q
SDP                 : 1
IfIndex             : 1526726659
VC-Id               : 3
Description         : 1-Gig Ethernet dual fiber
===============================================================================
*A:ALA-48>config>pw-port$ show pw-port sdp none

===============================================================================
PW Port Information
===============================================================================
PW Port    Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
5          dot1q                     1526726661
===============================================================================


*A:ALA-48>config>pw-port$ show pw-port sdp 1

===============================================================================
PW Port Information
===============================================================================
PW Port    Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
```

```
1          dot1q          1          1526726657          1
2          qinq           1          1526726658          2
3          dot1q          1          1526726659          3
4          qinq           1          1526726660          4
===============================================================================
```

Table 32 **Show PW-Port Output Fields**

Label	Description
PW Port	The PW Port identifier.
Encap	The encapsulation type of the PW Port.
SDP	The SDP identifier.
IfIndex	The interface index used for the PW Port.
VC-Id	The Virtual Circuit identifier.
Description	The description string for the PW Port.

2.18.2.2 VLL Clear Commands

id

Syntax **id** {*service-id* | *service-name*} **neighbor**

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

For the 7450 ESS or 7750 SR, it clears the discovered IPv6 address of the neighboring CE associated with an iPipe SAP. When IPv6CP comes back up following the execution of this command on an IPv6CP SAP, the node will check to see if an IPv6 address has been learned for the remote CE attached to the Ipipe service. If one has been learned, then this is used to bring up IPv6CP.

Parameters *service-id* — The ID that uniquely identifies a service.

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters long.

service-name — Neighboring IPv6 address, for the 7450 ESS or 7750 SR only.

arp

Syntax **arp**

Context clear>service>id

Description This command clears all ARP entries. This command is only valid for Ipipe services.

neighbor

Syntax **neighbor**

Context clear>service>id

Description This command clears the discovered IPv6 address of the neighboring CE associated with an iPipe SAP.

host-tracking

Syntax **host-tracking** [**statistics**]
host-tracking sap *sap-id* [**host** *ip-address*] [**statistics**]

Context clear>service>id

Description This command clears host tracking data.

Parameters *sap-id* — Specifies a SAP for which to clear host tracking data.

ip-address — Specifies the IP address of a host for which to clear tracking data.

Values a.b.c.d

statistics — Clears statistics.

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[:*vc-id*] **ingress-vc-label**
mesh-sdp *sdp-id*[:*vc-id*] **vccv-bfd** {**session** | **statistics**}

Context clear>service>id

Description This command clears and resets the mesh SDP binding.

Parameters *sdp-id* — The spoke-SDP ID for which to clear statistics.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 to 4294967295

ingress-vc-label — Specifies to clear the ingress VC label.

vccv-bfd session — Specifies to clear the session mismatch flag on the mesh-SDP binding after the flag was set to true by a detected mismatch between the configured parameters and the received parameters.

vccv-bfd statistics — Specifies to clear a VCCV BFD session statistics for a specified mesh-SDP.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* [**ingress-vc-label**] [**lt2pv3**] [**vccv-bfd** {**session** | **statistics**}]

Context clear>service>id

Description This command clears and resets the spoke-SDP bindings for the service.

Parameters *sdp-id* — The spoke-SDP ID to be reset.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 to 4294967295

ingress-vc-label — Specifies to clear the ingress VC label.

l2tpv3 — Specifies to clear the session mismatch flag on the spoke-SDP binding after the flag was set to true by a detected mismatch between the configured parameters and the received parameters.

vccv-bfd session — Specifies to clear a VCCV BFD session for a specified spoke-SDP. Clearing the VCCV BFD session for a specified spoke-SDP will cause the session to go down and restart.

vccv-bfd statistics — Specifies to clear a VCCV BFD session statistics for a specified spoke-SDP.

statistics

Syntax **statistics**

Context clear>service

Description This command clears the statistics for a service.

sap

Syntax **sap** *sap-id* {**all** | **cem** | **counters** | **stp** | **l2pt** | **mrp**}
sap *sap-id* **encap-group** *group-name* [**member** *encap-id*]

Context clear>service>statistics

Description This command clears SAP statistics for a SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

all — Clears all SAP queue statistics and STP statistics.

counters — Clears all queue statistics associated with the SAP.

stp — Clears all STP statistics associated with the SAP.

l2pt — Clears all L2PT statistics associated with the SAP.

mrp — Clears all MRP statistics associated with the SAP.

group-name — Specifies the group name, up to 32 characters.

encap-id — Specifies the encapsulation ID.

Values 0 to 16777215

sdp

Syntax **sdp** *sdp-id* **keep-alive**
sdp *sdp-id* **pw-port** [*1 to 10239*]

Context clear>service>statistics

Description This command clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — The SDP ID for which to clear keepalive statistics.

Values 1 to 17407

keep-alive — Clears the keep-alive history associated with the SDP ID.

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] {**all** | **counters** | **stp** | **12pt** | **mrp**}

Context clear>service>statistics>id

Description This command clears statistics for the spoke-SDP bound to the service.

Parameters *sdp-id* — The spoke-SDP ID for which to clear statistics.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 to 4294967295

all — Clears all queue statistics and STP statistics associated with the SDP.

counters — Clears all queue statistics associated with the SDP.

stp — Clears all STP statistics associated with the SDP.

stp

Syntax **stp**

Context clear>service>statistics>id

Description Clears all spanning tree statistics for the service ID.

2.18.2.3 VLL Debug Commands

id

Syntax **id** *service-id*

Context debug>service

Description This command debugs commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters in length

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id

Description This command enables debugging for a particular SAP.

Parameters *sap-id* — Specifies the SAP ID.

event-type

Syntax [**no**] **event-type** {**arp** | **config-change** | **oper-status-change** | **neighbor-discovery**}

Context debug>service>id>sap

Description This command enables a particular debugging event type.

The **no** form of the command disables the event type debugging.

Parameters **arp** — Displays ARP events.

config-change — Debugs configuration change events.

oper-status-change — Debugs service operational status changes.

neighbor-discovery — Displays the status of IPv6 neighbor discovery for the sap or the spoke-sdp for the 7450 ESS or 7750 SR only.

Output The following output is an example of event-type information.

Sample Output

```
A:bksim180# debug service id 1000 sap 1/7/1 event-type arp
DEBUG OUTPUT show on CLI is as follows:
3 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000 SAP
1/7/1 "Service 1000 SAP 1/7/1:
RX: ARP_REQUEST (0x0001)
hwType     : 0x0001
prType     : 0x0800
hwLength   : 0x06
prLength   : 0x04
srcMac     : 8c:c7:01:07:00:03
destMac    : 00:00:00:00:00:00
srcIp      : 200.1.1.2
destIp     : 200.1.1.1
"

4 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000
SAP 1/7/1 "Service 1000 SAP 1/7/1:
TX: ARP_RESPONSE (0x0002)
hwType     : 0x0001
prType     : 0x0800
hwLength   : 0x06
prLength   : 0x04
srcMac     : 00:03:0a:0a:0a:0a
destMac    : 8c:c7:01:07:00:03
srcIp      : 200.1.1.1
destIp     : 200.1.1.2
"
```

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id

Description This command enables debugging for a particular SDP.

Parameters *sdp-id* — Specifies the SDP ID.

Values 1 to 17407

vc-id — Specifies the virtual circuit ID.

Values 1 to 4294967295

event-type

Syntax [**no**] **event-type** {**config-change** | **oper-status-change** | **neighbor-discovery** | **control-channel-status**}

Context debug>service>id>sdp

Description This command enables a particular debugging event type.

The **no** form of the command disables the event type debugging.

Parameters **config-change** — Debugs configuration change events.

oper-status-change — Debugs service operational status changes.

neighbor-discovery — Displays the status of IPv6 neighbor discovery for the sap or the spoke-sdp for the 7450 ESS or 7750 SR only.

control-channel-status —

2.18.2.4 VLL Tools Commands

epipe-map-access-to-egress-port

Syntax **epipe-map-access-to-egress-port** {**service** *target-svc-id* [**end-service** *end-svc-id*]} | **lag** *lag-id*

Context tools>dump

Description This command will display the egress port that will be used to transmit traffic associated with the displayed Epipe service(s). The information displayed shows the egress port for traffic traveling from SAP to egress SDP or SAP.

This command will support Epipe services with the following combinations:

- SAP to SDP (with no endpoint configuration)
- SAP to SAP (with or without an ICB)

- SAP to SDP using endpoints with 1 or 2 SDPs

The command can be executed by specifying either a service ID, service-ID range or an ingress LAG ID.

This command will not display the egress port for traffic traveling from the SDP to egress SAP. This command also does not work with services that use policers or shared queues and also does not support PBB services.

This command replaces the command tools dump epipe-map-to-network, which has been deprecated.

Parameters **service** *service-id* — Identifies the service ID for which the command will return the egress port. If used in conjunction with the end-service parameter, this value represent the beginning of the service ID range for which the command will be executed against.

Values 1 to 2148278316, *svc-name: 64 characters max*

end-service *service-id* — This parameter is used to identify the end of the service ID range for which the command will be executed against.

Values 1 to 2148278316, *svc-name: 64 characters max*

lag-id — This parameter caused the command to return the egress port for all service with SAPs associated with the specified LAG ID.

Values 1 to 800

3 Virtual Private LAN Service

3.1 VPLS Service Overview

VPLS as described in RFC 4905, *Encapsulation methods for transport of layer 2 frames over MPLS*, is a class of virtual private network service that allows the connection of multiple sites in a single bridged domain over a provider-managed IP/MPLS network. The customer sites in a VPLS instance appear to be on the same LAN, regardless of their location. VPLS uses an Ethernet interface on the customer-facing (access) side, which simplifies the LAN/WAN boundary and allows for rapid and flexible service provisioning.

VPLS offers a balance between point-to-point Frame Relay service and outsourced routed services (VPRN). VPLS enables each customer to maintain control of their own routing strategies. All customer routers in the VPLS service are part of the same subnet (LAN), which simplifies the IP addressing plan, especially when compared to a mesh constructed from many separate point-to-point connections. The VPLS service management is simplified since the service is not aware of nor participates in the IP addressing and routing.

A VPLS service provides connectivity between two or more SAPs on one (which is considered a local service) or more (which is considered a distributed service) service routers. The connection appears to be a bridged domain to the customer sites so protocols, including routing protocols, can traverse the VPLS service.

Other VPLS advantages include:

- VPLS is a transparent, protocol-independent service.
- There is no Layer 2 protocol conversion between LAN and WAN technologies.
- There is no need to design, manage, configure, and maintain separate WAN access equipment, which eliminates the need to train personnel on WAN technologies such as Frame Relay.

3.1.1 VPLS Packet Walkthrough

This section provides an example of VPLS processing of a customer packet sent across the network from site A, which is connected to PE Router A, to site B, which is connected to PE Router C (see Figure 56).

Figure 56 **VPLS Service Architecture**

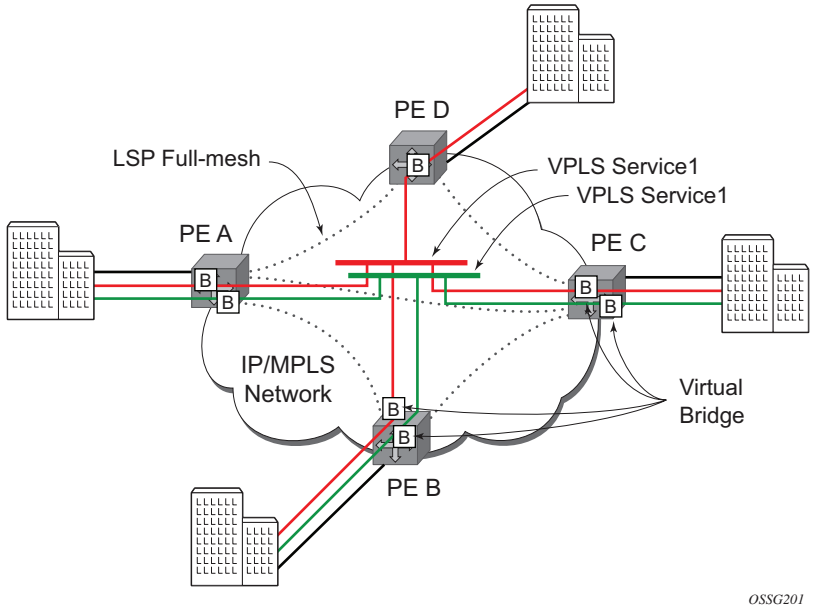

1. PE Router A (Figure 57)
   a. Service packets arriving at PE Router A are associated with a VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet.

Figure 57 **Access Port Ingress Packet Format and Lookup**

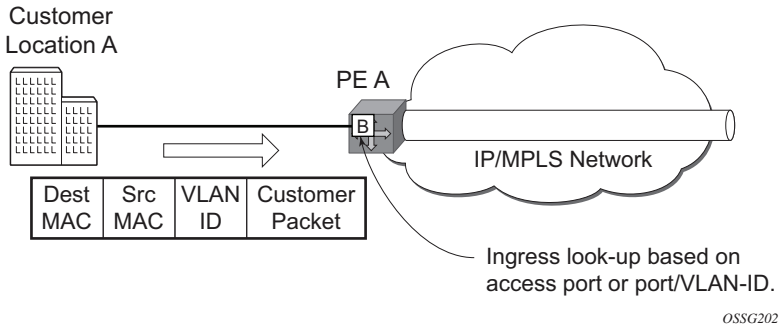

   b. PE Router A learns the source MAC address in the packet and creates an entry in the FDB table that associates the MAC address to the service access point (SAP) on which it was received.

c. The destination MAC address in the packet is looked up in the FDB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address is not yet learned (unknown MAC address).

   For a Known MAC Address (Figure 58)

d. If the destination MAC address has already been learned by PE Router A, an existing entry in the FDB table identifies the far-end PE-router and the service VC-label (inner label) to be used before sending the packet to far-end PE Router C.

e. PE Router A chooses a transport LSP to send the customer packets to PE Router C. The customer packet is sent on this LSP after the IEEE 802.1Q tag is stripped and the service VC-label (inner label) and the transport label (outer label) are added to the packet.

   For an Unknown MAC Address (Figure 58)

If the destination MAC address has not been learned, PE Router A will flood the packet to both PE Router B and PE Router C that are participating in the service by using the VC-labels that each PE Router previously added for the VPLS instance. The packet is not sent to PE Router D since this VPLS service does not exist on that PE-router.

Figure 58 **Network Port Egress Packet Format and Flooding**

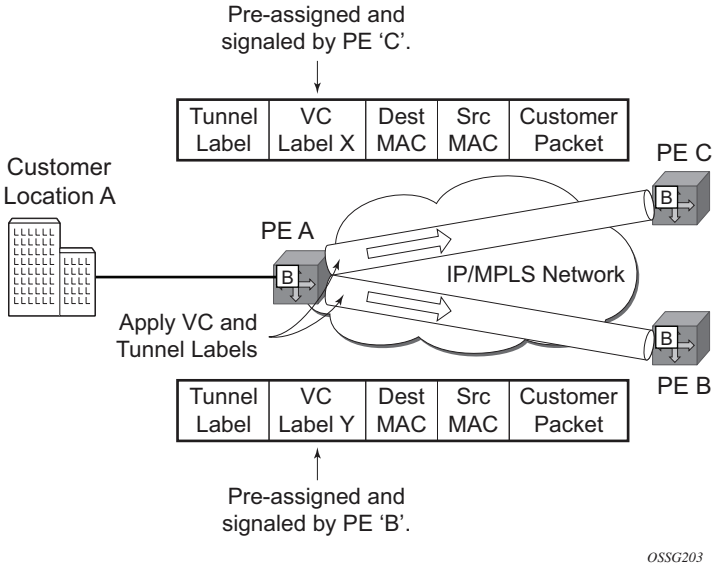

2. Core Router Switching

All the core routers (“P” routers in IETF nomenclature) between PE Router A and PE Router B and PE Router C are Label Switch Routers (LSRs) that switch the packet based on the transport (outer) label of the packet until the packet arrives at the far-end PE Router. All core routers are unaware that this traffic is associated with a VPLS service.

3. PE Router C
   a. PE Router C strips the transport label of the received packet to reveal the inner VC-label. The VC-label identifies the VPLS service instance to which the packet belongs.
   b. PE Router C learns the source MAC address in the packet and creates an entry in the FDB table that associates the MAC address to PE Router A and the VC-label that PE Router A added for the VPLS service on which the packet was received
   c. The destination MAC address in the packet is looked up in the FDB table for the VPLS instance. Again, there are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address has not been learned on the access side of PE Router-C (unknown MAC address).

      For a Known MAC address (Figure 59)

If the destination MAC address has been learned by PE Router C, an existing entry in the FDB table identifies the local access port and the IEEE 802.1Q tag to be added before sending the packet to customer Location-C. The egress Q tag may be different than the ingress Q tag.

Figure 59 **Access Port Egress Packet Format and Lookup**

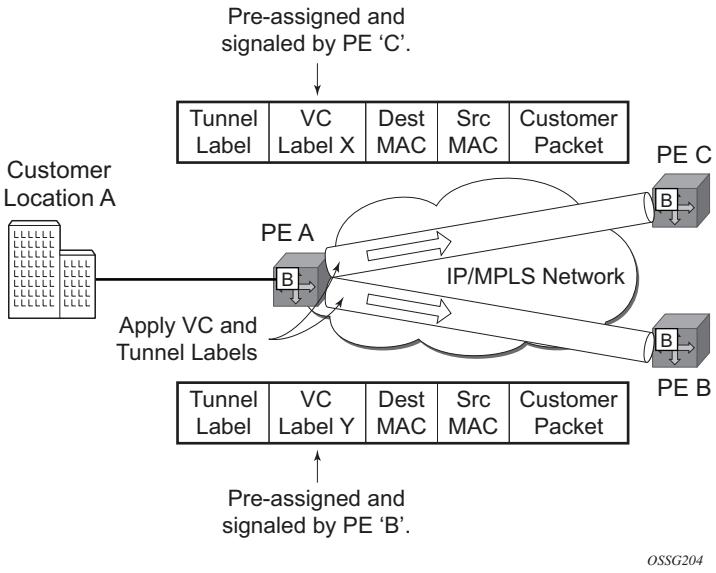

3.2 VPLS Features

This section provides information about VPLS features.

3.2.1 VPLS Enhancements

Nokia's VPLS implementation includes several enhancements beyond basic VPN connectivity. The following VPLS features can be configured individually for each VPLS service instance:

- Extensive MAC and IP filter support (up to Layer 4). Filters can be applied on a per-SAP basis.
- Forwarding Database (FDB) management features on a per service-level basis including:
  - Configurable FDB size limit. On the 7450 ESS, it can be configured on a per-VPLS, per-SAP, and per spoke-SDP basis.
  - FDB size alarms. On the 7450 ESS, it can be configured on a per-VPLS basis.
  - MAC learning disable. On the 7450 ESS, it can be configured on a per-VPLS, per-SAP, and per spoke-SDP basis.
  - Discard unknown. On the 7450 ESS, it can be configured on a per VPLS basis.
  - Separate aging timers for locally and remotely learned MAC addresses.
- Ingress rate limiting for broadcast, multicast, and destination unknown flooding on a per-SAP basis.
- Implementation of STP parameters on a per-VPLS, per-SAP, and per spoke-SDP basis.
- A split horizon group on a per-SAP and per-spoke-SDP basis.
- DHCP snooping and anti-spoofing on a per-SAP and per-SDP basis for the 7450 ESS or 7750 SR.
- IGMP snooping on a per-SAP and per-SDP basis.
- Optional SAP and/or spoke-SDP redundancy to protect against node failure.

3.2.2 VPLS over MPLS

The VPLS architecture proposed in RFC 4762, *Virtual Private LAN Services Using LDP Signaling* specifies the use of provider equipment (PE) that is capable of learning, bridging, and replication on a per-VPLS basis. The PE routers that participate in the service are connected using MPLS Label Switched Path (LSP) tunnels in a full-mesh composed of mesh SDPs or based on an LSP hierarchy (Hierarchical VPLS (H-VPLS)) composed of mesh SDPs and spoke-SDPs.

Multiple VPLS services can be offered over the same set of LSP tunnels. Signaling specified in RFC 4905, *Encapsulation methods for transport of layer 2 frames over MPLS* is used to negotiate a set of ingress and egress VC labels on a per-service basis. The VC labels are used by the PE routers for demultiplexing traffic arriving from different VPLS services over the same set of LSP tunnels.

VPLS is provided over MPLS by:

- connecting bridging-capable provider edge routers with a full mesh of MPLS LSP tunnels
- negotiating per-service VC labels using *draft-Martini* encapsulation
- replicating unknown and broadcast traffic in a service domain
- enabling MAC learning over tunnel and access ports (see VPLS MAC Learning and Packet Forwarding)
- using a separate FDB per VPLS service

3.2.3 VPLS Service Pseudowire VLAN Tag Processing

VPLS services can be connected using pseudowires that can be provisioned statically or dynamically and are represented in the system as either a mesh or a spoke-SDP. The mesh and spoke-SDP can be configured to process zero, one, or two VLAN tags as traffic is transmitted and received. In the transmit direction, VLAN tags are added to the frame being sent, and in the received direction, VLAN tags are removed from the frame being received. This is analogous to the SAP operations on a null, dot1q, and QinQ SAP.

The system expects a symmetrical configuration with its peer; specifically, it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. When removing VLAN tags from a mesh or spoke-SDP, the system attempts to remove the configured number of VLAN tags (see below for the configuration details); if fewer tags are found, the system removes the VLAN tags

found and forwards the resulting packet. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. With an asymmetrical behavior, protocol extractions will not necessarily function as they would with a symmetrical configuration, resulting in an unexpected operation.

The VLAN tag processing is configured as follows on a mesh or spoke-SDP in a VPLS service:

- Zero VLAN tags processed—This requires the configuration of **vc-type ether** under the mesh-SDP or spoke-SDP, or in the related **pw-template**.
- One VLAN tag processed—This requires one of the following configurations:
  - **vc-type vlan** under the mesh-SDP or spoke-SDP, or in the related **pw-template.**
  - **vc-type ether** and **force-vlan-vc-forwarding** under the mesh-SDP or spoke-SDP, or in the related **pw-template.**
- Two VLAN tags processed—This requires the configuration of **force-qinq-vc-forwarding** under the mesh-SDP or spoke-SDP, or in the related **pw-template**.

The **pw-template** configuration provides support for BGP VPLS services and LDP VPLS services using BGP Auto-Discovery.

The following restrictions apply to VLAN tag processing:

- The configuration of **vc-type vlan** and **force-vlan-vc-forwarding** is mutually exclusive.
- BGP VPLS services operate in a mode equivalent to **vc-type ether**; consequently, the configuration of **vc-type vlan** in a **pw-template** for a BGP VPLS service is ignored.
- **force-qinq-vc-forwarding** can be configured with the mesh-SDP or spoke-SDP signaled as either **vc-type ether** or **vc-type vlan**.
- The following are not supported with **force-qinq-vc-forwarding** configured under the mesh-SDP or spoke-SDP, or in the related **pw-template**:
  - Routed, E-Tree, or PBB VPLS services.
  - L2PT termination on QinQ mesh-SDP or spoke-SDPs.
  - IGMP/MLD/PIM snooping within the VPLS service.
  - ETH-CFM MIPs and MEPs are not supported on dynamically signaled BGP QinQ PWs.

Table 33 and Table 34 describe the VLAN tag processing with respect to the zero, one, and two VLAN tag configuration described for the VLAN identifiers, Ethertype, ingress QoS classification (dot1p/DE), and QoS propagation to the egress (which can be used for egress classification and/or to set the QoS information in the innermost egress VLAN tag).

Table 33 **VPLS Mesh and Spoke SDP VLAN Tag Processing: Ingress**

Ingress (received on mesh or spoke-SDP)	**Zero VLAN tags**	**One VLAN tag**	**Two VLAN tags**
VLAN identifiers	N/A	Ignored	Both inner and outer ignored
Ethertype (to determine the presence of a VLAN tag)	N/A	0x8100 or value configured under **sdp vlan-vc-etype**	Both inner and outer VLAN tags: 0x8100, or outer VLAN tag value configured under **sdp vlan-vc-etype** (inner VLAN tag value must be 0x8100)
Ingress QoS (dot1p/ DE) classification	N/A	Ignored	Both inner and outer ignored
QoeE (dot1p/DE) propagation to egress	Dot1p/DE=0	Dot1p/DE taken from received VLAN tag	Dot1p/DE taken from inner received VLAN tag

Table 34 **VPLS Mesh and Spoke SDP VLAN Tag Processing: Egress**

Egress (sent on mesh or spoke-SDP)	Zero VLAN tags	One VLAN tag	Two VLAN tags
VLAN identifiers (set in VLAN tags)	N/A	For one VLAN tag, one of the following applies: • the **vlan-vc-tag** value configured in **pw-template** or under the mesh/spoke-SDvalue fromP • value from the inner tag received on a QinQ SAP or QinQ mesh/spoke-SDP • value from the VLAN tag received on a dot1q SAP or mesh/spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or mesh/spoke-SDP	For both inner and outer VLAN tags, one of the following applies: • the **vlan-vc-tag** value configured in **pw-template** or under the mesh/spoke-SDP • value from the inner tag received on a QinQ SAP or QinQ mesh/spoke-SDP • value from the VLAN tag received on a dot1q SAP or mesh/spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or mesh/spoke-SDP
Ethertype (set in VLAN tags)	N/A	0x8100 or value configured under **sdp vlan-vc-etype**	Both inner and outer VLAN tags: 0x8100, or outer VLAN tag value configured under **sdp vlan-vc-etype** (inner VLAN tag value will be 0x8100)

Table 34 **VPLS Mesh and Spoke SDP VLAN Tag Processing: Egress (Continued)**

Egress (sent on mesh or spoke-SDP)	Zero VLAN tags	One VLAN tag	Two VLAN tags
Egress QoS (dot1p/DE) (set in VLAN tags)	N/A	Taken from the innermost ingress service delimiting tag, one of the following applies: • the inner tag received on a QinQ SAP or QinQ mesh/spoke-SDP • value from the VLAN tag received on a dot1q SAP or mesh/spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or mesh/spoke-SDP **Note**: neither the inner nor outer dot1p/DE values can be explicitly set.	For both inner and outer dot1p/DE, the value is taken from the innermost ingress service delimiting tag. One of the following applies: • the inner tag received on a QinQ SAP or QinQ mesh/spoke-SDP • value from the VLAN tag received on a dot1q SAP or mesh/spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**) • value from the outer tag received on a qtag.* SAP • 0 if there is no service delimiting VLAN tag at the ingress SAP or mesh/spoke-SDP **Note**: neither the inner nor outer dot1p/DE values can be explicitly set.

Any non-service delimiting VLAN tags are forwarded transparently through the VPLS service. SAP egress classification is possible on the outermost customer VLAN tag received on a mesh or spoke-SDP using the **ethernet-ctag** parameter in the associated SAP egress QoS policy.

3.2.4 VPLS MAC Learning and Packet Forwarding

The 7950 XRS, 7750 SR, and 7450 ESS perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the router to reduce the amount of unknown destination MAC address flooding.

The 7450 ESS, 7750 SR, and 7950 XRS routers learn the source MAC addresses of the traffic arriving on their access and network ports.

Each router maintains a Forwarding Database (FDB) for each VPLS service instance and learned MAC addresses are populated in the FDB table of the service. All traffic is switched based on MAC addresses and forwarded between all objects in the VPLS service. Unknown destination packets (for example, the destination MAC address has not been learned) are forwarded on all objects to all participating nodes for that service until the target station responds and the MAC address is learned by the routers associated with that service.

3.2.4.1 MAC Learning Protection

In a Layer 2 environment, subscribers or customers connected to SAPs A, B, C can create a denial of service attack by sending packets sourcing the gateway MAC address. This will move the learned gateway MAC from the uplink SDP/SAP to the subscriber's or customer's SAP causing all communication to the gateway to be disrupted. If local content is attached to the same VPLS (D), a similar attack can be launched against it. Communication between subscribers or customers is also disallowed but split horizon will not be sufficient in the topology shown in Figure 60.

Figure 60 **MAC Learning Protection**

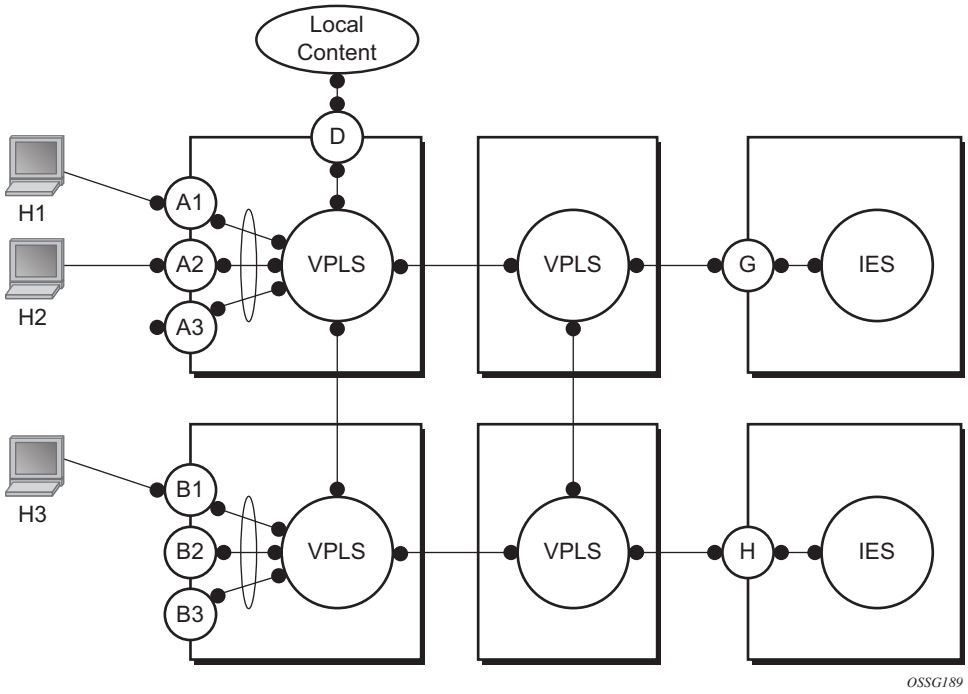

The 7450 ESS, 7750 SR, and 7950 XRS routers enable MAC learning protection capability for SAPs and SDPs. With this mechanism, forwarding and learning rules apply to the non-protected SAPs. Assume hosts H1, H2, and H3 (Figure 60) are non-protected while IES interfaces G and H are protected. When a frame arrives at a protected SAP/SDP, the MAC is learned as usual. When a frame arrives from a non-protected SAP or SDP, the frame must be dropped if the source MAC address is protected and the MAC address is not relearned. The system allows only packets with a protected MAC destination address.

The system can be configured statically. The addresses of all protected MACs are configured. Only the IP address can be included and use a dynamic mechanism to resolve the MAC address (**cpe-ping**). All protected MACs in all VPLS instances in the network must be configured.

To eliminate the ability of a subscriber or customer to cause a DoS attack, the node restricts the learning of protected MAC addresses based on a statically defined list. Also, the destination MAC address is checked against the protected MAC list to verify that a packet entering a restricted SAP has a protected MAC as a destination.

3.2.4.2 DEI in IEEE 802.1ad

The IEEE 802.1ad-2005 standard allows drop eligibility to be conveyed separately from priority in Service VLAN TAGs (S-TAGs) so that all of the previously introduced traffic types can be marked as drop eligible. The Service VLAN TAG has a new format where the priority and discard eligibility parameters are conveyed in the 3-bit Priority Code Point (PCP) field and, respectively, in the DE Bit (Figure 61).

Figure 61 **DE Bit in the 802.1ad S-TAG**

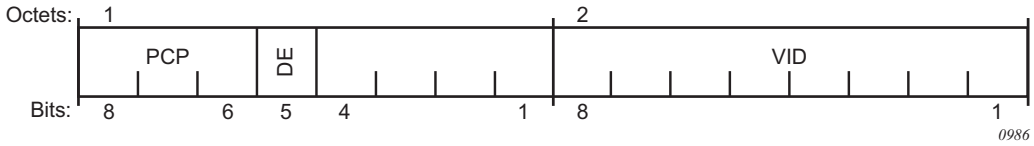

The DE bit allows the S-TAG to convey eight forwarding classes/distinct emission priorities, each with a drop eligible indication.

When the DE bit is set to 0 (DE=FALSE), the related packet is not discard eligible. This is the case for the packets that are within the CIR limits and must be given priority in case of congestion. If the DEI is not used or backwards compliance is required, the DE bit should be set to zero on transmission and ignored on reception.

When the DE bit is set to 1 (DE=TRUE), the related packet is discard eligible. This is the case for the packets that are sent above the CIR limit (but below the PIR). In case of congestion, these packets will be the first ones to be dropped.

3.2.5 VPLS Using G.8031 Protected Ethernet Tunnels

The use of MPLS tunnels provides a way to scale the core while offering fast failover times using MPLS FRR. In environments where Ethernet services are deployed using native Ethernet backbones, Ethernet tunnels are provided to achieve the same fast failover times as in the MPLS FRR case. There are still service provider environments where Ethernet services are deployed using native Ethernet backbones.

The Nokia VPLS implementation offers the capability to use core Ethernet tunnels compliant with ITU-T G.8031 specification to achieve 50 ms resiliency for backbone failures. This is required to comply with the stringent SLAs provided by service providers in the current competitive environment. The implementation also allows a LAG-emulating Ethernet tunnel providing a complimentary native Ethernet E-LAN capability. The LAG-emulating Ethernet tunnels and G.8031 protected Ethernet tunnels operate independently (refer to *"LAG emulation using Ethernet Tunnels"* in the *Services Overview* Guide).

When using Ethernet tunnels, the Ethernet tunnel logical interface is created first. The Ethernet tunnel has member ports that are the physical ports supporting the links. The Ethernet tunnel controls SAPs that carry G.8031 and 802.1ag control traffic and user data traffic. Ethernet Service SAPs are configured on the Ethernet tunnel. Optionally, when tunnels follow the same paths end-to-end services are configured with same-fate Ethernet tunnel SAPs which carry only user data traffic and shares the fate of the Ethernet tunnel port (if properly configured).

When configuring VPLS and B-VPLS using Ethernet tunnels, the services are very similar.

Refer to the *IEEE 802.1ah PBB Guide* for examples.

3.2.6 Pseudowire Control Word

The **control-word** command enables the use of the control word individually on each mesh SDP or spoke-SDP. By default, the control word is disabled. When the control word is enabled, all VPLS packets, including the BPDU frames, are encapsulated with the control word. The Targeted LDP (T-LDP) control plane behavior will be the same as the control word for VLL services. The configuration for the two directions of the Ethernet pseudowire should match.

3.2.7 Table Management

The following sections describe VPLS features related to management of the Forwarding Database (FDB).

3.2.7.1 Selective MAC Address Learning

Source MAC addresses are learned in a VPLS service by default with an entry allocated in the FDB for each address on all line cards. Therefore, all MAC addresses are considered to be global. This operation can be modified so that the line card allocation of some MAC addresses is selective, based on where the service has a configured object.

An example of the advantage of selective MAC address learning is for services to benefit from the higher MAC address scale of some line cards (particularly for network interfaces used by mesh or spoke-SDPs, EVPN-VXLAN tunnels, and EVPN-MPLS destinations) while using lower MAC address scale cards for the SAPs.

Selective MAC addresses are those learned locally and dynamically in the data path (displayed in the **show** output with type “L”) or by EVPN (displayed in the **show** output with type “Evpn”, excluding those with the sticky bit set, which are displayed with type “EvpnS”). An exception is when a MAC address configured as a conditional static MAC address is learned dynamically on an object other than its monitored object; this can be displayed with type “L” or “Evpn” but is learned as global because of the conditional static MAC configuration.

Selective MAC addresses have FDB entries allocated on line cards where the service has a configured object. When a MAC address is learned, it is allocated an FDB entry on all line cards on which the service has a SAP configured (for LAG or Ethernet tunnel SAPs, the MAC address is allocated an FDB entry on all line cards on which that LAG or Ethernet tunnel has configured ports) and on all line cards that have a network interface port if the service is configured with VXLAN, EVPN-MPLS, or a mesh or spoke-SDP.

When using selective learning in an I-VPLS service, the learned CMACs are allocated FDB entries on all the line cards where the I-VPLS service has a configured object and on the line cards on which the associated B-VPLS has a configured object. When using selective learning in a VPLS service with **allow-ip-intf-bind** configured (for it to become a routed VPLS), FDB entries are allocated on all line cards on which there is an IES or VPRN interface.

If a new configured object is added to a service and there are sufficient MAC FDB resources available on the new line cards, the selective MAC addresses present in the service are allocated on the new line cards. Otherwise, if any of the selective MAC addresses currently learned in the service cannot be allocated an FDB entry on the new line cards, those MAC addresses will be deleted from all line cards. Such a deletion increments the FailedMacCmplxMapUpdts statistic displayed in the **tools dump service vpls-fdb-stats** output.

When the set of configured objects changes for a service using selective learning, the system must reallocate its FDB entries accordingly, which can cause FDB entry "allocate" or "free" operations to pend temporarily. The pending operations can be displayed using the **tools dump service** *id* **fdb** command.

When a global MAC address is to be learned, there must be a free FDB entry in the service and system FDBs and on all line cards in the system for it to be accepted. When a selective MAC address is to be learned, there must be a free FDB entry in the service and system FDBs and on all line cards where the service has a configured object for it to be accepted.

To demonstrate the selective MAC address learning logic, consider the following:

- a system has three line cards: 1, 2, and 3
- two VPLS services are configured on the system:
  - VPLS 1 having learned MAC addresses M1, M2, and M3 and has configured SAPs 1/1/1 and 2/1/1
  - VPLS 2 having learned MAC addresses M4, M5, and M6 and has configured SAPs 2/1/2 and 3/1/1

This is shown in Table 35.

Table 35 **MAC Address Learning Logic Example**

	Learned MAC addresses	Configured SAPs
VPLS1	M1, M2, M3	sap 1/1/1 sap 2/1/1
VPLS2	M4, M5, M6	sap 2/1/2 sap 3/1/1

Figure 62 shows the FDB entry allocation when the MAC addresses are global and when they are selective. Notice that in the selective case, all MAC addresses are allocated FDB entries on line card 2, but line card 1 and 3 only have FDB entries allocated for services VPLS 1 and VPLS 2, respectively.

Figure 62 **MAC FDB Entry Allocation: Global versus Selective**

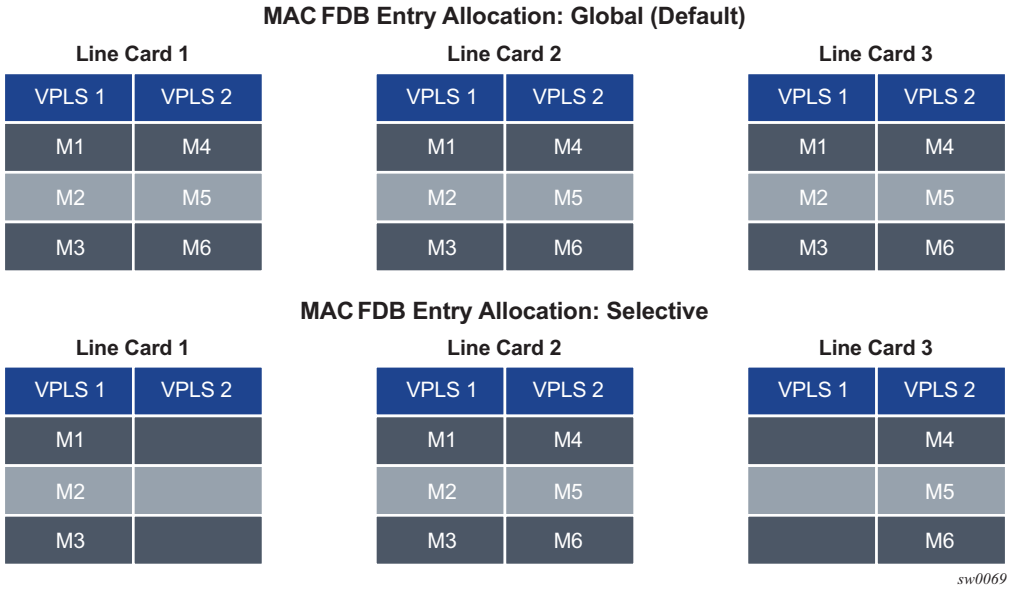

Selective MAC address learning can be enabled as follows within any VPLS service, except for B-VPLS and routed VPLS services:

```
configure
    service
        vpls <service-id> create
            [no] selective-learned-fdb
```

Enabling selective MAC address learning has no effect on single line card systems.

When selective learning is enabled or disabled in a VPLS service, the system may need to reallocate FDB entries; this can cause temporary pending FDB entry allocate or free operations. The pending operations can be displayed using the **tools dump service** *id* **fdb** command.

3.2.7.1.1 Example Operational Information

The **show** and **tools dump** command output can display the global and selective MAC addresses along with the MAC address limits and the number of allocated and free MAC addresses FDB entries. The **show** output displays the system and card FDB usage, while the **tools** output displays the FDB per service with respect to MAC addresses and cards.

The configuration for the following output is similar to the simple example above:

- the system has three line cards: 1, 2, and 5
- the system has two VPLS services:
  - VPLS 1 is an EVPN-MPLS service with a SAP on 5/1/1:1 and uses a network interface on 5/1/5.
  - VPLS 2 has two SAPs on 2/1/1:2 and 2/1/2:2.

The first output shows the default where all MAC addresses are global. The second enables selective learning in the two VPLS services.

Global MAC address learning only (default)

By default, VPLS 1 and 2 are not configured for selective learning, so all MAC addresses are global:

```
*A:PE1# show service id [1,2] fdb | match expression ", Service|Sel Learned FDB"
Forwarding Database, Service 1
Sel Learned FDB   : Disabled
Forwarding Database, Service 2
Sel Learned FDB   : Disabled
*A:PE1#
```

Traffic is sent into the services, resulting in the following MAC addresses being learned:

```
*A:PE1# show service fdb-mac
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
```

```
-------------------------------------------------------------------------------
1         00:00:00:00:01:01 sap:5/1/1:1             L/0      01/31/17 08:44:37
1         00:00:00:00:01:02 sap:5/1/1:1             L/0      01/31/17 08:44:37
1         00:00:00:00:01:03 eMpls:                  EvpnS    01/31/17 08:41:38
                                                    P
                            10.251.72.58:262142
1         00:00:00:00:01:04 eMpls:                  EvpnS    01/31/17 08:41:38
                                                    P
                            10.251.72.58:262142
2         00:00:00:00:02:01 sap:2/1/2:2             L/0      01/31/17 08:44:37
2         00:00:00:00:02:02 sap:2/1/2:2             L/0      01/31/17 08:44:37
2         00:00:00:02:02:03 sap:2/1/1:2             L/0      01/31/17 08:44:37
2         00:00:00:02:02:04 sap:2/1/1:2             L/0      01/31/17 08:44:37
-------------------------------------------------------------------------------
No. of Entries: 8
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
*A:PE1#
```

A total of eight MAC addresses are learned. There are two MAC addresses learned locally on SAP 5/1/1:1 in service VPLS 1 (type “L”), another two MAC addresses learned using EVPN with the sticky bit set, also in service VPLS 1 (type “EvpnS”). A further two sets of two MAC addresses are learned on SAP 2/1/1:2 and 2/1/2:2 in service VPLS 2 (type “L”).

The system and line card FDB usage is shown as follows:

```
*A:PE1# show service system fdb-usage
===============================================================================
FDB Usage
===============================================================================
System
-------------------------------------------------------------------------------
Limit:     511999
Allocated: 8
Free:      511991
Global:    8
-------------------------------------------------------------------------------
Line Cards
-------------------------------------------------------------------------------
Card        Selective         Allocated         Limit             Free
-------------------------------------------------------------------------------
1           0                 8                 511999            511991
2           0                 8                 511999            511991
5           0                 8                 511999            511991
-------------------------------------------------------------------------------
===============================================================================
*A:PE1#
```

The system MAC address limit is 511999, of which eight are allocated, and the rest are free. All eight MAC addresses are global and are allocated on cards 1, 2, and 5. There are no selective MAC addresses. This output can be reduced to specific line cards by specifying the card’s slot ID as a parameter to the command.

To see the MAC address information per service, **tools dump** commands can be used, as follows for VPLS 1. The following output displays the card status:

```
*A:PE1# tools dump service id 1 fdb card-status
===============================================================================
VPLS FDB Card Status at 01/31/2017 08:44:38
===============================================================================
Card                Allocated           PendAlloc           PendFree
-------------------------------------------------------------------------------
1                   4                   0                   0
2                   4                   0                   0
5                   4                   0                   0
===============================================================================
*A:PE1#
```

All of the line cards have four FDB entries allocated in VPLS 1. The “PendAlloc” and “PendFree” columns show the number of pending MAC address allocate and free operations, which are all zero.

The following output displays the MAC address status for VPLS 1:

```
*A:PE1# tools dump service id 1 fdb mac-status
===============================================================================
VPLS FDB MAC status at 01/31/2017 08:44:38
===============================================================================
MAC Address         Type                Status : Card list
-------------------------------------------------------------------------------
00:00:00:00:01:01   Global              Allocated : All
00:00:00:00:01:02   Global              Allocated : All
00:00:00:00:01:03   Global              Allocated : All
00:00:00:00:01:04   Global              Allocated : All
===============================================================================
*A:PE1#
```

The type and card list for each MAC address in VPLS 1 is displayed. VPLS 1 has learned four MAC addresses; the two local MAC addresses on SAP 5/1/1:1 and the two EvpnS MAC addresses. Each MAC address has an FDB entry allocated on all line cards. This output can be further reduced by optionally specifying a specified MAC address, a specific card, and the operational pending state.

Selective and global MAC address learning

Selective MAC address learning is now enabled in VPLS 1 and VPLS 2, as follows:

```
*A:PE1# show service id [1,2] fdb | match expression ", Service|Sel Learned FDB"
Forwarding Database, Service 1
Sel Learned FDB   : Enabled
Forwarding Database, Service 2
Sel Learned FDB   : Enabled
*A:PE1#
```

The MAC addresses learned are the same, with the same traffic being sent; however, there are now selective MAC addresses that are allocated FDB entries on different line cards.

The system and line card FDB usage is as follows:

```
*A:PE1# show service system fdb-usage
===============================================================================
FDB Usage
===============================================================================
System
-------------------------------------------------------------------------------
Limit:     511999
Allocated: 8
Free:      511991
Global:    2
-------------------------------------------------------------------------------
Line Cards
-------------------------------------------------------------------------------
Card        Selective         Allocated         Limit             Free
-------------------------------------------------------------------------------
1           0                 2                 511999            511997
2           4                 6                 511999            511993
5           2                 4                 511999            511995
-------------------------------------------------------------------------------
===============================================================================
*A:PE1#
```

The system MAC address limit and allocated numbers have not changed but now there are only two global MAC addresses; these are the two EvpnS MAC addresses.

There are two FDB entries allocated on card 1, which are the global MAC addresses; there are no services or network interfaces configured on card 1, so the FDB entries allocated are for the global MAC addresses.

Card 2 has six FDB entries allocated in total: two for the global MAC addresses plus four for the selective MAC addresses in VPLS 2 (these are the two sets of two local MAC addresses in VPLS 2 on SAP 2/1/1:2 and 2/1/2:2).

Card 5 has four FDB entries allocated in total; two for the global MAC addresses plus two for the selective MAC addresses in VPLS 1 (these are the two local MAC addresses in VPLS 1 on SAP 5/1/1:1).

This output can be reduced to specific line cards by specifying the card's slot ID as a parameter to the command.

To see the MAC address information per service, **tools dump** commands can be used for VPLS 1.

The following output displays the card status:

```
*A:PE1# tools dump service id 1 fdb card-status
```

```
===============================================================================
VPLS FDB Card Status at 01/31/2017 08:44:39
===============================================================================
Card                Allocated           PendAlloc           PendFree
-------------------------------------------------------------------------------
1                   2                   0                   0
2                   2                   0                   0
5                   4                   0                   0
===============================================================================
*A:PE1#
```

There are two FDB entries allocated on line card 1, two on line card 2, and four on line card 5. The “PendAlloc” and “PendFree” columns are all zeros.

The following output displays the MAC address status for VPLS 1:

```
*A:PE1# tools dump service id 1 fdb mac-status
===============================================================================
VPLS FDB MAC status at 01/31/2017 08:44:39
===============================================================================
MAC Address         Type                Status : Card list
-------------------------------------------------------------------------------
00:00:00:00:01:01   Select              Allocated : 5
00:00:00:00:01:02   Select              Allocated : 5
00:00:00:00:01:03   Global              Allocated : All
00:00:00:00:01:04   Global              Allocated : All
===============================================================================
*A:PE1#
```

The type and card list for each MAC address in VPLS 1 is displayed. VPLS 1 has learned four MAC addresses: the two local MAC addresses on SAP 5/1/1:1 and the two EvpnS MAC addresses. The local MAC addresses are selective and have FDB entries allocated only on card 5. The global MAC addresses are allocated on all line cards. This output can be further reduced by optionally specifying a specified MAC address, a specific card, and the operational pending state.

3.2.7.2 System FDB Size

The system FDB table size is configurable as follows:

```
configure
    service
        system
            fdb-table-size table-size
```

where table-size can have values in the range from 255999 to 2047999 (2000k).

The default, minimum, and maximum values for the table size are dependent on the chassis type. To support more than 500k MAC addresses, the CPMs provisioned in the system must have at least 16 GB memory. The maximum system FDB table size also limits the maximum FDB table size of any card within the system.

The actual achievable maximum number of MAC addresses depends on the MAC address scale supported by the active cards and whether selective learning is enabled.

If an attempt is made to configure the system FDB table size such that:

- the new size is greater than or equal to the current number of allocated FDB entries, the command succeeds and the new system FDB table size is used
- the new size is less than the number of allocated FDB entries, the command fails with an error message. In this case, the user is expected to reduce the current FDB usage (for example, by deleting statically configured MAC addresses, shutting down EVPN, clearing learned MACs, and so on) to lower the number of allocated MAC addresses in the FDB so that it does not exceed the system FDB table size being configured.

The logic when attempting a rollback is similar; however, when rolling back to a configuration where the system FDB table size is smaller than the current system FDB table size, the system will flush all learned MAC addresses (by performing a **shutdown** then **no shutdown** in all VPLS services) to allow the rollback to continue.

The system FDB table size can be larger than some of the line card FDB sizes, resulting in the possibility that the current number of allocated global MAC addresses is larger than the maximum FDB size supported on some line cards. When a new line card is provisioned, the system checks whether the line card's FDB can accommodate all of the currently allocated global MAC addresses. If it can, then the provisioning succeeds; if it cannot, then the provisioning fails and an error is reported. If the provisioning fails, the number of global MACs allocated must be reduced in the system to a number that the new line card can accommodate, then the **card-type** must be reprovisioned.

3.2.7.3 Per-VPLS Service FDB Size

The following MAC table management features are available for each instance of a SAP or spoke-SDP within a particular VPLS service instance:

- MAC FDB size limits — Allows users to specify the maximum number of MAC FDB entries that are learned locally for a SAP or remotely for a spoke-SDP. If the configured limit is reached, no new addresses will be learned from the SAP or spoke-SDP until at least one FDB entry is aged out or cleared.

- When the limit is reached on a SAP or spoke-SDP, packets with unknown source MAC addresses are still forwarded (this default behavior can be changed by configuration). By default, if the destination MAC address is known, it is forwarded based on the FDB, and if the destination MAC address is unknown, it will be flooded. Alternatively, if discard unknown is enabled at the VPLS service level, any packets from unknown source MAC addresses are discarded at the SAP.
- The log event SAP MAC Limit Reached is generated when the limit is reached. When the condition is cleared, the log event SAP MAC Limit Reached Condition Cleared is generated.
- Disable learning allows users to disable the dynamic learning function on a SAP or a spoke-SDP of a VPLS service instance.
- Disable aging allows users to turn off aging for learned MAC addresses on a SAP or a spoke-SDP of a VPLS service instance.

3.2.7.4 System FDB Size Alarms

High and low watermark alarms give warning when the system MAC FDB usage is high. An alarm is generated when the number of FDB entries allocated in the system FDB reaches 95% of the total system FDB table size and is cleared when it reduces to 90% of the system FDB table size. These percentages are not configurable.

3.2.7.5 Line Card FDB Size Alarms

High and low watermark alarms give warning when a line card's MAC FDB usage is high. An alarm is generated when the number of FDB entries allocated in a line card FDB reaches 95% of its maximum FDB table size and is cleared when it reduces to 90% of its maximum FDB table size. These percentages are not configurable.

3.2.7.6 Per VPLS FDB Size Alarms

The size of the VPLS FDB can be configured with a low watermark and a high watermark, expressed as a percentage of the total FDB size limit. If the actual FDB size grows above the configured high watermark percentage, an alarm is generated. If the FDB size falls below the configured low watermark percentage, the alarm is cleared by the system.

3.2.7.7 Local and Remote Aging Timers

Like a Layer 2 switch, learned MACs within a VPLS instance can be aged out if no packets are sourced from the MAC address for a specified period of time (the aging time). In each VPLS service instance, there are independent aging timers for locally learned MAC and remotely learned MAC entries in the FDB. A local MAC address is a MAC address associated with a SAP because it ingressed on a SAP. A remote MAC address is a MAC address received by an SDP from another router for the VPLS instance. The local-age timer for the VPLS instance specifies the aging time for locally learned MAC addresses, and the remote-age timer specifies the aging time for remotely learned MAC addresses.

In general, the remote-age timer is set to a longer period than the local-age timer to reduce the amount of flooding required for destination unknown MAC addresses. The aging mechanism is considered a low priority process. In most situations, the aging out of MAC addresses happens within tens of seconds beyond the age time. However, it, can take up to two times their respective age timer to be aged out.

3.2.7.8 Disable MAC Aging

The MAC aging timers can be disabled, which will prevent any learned MAC entries from being aged out of the FDB. When aging is disabled, it is still possible to manually delete or flush learned MAC entries. Aging can be disabled for learned MAC addresses on a SAP or a spoke-SDP of a VPLS service instance.

3.2.7.9 Disable MAC Learning

When MAC learning is disabled for a service, new source MAC addresses are not entered in the VPLS FDB, whether the MAC address is local or remote. MAC learning can be disabled for individual SAPs or spoke-SDPs.

3.2.7.10 Unknown MAC Discard

Unknown MAC discard is a feature that discards all packets ingressing the service where the destination MAC address is not in the FDB. The normal behavior is to flood these packets to all end points in the service.

Unknown MAC discard can be used with the disable MAC learning and disable MAC aging options to create a fixed set of MAC addresses allowed to ingress and traverse the service.

3.2.7.11 VPLS and Rate Limiting

Traffic that is normally flooded throughout the VPLS can be rate limited on SAP ingress through the use of service ingress QoS policies. In a service ingress QoS policy, individual queues can be defined per forwarding class to provide shaping of broadcast traffic, MAC multicast traffic, and unknown destination MAC traffic.

3.2.7.12 MAC Move

The MAC move feature is useful to protect against undetected loops in a VPLS topology as well as the presence of duplicate MACs in a VPLS service.

If two clients in the VPLS have the same MAC address, the VPLS will experience a high relearn rate for the MAC. When MAC move is enabled, the 7450 ESS, 7750 SR, or 7950 XRS will shut down the SAP or spoke-SDP and create an alarm event when the threshold is exceeded.

MAC move allows sequential order port blocking. By configuration, some VPLS ports can be configured as “non-blockable”, which allows a simple level of control of which ports are being blocked during loop occurrence. There are two sophisticated control mechanisms that allow blocking of ports in a sequential order:

1. Configuration capabilities to group VPLS ports and to define the order in which they should be blocked
2. Criteria defining when individual groups should be blocked

For the first, configuration CLI is extended by definition of “primary” and “secondary” ports. Per default, all VPLS ports are considered “tertiary” ports unless they are explicitly declared primary or secondary. The order of blocking will always follow a strict order starting from tertiary to secondary, and then primary.

The definition of criteria for the second control mechanism is the number of periods during which the specified relearn rate has been exceeded. The mechanism is based on the cumulative factor for every group of ports. Tertiary VPLS ports are blocked if the relearn rate exceeds the configured threshold during one period while secondary ports are blocked only when relearn rates are exceeded during two consecutive

periods, and so forth. The retry timeout period must be larger than the period before blocking the highest priority port so it sufficiently spans across the period required to block all ports in sequence. The period before blocking the highest priority port is the cumulative factor of the highest configured port multiplied by 5 seconds (the retry timeout can be configured through the CLI).

3.2.7.13 Auto-Learn MAC Protect

This section provides information about auto-learn-mac-protect and restrict-protected-src discard-frame features.

VPLS solutions usually involve learning of MAC addresses in order for traffic to be forwarded to the correct SAP/SDP. If a MAC address is learned on the wrong SAP/SDP, traffic would be re-directed away from its intended destination. This could occur through a mis-configuration, a problem in the network, or by a malicious source creating a DoS attack, and is applicable to any type of VPLS network; for example, mobile backhaul or residential service delivery networks. The auto-learn-mac-protect feature can be used to safeguard against the possibility of MAC addresses being learned on the wrong SAP/SDP.

This feature provides the ability to automatically protect source MAC addresses that have been learned on a SAP or a spoke/mesh SDP and prevent frames with the same protected source MAC address from entering into a different SAP/spoke or mesh SDP.

This is a complementary solution to features such as mac-move and mac-pinning, but has the advantage that MAC moves are not seen and it has a low operational complexity. If a MAC is initially learned on the wrong SAP/SDP, the operator can clear the MAC from the MAC FDB in order for it to be relearned on the correct SAP/SDP.

Two separate commands are used, which provide the configuration flexibility of separating the identification (learning) function from the application of the restriction (discard).

The **auto-learn-mac-protect** and **restrict-protected-src** commands allow the following functions:

- The ability to enable the automatic protection of a learned MAC using the **auto-learn-mac-protect** command under a SAP/spoke or mesh SDP/SHG context

- The ability to discard frames associated with automatically protected MACs instead of shutting down the entire SAP/SDP as with the restrict-protected-src feature. This is enabled using a r**estrict-protected-src discard-frame** command in the SAP/spoke or mesh SDP/SHG context. An optimized alarm mechanism is used to generate alarms related to these discards. The frequency of alarm generation is fixed to be at most one alarm per MAC address per forwarding complex per 10 minutes in a VPLS service.

If auto-learn-mac-protect or restrict-protected-src discard-frame feature is configured under an SHG, the operation applies only to SAPs in the SHG, not to spoke-SDPs in the SHG. If required, these parameters can also be enabled explicitly under specific SAPs/spoke-SDPs within the SHG.

Applying or removing auto-learn-mac-protect or restrict-protected-src discard-frame to/from a SAP, spoke or mesh SDP, or SHG, will clear the MACs on the related objects (for the SHG, this results in clearing the MACs only on the SAPs within the SHG).

The use of restrict-protected-src discard-frame is mutually exclusive with both the **restrict-protected-src [alarm-only]** command and with the configuration of manually protected MAC addresses, using the **mac-protect** command, within a specified VPLS.

The following rules govern the changes to the state of protected MACs:

- Automatically learned protected MACs are subject to normal removal, aging (unless disabled), and flushing, at which time the associated entries are removed from the FDB.
- Automatically learned protected MACs can only move from their learned SAP/spoke or mesh SDP if they enter a SAP/spoke or mesh SDP without restrict-protected-src enabled.

If a MAC address does legitimately move between SAPs/spoke or mesh SDPs after it has been automatically protected on a specified SAP/spoke or mesh SDP (thereby causing discards when received on the new SAP/spoke or mesh SDP), the operator must manually clear the MAC from the FDB for it to be learned in the new/correct location.

MAC addresses that are manually created (using static-mac, static-host with a MAC address specified, or oam mac-populate) will not be protected even if they are configured on a SAP/x SDP that has auto-learn-mac-protect enabled on it. Also, the MAC address associated with a routed VPLS IP interface is protected within its VPLS service such that frames received with this MAC address as the source address are discarded (this is not based on the auto-learn MAC protect function). However, VRRP MAC addresses associated with a routed VPLS IP interface are not protected either in this way or using the auto-learn MAC protect function.

MAC addresses that are dynamically created (learned, using static-host with no MAC address specified, or lease-populate) will be protected when the MAC address is learned on a SAP/x- SDP that has auto-learn-mac-protect enabled on it.

The actions of the following features are performed in the order listed.

1. Restrict-protected-src
2. MAC-pinning
3. MAC-move

3.2.7.13.1 Operation

Figure 63 shows a specific configuration using auto-learn-mac-protect and restrict-protected-src discard-frame in order to describe their operation for the 7750 SR, 7450 ESS, or 7950 XRS.

Figure 63 **Auto-Learn-Mac-Protect Operation**

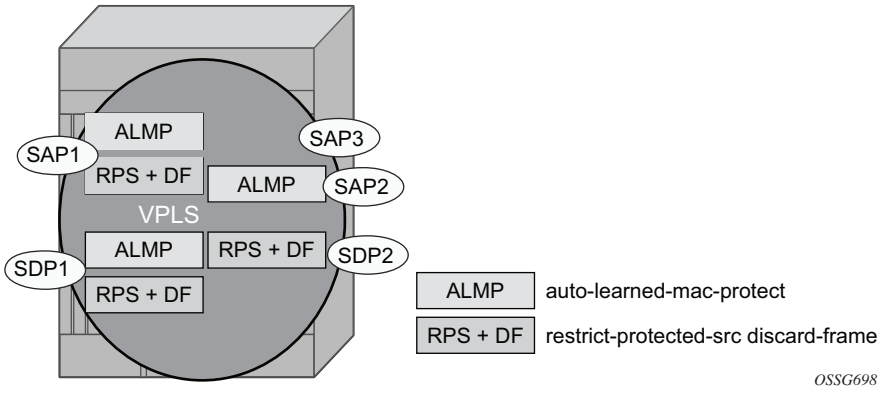

A VPLS service is configured with SAP1 and SDP1 connecting to access devices and SAP2, SAP3, and SDP2 connecting to the core of the network. The auto-learn-mac-protect feature is enabled on SAP1, SAP3, and SDP1 and restrict-protected-src discard-frame is enabled on SAP1, SDP1, and SDP2. The following series of events describes the details of the functionality:

Assume that the FDB is empty at the start of each sequence.

Sequence 1:

1. A frame with source MAC A enters SAP1, MAC A is learned on SAP1, and MAC-A/SAP1 is protected because of the presence of the auto-learn-mac-protect on SAP1.
2. All subsequent frames with source MAC A entering SAP1 are forwarded into the VPLS.
3. Frames with source MAC A enter either SDP1 or SDP2, these frames are discarded, and an alarm indicating MAC A and SDP1/SDP2 is initiated because of the presence of the restrict-protected-src discard-frame on SDP1/SDP2.
4. The above continues, with MAC-A/SAP1 protected in the FDB until MAC A on SAP1 is removed from the FDB.

Sequence 2:

1. A frame with source MAC A enters SAP1, MAC A is learned on SAP1, and MAC-A/SAP1 is protected because of the presence of the auto-learn-mac-protect on SAP1.
2. A frame with source MAC A enters SAP2. As restrict-protected-src is not enabled on SAP2, MAC A is relearned on SAP2 (but not protected), replacing the MAC-A/SAP1 entry in the FDB.
3. All subsequent frames with source MAC A entering SAP2 are forwarded into the VPLS. This is because restrict-protected-src is not enabled SAP2 and auto-learn-mac-protect is not enabled on SAP2, so the FDB is not changed.
4. A frame with source MAC A enters SAP1, MAC A is relearned on SAP1, and MAC-A/SAP1 is protected because of the presence of the auto-learn-mac-protect on SAP1.

Sequence 3:

1. A frame with source MAC A enters SDP2, MAC A is learned on SDP2, but is not protected as auto-learn-mac-protect is not enabled on SDP2.
2. A frame with source MAC A enters SDP1, and MAC A is relearned on SDP1 because previously it was not protected. Consequently, MAC-A/SDP1 is protected because of the presence of the auto-learn-mac-protect on SDP1.

Sequence 4:

1. A frame with source MAC A enters SAP1, MAC A is learned on SAP1, and MAC-A/SAP1 is protected because of the presence of the auto-learn-mac-protect on SAP1.
2. A frame with source MAC A enters SAP3. As restrict-protected-src is not enabled on SAP3, MAC A is relearned on SAP3 and the MAC-A/SAP1 entry is removed from the FDB with MAC-A/SAP3 being added as protected to the FDB (because auto-learn-mac-protect is enabled on SAP3).

3. All subsequent frames with source MAC A entering SAP3 are forwarded into the VPLS.
4. A frame with source MAC A enters SAP1, these frames are discarded and an alarm indicating MAC A and SAP1 is initiated because of the presence of the restrict-protected-src discard-frame on SAP1.

Example Use

Figure 64 shows a possible configuration using auto-learn-mac-protect and restrict-protected-src discard-frame in a mobile backhaul network, with the focus on PE1 for the 7750 SR or 7950 XRS.

Figure 64 **Auto-Learn-Mac-Protect Example**

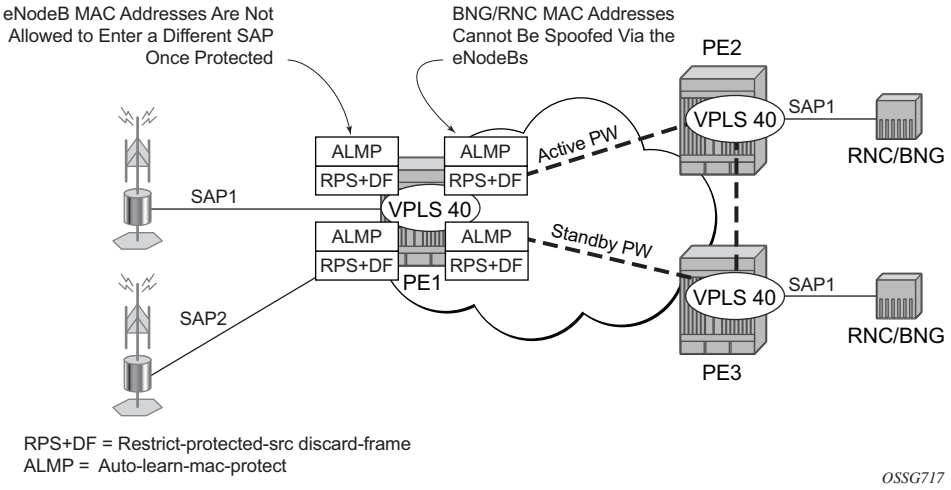

To protect the MAC addresses of the BNG/RNCs on PE1, the **auto-learn-mac-protect** command is enabled on the pseudowires connecting PE1 to PE2 and PE3. Enabling the **restrict-protected-src discard-frame** command on the SAPs toward the eNodeBs will prevent frames with the source MAC addresses of the BNG/RNCs from entering PE1 from the eNodeBs.

The MAC addresses of the eNodeBs are protected in two ways. In addition to the above commands, enabling the **auto-learn-mac-protect** command on the SAPs toward the eNodeBs will prevent the MAC addresses of the eNodeBs being learned on the wrong eNodeB SAP. Enabling the **restrict-protected-src discard-frame** command on the pseudowires connecting PE1 to PE2 and PE3 will protect the eNodeB MAC addresses from being learned on the pseudowires. This may happen if their MAC addresses are incorrectly injected into VPLS 40 on PE2/PE3 from another eNodeB aggregation PE.

The above configuration is equally applicable to other Layer 2 VPLS-based aggregation networks; for example, to business or residential service networks.

3.2.8 Split Horizon SAP Groups and Split Horizon Spoke SDP Groups

Within the context of VPLS services, a loop-free topology within a fully meshed VPLS core is achieved by applying a split horizon forwarding concept that packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications such as DSL aggregation, it is useful to extend this split horizon concept also to groups of SAPs and/or spoke-SDPs. This extension is referred to as a split horizon SAP group or residential bridging.

Traffic arriving on a SAP or a spoke-SDP within a split horizon group will not be copied to other SAPs and spoke-SDPs in the same split horizon group (but will be copied to SAPs/spoke-SDPs in other split horizon groups if these exist within the same VPLS).

3.2.9 VPLS and Spanning Tree Protocol

Nokia’s VPLS service provides a bridged or switched Ethernet Layer 2 network. Equipment connected to SAPs forward Ethernet packets into the VPLS service. The 7450 ESS, 7750 SR, or 7950 XRS participating in the service learns where the customer MAC addresses reside, on ingress SAPs or ingress SDPs.

Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs in the service. If SAPs are connected together, either through misconfiguration or for redundancy purposes, loops can form and flooded packets can keep flowing through the network. The Nokia implementation of the STP is designed to remove these loops from the VPLS topology. This is done by putting one or several SAPs and/or spoke-SDPs in the discarding state.

Nokia’s implementation of the STP incorporates some modifications to make the operational characteristics of VPLS more effective.

The STP instance parameters allow the balancing between resiliency and speed of convergence extremes. Modifying particular parameters can affect the behavior. For information on command usage, descriptions, and CLI syntax, see Configuring a VPLS Service with CLI.

3.2.9.1 Spanning Tree Operating Modes

Per VPLS instance, a preferred STP variant can be configured. The STP variants supported are:

- `rstp` — Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode
- `dot1w` — Compliant with IEEE 802.1w
- `comp-dot1w` — Operation as in RSTP but backwards compatible with IEEE 802.1w (this mode allows interoperability with some MTU types)
- `mstp` — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q-REV/D5.0-09/2005. This mode of operation is only supported in a Management VPLS (M-VPLS).

While the 7450 ESS, 7750 SR, or 7950 XRS initially use the mode configured for the VPLS, it will dynamically fall back (on a per-SAP basis) to STP (IEEE 802.1D-1998) based on the detection of a BPDU of a different format. A trap or log entry is generated for every change in spanning tree variant.

Some older 802.1W compliant RSTP implementations may have problems with some of the features added in the 802.1D-2004 standard. Interworking with these older systems is improved with the comp-dot1w mode. The differences between the RSTP mode and the comp-dot1w mode are:

- The RSTP mode implements the improved convergence over shared media feature, for example, RSTP will transition from discarding to forwarding in 4 seconds when operating over shared media. The comp-dot1w mode does not implement this 802.1D-2004 improvement and transitions conform to 802.1w in 30 seconds (both modes implement fast convergence over point-to-point links).
- In the RSTP mode, the transmitted BPDUs contain the port's designated priority vector (DPV) (conforms to 802.1D-2004). Older implementations may be confused by the DPV in a BPDU and may fail to recognize an agreement BPDU correctly. This would result in a slow transition to a forwarding state (30 seconds). For this reason, in the comp-dot1w mode, these BPDUs contain the port's port priority vector (conforms to 802.1w).

The 7450 ESS, 7750 SR, and 7950 XRS support two BDPU encapsulation formats, and can dynamically switch between the following supported formats (on a per-SAP basis):

- IEEE 802.1D STP
- Cisco PVST

3.2.9.2 Multiple Spanning Tree

The Multiple Spanning Tree Protocol (MSTP) extends the concept of IEEE 802.1w RSTP by allowing grouping and associating VLANs to Multiple Spanning Tree Instances (MSTI). Each MSTI can have its own topology, which provides architecture enabling load balancing by providing multiple forwarding paths. At the same time, the number of STP instances running in the network is significantly reduced as compared to Per VLAN STP (PVST) mode of operation. Network fault tolerance is also improved because a failure in one instance (forwarding path) does not affect other instances.

The Nokia implementation of M-VPLS is used to group different VPLS instances under a single RSTP instance. Introducing MSTP into the M-VPLS allows interoperating with traditional Layer 2 switches in an access network and provides an effective solution for dual homing of many business Layer 2 VPNs into a provider network.

3.2.9.2.1 Redundancy Access to VPLS

The GigE MAN portion of the network is implemented with traditional switches. Using MSTP running on individual switches facilitates redundancy in this part of the network. To provide dual homing of all VPLS services accessing from this part of the network, the VPLS PEs must participate in MSTP.

This can be achieved by configuring M-VPLS on VPLS-PEs (only PEs directly connected to the GigE MAN network), and then assigning different managed-VLAN ranges to different MSTP instances. Typically, the M-VPLS would have SAPs with null encapsulations (to receive, send, and transmit MSTP BPDUs) and a mesh SDP to interconnect a pair of VPLS PEs.

Different access scenarios are displayed in Figure 65 as an example of network diagrams dually connected to the PBB PEs:

- Access Type A — Source devices connected by null or dot1q SAPs
- Access Type B — One QinQ switch connected by QinQ/801ad SAPs

- Access Type C — Two or more ES devices connected by QinQ/802.1ad SAPs

Figure 65 **Access Resiliency**

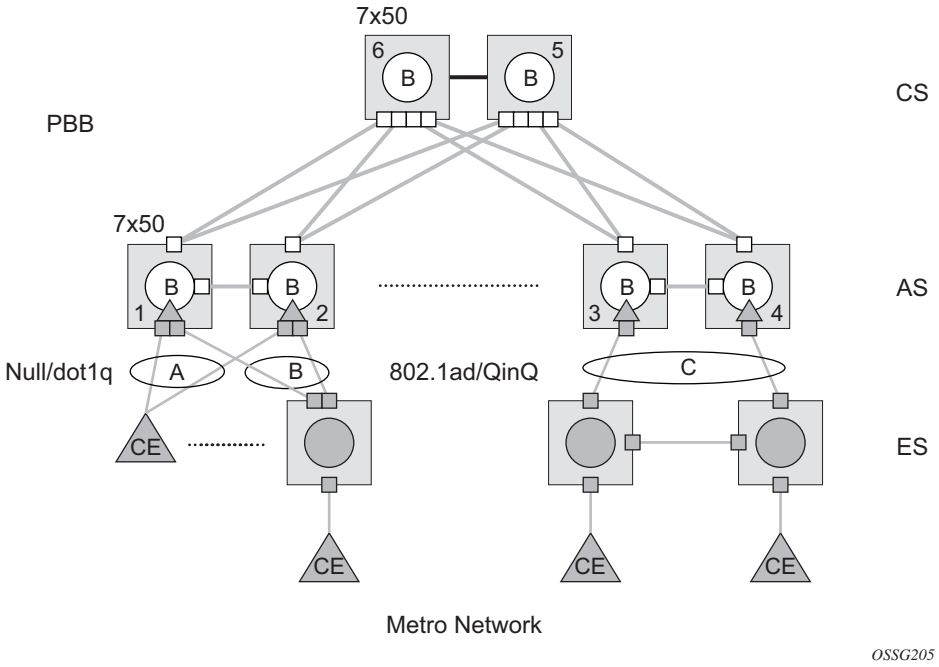

The following mechanisms are supported for the I-VPLS:

- STP/RSTP can be used for all access types.
- M-VPLS with MSTP can be used as is just for access type A. MSTP is required for access type B and C.
- LAG and MC-LAG can be used for access type A and B.
- Split-horizon-group does not require residential.

PBB I-VPLS inherits current STP configurations from the regular VPLS and M-VPLS.

3.2.9.3 MSTP for QinQ SAPs

MSTP runs in a M-VPLS context and can control SAPs from source VPLS instances. QinQ SAPs are supported. The outer tag is considered by MSTP as part of VLAN range control.

3.2.9.4 Provider MSTP

Provider MSTP is specified in IEEE-802.1ad-2005. It uses a provider bridge group address instead of a regular bridge group address used by STP, RSTP, and MSTP BPDUs. This allows for implicit separation of source and provider control planes.

The 802.1ad access network sends PBB PE P-MSTP BPDUs using the specified MAC address and also works over QinQ interfaces. P-MSTP mode is used in PBBN for core resiliency and loop avoidance.

Similar to regular MSTP, the STP mode (for example, PMSTP) is only supported in VPLS services where the m-VPLS flag is configured.

3.2.9.4.1 MSTP General Principles

MSTP represents modification of RSTP that allows the grouping of different VLANs into multiple MSTIs. To enable different devices to participate in MSTIs, they must be consistently configured. A collection of interconnected devices that have the same MST configuration (region-name, revision, and VLAN-to-instance assignment) comprises an MST region.

There is no limit to the number of regions in the network, but every region can support a maximum of 16 MSTIs. Instance 0 is a special instance for a region, known as the Internal Spanning Tree (IST) instance. All other instances are numbered from 1 to 4094. IST is the only spanning-tree instance that sends and receives BPDUs (typically, BPDUs are untagged). All other spanning-tree instance information is included in MSTP records (M-records), which are encapsulated within MSTP BPDUs. This means that a single BPDU carries information for multiple MSTIs, which reduces overhead of the protocol.

Any MSTI is local to an MSTP region and completely independent from an MSTI in other MST regions. Two redundantly connected MST regions will use only a single path for all traffic flows (no load balancing between MST regions or between MST and SST region).

Traditional Layer 2 switches running MSTP protocol assign all VLANs to the IST instance per default. The operator may then “re-assign” individual VLANs to a specified MSTI by configuring per VLAN assignment. This means that as SR-series PE can be considered as the part of the same MST region only if the VLAN assignment to IST and MSTIs is identical to the one of Layer 2 switches in the access network.

3.2.9.4.2 MSTP in the SR-series Platform

The SR-series platform uses a concept of M-VPLS to group different SAPs under a single STP instance. The VLAN range covering SAPs to be managed by a specified M-VPLS is declared under a specific M-VPLS SAP definition. MSTP mode-of-operation is only supported in an M-VPLS.

When running MSTP, by default, all VLANs are mapped to the CIST. At the VPLS level, VLANs can be assigned to specific MSTIs. When running RSTP, the operator must explicitly indicate, per SAP, which VLANs are managed by that SAP.

3.2.9.5 Enhancements to the Spanning Tree Protocol

To interconnect 7450 ESS or 7750 SR OS (PE devices) across the backbone, service tunnels (SDPs) are used. These service tunnels are shared among multiple VPLS instances. The Nokia implementation of the STP incorporates some enhancements to make the operational characteristics of VPLS more effective. The implementation of STP on the router is modified in order to guarantee that service tunnels will not be blocked in any circumstance without imposing artificial restrictions on the placement of the root bridge within the network. The modifications introduced are fully compliant with the 802.1D-2004 STP specification.

When running MSTP, spoke-SDPs cannot be configured. Also, ensure that all bridges connected by mesh SDPs are in the same region. If not, the mesh will be prevented from becoming active (trap is generated).

To achieve this, all mesh SDPs are dynamically configured as either root ports or designated ports. The PE devices participating in each VPLS mesh determine (using the root path cost learned as part of the normal protocol exchange) which of the 7450 ESS, 7750 SR, or 7950 XRS devices is closest to the root of the network. This PE device is internally designated as the primary bridge for the VPLS mesh. As a result of this, all network ports on the primary bridges are assigned the designated port role and therefore remain in the forwarding state.

The second part of the solution ensures that the remaining PE devices participating in the STP instance see the SDP ports as a lower-cost path to the root rather than a path that is external to the mesh. Internal to the PE nodes participating in the mesh, the SDPs are treated as zero-cost paths toward the primary bridge. As a consequence, the path through the mesh is seen as lower cost than any alternative and the PE node will designate the network port as the root port. This approach ensures that network ports always remain in forwarding state.

In combination, these two features ensure that network ports will never be blocked and will maintain interoperability with bridges external to the mesh that are running STP instances.

3.2.9.5.1 L2PT Termination

L2PT is used to transparently transport protocol data units (PDUs) of Layer 2 protocols such as STP, CDP, VTP, PAGP, and UDLD. This allows running these protocols between customer CPEs without involving backbone infrastructure.

The 7450 ESS, 7750 SR, and 7950 XRS routers allow transparent tunneling of PDUs across the VPLS core. However, in some network designs, the VPLS PE is connected to CPEs through a legacy Layer 2 network, rather than having direct connections. In such environments, termination of tunnels through such infrastructure is required.

L2PT tunnels protocol PDUs by overwriting MAC destination addresses at the ingress of the tunnel to a proprietary MAC address such as 01-00-0c-cd-cd-d0. At the egress of the tunnel, this MAC address is then overwritten back to the MAC address of the respective Layer 2 protocol.

The 7450 ESS, 7750 SR, and 7950 XRS routers support L2PT termination for STP BPDUs. More specifically:

- At ingress of every SAP/spoke-SDP that is configured as L2PT termination, all PDUs with a MAC destination address of, 01-00-0c-cd-cd-d0 will be intercepted and their MAC destination address will be overwritten to the MAC destination address used for the corresponding protocol (PVST, STP, RSTP). The type of the STP protocol can be derived from LLC and SNAP encapsulation.
- In egress direction, all STP PDUs received on all VPLS ports will be intercepted and L2PT encapsulation will be performed for SAP/spoke-SDPs configured as L2PT termination points. Because of the implementation reasons, PDU interception and re-direction to CPM can be performed only at ingress. Therefore, to comply with the above requirement, as soon as at least one port of a specified VPLS service is configured as L2PT termination port, redirection of PDUs to CPM will be set on all other ports (SAPs, spoke-SDPs, and mesh SDPs) of the VPLS service.

L2PT termination can be enabled only if STP is disabled in a context of the specified VPLS service.

3.2.9.5.2 BPDU Translation

VPLS networks are typically used to interconnect different customer sites using different access technologies such as Ethernet and bridged-encapsulated ATM PVCs. Typically, different Layer 2 devices can support different types of STP, even if they are from the same vendor. In some cases, it is necessary to provide BPDU translation in order to provide an interoperable e2e solution.

To address these network designs, BPDU format translation is supported on 7450 ESS, 7750 SR, and 7950 XRS devices. If enabled on a specified SAP or spoke-SDP, the system will intercept all BPDUs destined for that interface and perform required format translation such as STP-to-PVST or vice versa.

Similarly, BPDU interception and redirection to the CPM is performed only at ingress, meaning that as soon as at least one port within a specified VPLS service has BPDU translation enabled, all BPDUs received on any of the VPLS ports will be redirected to the CPM.

BPDU translation involves all encapsulation actions that the data path would perform for a specified outgoing port (such as adding VLAN tags depending on the outer SAP and the SDP encapsulation type) and adding or removing all the required VLAN information in a BPDU payload.

This feature can be enabled on a SAP only if STP is disabled in the context of the specified VPLS service.

3.2.9.5.3 L2PT and BPDU Translation

Cisco Discovery Protocol (CDP), Digital Trunking Protocol (DTP), Port Aggregation Protocol (PAGP), Uni-directional Link Detection (UDLD), and Virtual Trunk Protocol (VTP) are supported. These protocols automatically pass the other protocols tunneled by L2PT toward the CPM and all carry the same specific Cisco MAC.

The existing L2PT limitations apply.

- The protocols apply only to VPLS.
- The protocols are mutually exclusive with running STP on the same VPLS as soon as one SAP has L2PT enabled.
- Forwarding occurs on the CPM.

3.2.10 VPLS Redundancy

The VPLS standard (RFC 4762, *Virtual Private LAN Services Using LDP Signaling*) includes provisions for hierarchical VPLS, using point-to-point spoke-SDPs. Two applications have been identified for spoke-SDPs:

- To connect to Multi-Tenant Units (MTUs) to PEs in a metro area network
- To interconnect the VPLS nodes of two metro networks.

In both applications, the spoke-SDPs serve to improve the scalability of VPLS. While node redundancy is implicit in non-hierarchical VPLS services (using a full mesh of SDPs between PEs), node redundancy for spoke-SDPs needs to be provided separately.

Nokia routers have implemented special features for improving the resilience of hierarchical VPLS instances, in both MTU and inter-metro applications.

3.2.10.1 Spoke SDP Redundancy for Metro Interconnection

When two or more meshed VPLS instances are interconnected by redundant spoke-SDPs (as shown in Figure 66), a loop, in the topology results. To remove such a loop from the topology, STP can be run over the SDPs (links) that form the loop such that one of the SDPs is blocked. As running STP in each and every VPLS in this topology is not efficient, the node includes functionality that can associate a number of VPLSs to a single STP instance running over the redundant SDPs. Therefore, node redundancy is achieved by running STP in one VPLS, and applying the conclusions of this STP to the other VPLS services. The VPLS instance running STP is referred to as the “management VPLS” or M-VPLS.

If the active node fails, STP on the management VPLS in the standby node will change the link states from disabled to active. The standby node will then broadcast a MAC flush LDP control message in each of the protected VPLS instances, so that the address of the newly active node can be relearned by all PEs in the VPLS.

It is possible to configure two management VPLS services, where both VPLS services have different active spokes (this is achieved by changing the path cost in STP). By associating different user VPLSs with the two management VPLS services, load balancing across the spokes can be achieved.

Figure 66 **HVPLS with Spoke Redundancy**

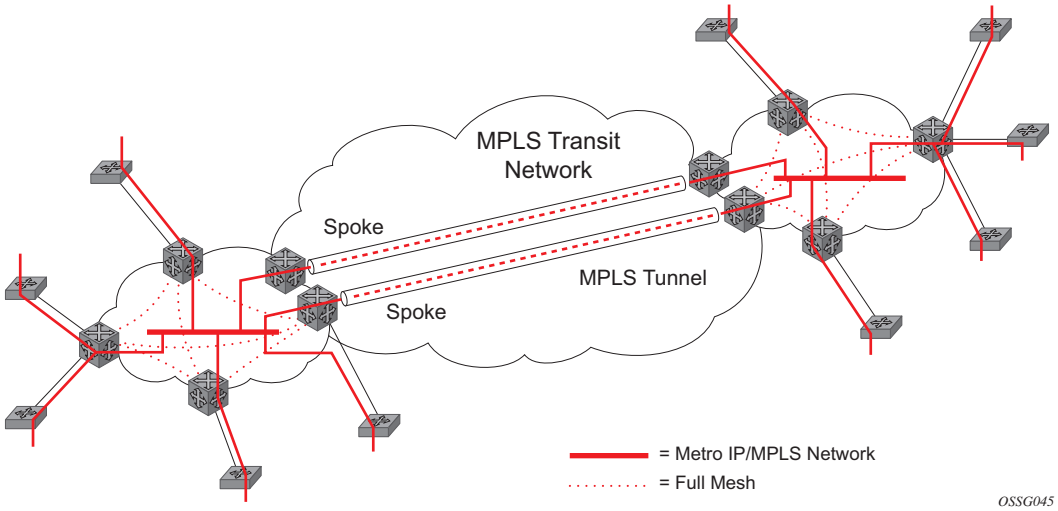

3.2.10.2 Spoke SDP Based Redundant Access

This feature provides the ability to have a node deployed as MTUs (Multi-Tenant Unit Switches) to be multi-homed for VPLS to multiple routers deployed as PEs without requiring the use of M-VPLS.

In the configuration example displayed in Figure 66, the MTUs have spoke-SDPs to two PE devices. One is designated as the primary and one as the secondary spoke-SDP. This is based on a precedence value associated with each spoke.

The secondary spoke is in a blocking state (both on receive and transmit) as long as the primary spoke is available. When the primary spoke becomes unavailable (due to link failure, PEs failure, and so on), the MTU immediately switches traffic to the backup spoke and starts receiving traffic from the standby spoke. Optional revertive operation (with configurable switch-back delay) is supported. Forced manual switchover is also supported.

To speed up the convergence time during a switchover, MAC flush is configured. The MTUs generate a MAC flush message over the newly unblocked spoke when a spoke change occurs. As a result, the PEs receiving the MAC flush will flush all MACs associated with the impacted VPLS service instance and forward the MAC flush to the other PEs in the VPLS network if “propagate-mac-flush” is enabled.

3.2.10.3 Inter-Domain VPLS Resiliency Using Multi-Chassis Endpoints

Inter-domain VPLS refers to a VPLS deployment where sites may be located in different domains. An example of inter-domain deployment can be where different Metro domains are interconnected over a Wide Area Network (Metro1-WAN-Metro2) or where sites are located in different autonomous systems (AS1-ASBRs-AS2).

Multi-chassis endpoint (MC-EP) provides an alternate solution that does not require RSTP at the gateway VPLS PEs while still using pseudowires to interconnect the VPLS instances located in the two domains. It is supported in both VPLS and PBB-VPLS on the B-VPLS side.

MC-EP expands the single chassis endpoint based on active-standby pseudowires for VPLS shown in Figure 67.

Figure 67 **HVPLS Resiliency Based on AS Pseudowires**

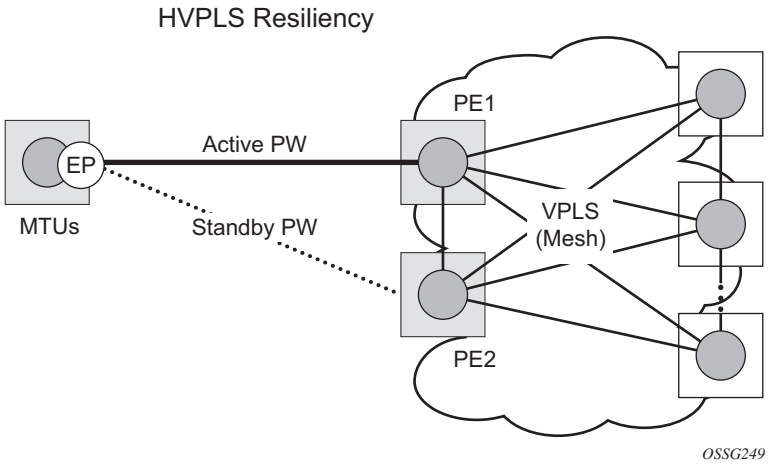

The active-standby pseudowire solution is appropriate for the scenario when only one VPLS PE (MTU-s) needs to be dual-homed to two core PEs (PE1 and PE2). When multiple VPLS domains need to be interconnected, the above solution provides a single point of failure at the MTU-s. The example shown in Figure 68 can be used.

Figure 68 **Multi-Chassis Pseudowire Endpoint for VPLS**

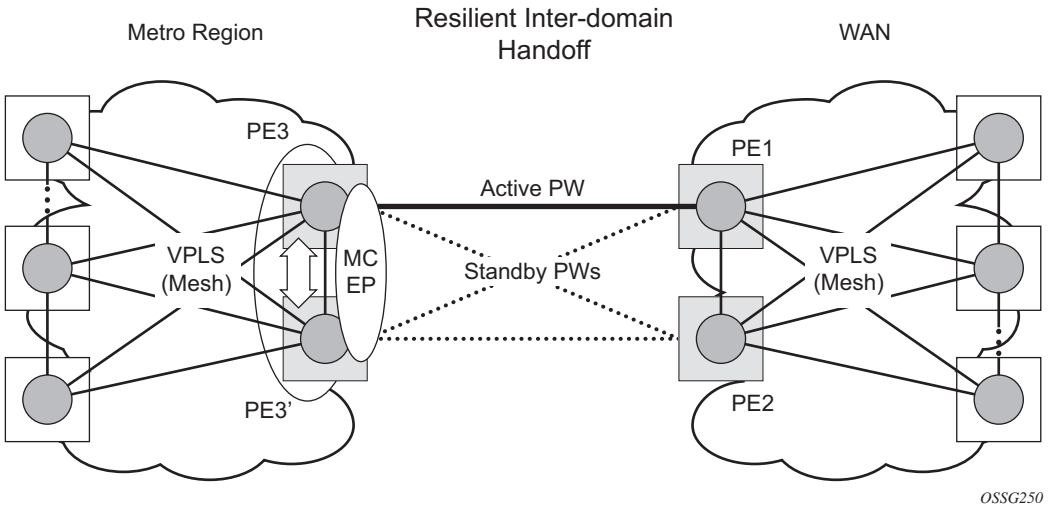

The two gateway pairs, PE3-PE3 and PE1-PE2, are interconnected using a full mesh of four pseudowires out of which only one pseudowire is active at any time.

The concept of pseudowire endpoint for VPLS provides multi-chassis resiliency controlled by the MC-EP pair, PE3-PE3 in this example. This scenario, referred to as multi-chassis pseudowire endpoint for VPLS, provides a way to group pseudowires distributed between PE3 and PE3 chassis in a virtual endpoint that can be mapped to a VPLS instance.

The MC-EP inter-chassis protocol is used to ensure configuration and status synchronization of the pseudowires that belong to the same MC-EP group on PE3 and PE3. Based on the information received from the peer shelf and the local configuration, the master shelf will make a decision on which pseudowire will become active.

The MC-EP solution is built around the following components:

- Multi-chassis protocol used to perform the following functions:
  - Selection of master chassis.
  - Synchronization of the pseudowire configuration and status.
  - Fast detection of peer failure or communication loss between MC-EP peers using either centralized BFD, if configured, or its own keep-alive mechanism.
- T-LDP signaling of pseudowire status:
  - Informs the remote PEs about the choices made by the MC-EP pair.

- Pseudowire data plane — Represented by the four pseudowires interconnecting the gateway PEs.
  - Only one of the pseudowires is activated based on the primary/secondary, preference configuration, and pseudowire status. In case of a tie, the pseudowire located on the master chassis will be chosen.
  - The rest of the pseudowires are blocked locally on the MC-EP pair and on the remote PEs as long as they implement the pseudowire active/standby status.

3.2.10.3.1 Fast Detection of Peer Failure using BFD

Although the MC-EP protocol has its own keep-alive mechanisms, sharing a common mechanism for failure detection with other protocols (for example, BGP, RSVP-TE) scales better. MC-EP can be configured to use the centralized BFD mechanism.

Similar to other protocols, MC-EP will register with BFD if the **bfd-enable** command is active under the **config**>**redundancy**>**multi-chassis**>**peer**>**mc-ep** context. As soon as the MC-EP application is activated using no shutdown, it tries to open a new BFD session or register automatically with an existing one. The source-ip configuration under redundancy multi-chassis peer-ip is used to determine the local interface while the peer-ip is used as the destination IP for the BFD session. After MC-EP registers with an active BFD session, it will use it for fast detection of MC-EP peer failure. If BFD registration or BFD initialization fails, the MC-EP will keep using its own keep-alive mechanism and it will send a trap to the NMS signaling the failure to register with/open a BFD session.

To minimize operational mistakes and wrong peer interpretation for the loss of BFD session, the following additional rules are enforced when the MC-EP is registering with a certain BFD session:

- Only the centralized BFD sessions using system or loopback IP interfaces (source-ip parameter) are accepted in order for MC-EP to minimize the false indication of peer loss.
- If the BFD session associated with MC-EP protocol is using a certain interface (system/loopback), the following actions are not allowed under the interface: IP address change, “shutdown”, “no bfd” commands. If one of these actions is required under the interface, the operator needs to disable BFD using the following procedures:
  - The **no bfd-enable** command in the **config**>**redundancy**>**multi-chassis**>**peer**>**mc-ep** context — this is the recommended procedure.

– The **shutdown** command in the **config**>**redundancy**>**multi-chassis**>**peer**>**mc-ep** or from under **config**>**redundancy**>**multi-chassis**>**peer** contexts.

MC-EP keep-alives are still exchanged for the following reasons:

- As a backup; if the BFD session does not come up or is disabled, the MC-EP protocol will use its own keep-alives for failure detection.
- To ensure the database is cleared if the remote MC-EP peer is shutdown or miss-configured (each x seconds; one second suggested as default).

If MC-EP de-registers with BFD using the “no bfd-enable” command, the following processing steps occur:

- Local peer indicates to the MC-EP peer the fact that the local BFD is being disabled using MC-EP peer-config-TLV fields ([BFD local: BFD remote]). This is done to avoid wrong interpretation of BFD session loss.
- Remote peer acknowledges reception indicating through the same peer-config-TLV fields that it is de-registering with the BFD session.
- Both MC-EP peers de-register and are going to use only keep-alives for failure detection.
- There should be no pseudowire status change during this process.

Traps are sent when the status of the monitoring of the MC-EP session through BFD changes in the following instances:

- When red/mc/peer is no shutdown and BFD is not enabled, send a notification indicating BFD is not monitoring MC-EP peering session.
- When BFD changes to open, send a notification indicating BFD is monitoring MC-EP peering session.
- When BFD changes to down/close, send a notification indicating BFD is not monitoring MC-EP peering session.

3.2.10.3.2 MC-EP Passive Mode

The MC-EP mechanisms are built to minimize the possibility of loops. It is possible that human error could create loops through the VPLS service. One way to prevent loops is to enable the MAC move feature in the gateway PEs (PE3, PE3', PE1 and PE2).

An MC-EP passive mode can also be used on the second PE pair, PE1 and PE2, as a second layer of protection to prevent any loops from occurring if the operator introduces operational errors on the MC-EP PE3, PE3 pair.

Figure 69 **MC-EP in Passive Mode**

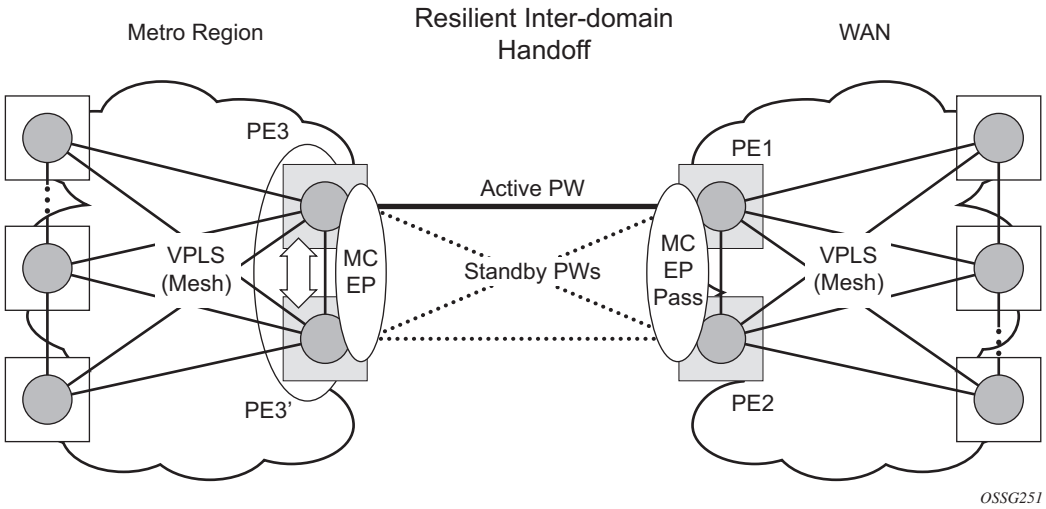

When in passive mode, the MC-EP peers stay dormant as long as one active pseudowire is signaled from the remote end. If more than one pseudowire belonging to the passive MC-EP becomes active, then the PE1 and PE2 pair applies the MC-EP selection algorithm to select the best choice and blocks all others. No signaling is sent to the remote pair to avoid flip-flop behavior. A trap is generated each time MC-EP in passive mode activates. Every occurrence of this kind of trap should be analyzed by the operator as it is an indication of possible misconfiguration on the remote (active) MC-EP peering.

In order for the MC-EP passive mode to work, the pseudowire status signaling for active/standby pseudowires should be enabled. This involves the following CLI configurations:

For the remote MC-EP PE3, PE3 pair:

```
config>service>vpls>endpoint# no suppress-standby-signaling
```

When MC-EP passive mode is enabled on the PE1 and PE2 pair, the following command is always enabled internally, regardless of the actual configuration:

```
config>service>vpls>endpoint no ignore-standby-signaling
```

3.2.10.4 Support for Single Chassis Endpoint Mechanisms

In cases of SC-EP, there is a consistency check to ensure that the configuration of the member pseudowires is the same. For example, mac-pining, mac-limit, and ignore standby signaling must be the same. In the MC-EP case, there is no consistency check between the member endpoints located on different chassis. The operator must carefully verify the configuration of the two endpoints to ensure consistency.

The following rules apply for suppress-standby-signaling and ignore-standby parameters:

- Regular MC-EP mode (non-passive) will follow the suppress-standby-signaling and ignore-standby settings from the related endpoint configuration.
- For MC-EP configured in passive mode, the following settings will be used, regardless of previous configuration: **suppress-standby-sig** and **no ignore-standby-sig**. It is expected that when passive mode is used at one side, the regular MC-EP side will activate signaling with **no suppress-stdby-sig**.
- When passive mode is configured in just one of the nodes in the MC-EP peering, the other node will be forced to change to passive mode. A trap is sent to the operator to signal the wrong configuration.

This section also describes how the main mechanisms used for single chassis endpoint are adapted for the MC-EP solution.

3.2.10.4.1 MAC Flush Support in MC-EP

In an MC-EP scenario, failure of a pseudowire or gateway PE will determine activation of one of the next best pseudowires in the MC-EP group. This section describes the MAC flush procedures that can be applied to ensure black-hole avoidance.

Figure 70 shows a pair of PE gateways (PE3 and PE3) running MC-EP toward PE1 and PE2 where F1 and F2 are used to indicate the possible direction of the MAC flush signaled using T-LDP MAC withdraw message. PE1 and PE2 can only use regular VPLS pseudowires and do not have to use an MC-EP or a regular pseudowire endpoint.

Figure 70 **MAC Flush in the MC-EP Solution**

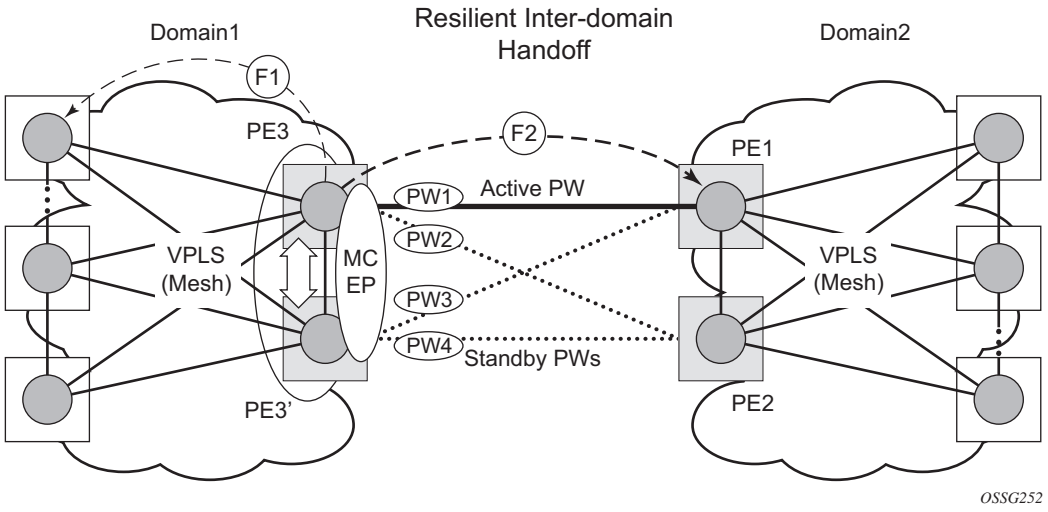

Regular MAC flush behavior will apply for the LDP MAC withdraw sent over the T-LDP sessions associated with the active pseudowire in the MC-EP; for example PE3 to PE1. That is, for any TCN events or failures associated with SAPs or pseudowires not associated with the MC-EP.

The following MAC flush behaviors apply to changes in the MC-EP pseudowire selection:

- If the local PW2 becomes active on PE3:
  - On PE3, the MACs mapped to PW1 are moved to PW2.
  - A T-LDP “flush-all-but-mine” message is sent toward PE2 in F2 direction and is propagated by PE2 in the local VPLS mesh.
  - No MAC flush is sent to F1 direction from PE3.
- If one of the pseudowires on the pair PE3 becomes active; for example PW4:
  - On PE3, the MACs mapped to PW1 are flushed, same as a regular endpoint.
  - PE3 must be configured with **send-flush-on-failure** to send a T-LDP “flush-all-from-me” message toward VPLS mesh in the F1 direction.
  - PE3 sends a T-LDP **flush-all-but-mine** message toward PE2 in the F2 direction, which is propagated by PE2 in the local VPLS mesh. When MC-EP is in passive mode and the first spoke becomes active, a **no mac flush-all-but-mine** message will be generated.

3.2.10.4.2 Block-on-Mesh-Failure Support in MC-EP Scenario

The following rules describe how the block-on-mesh-failure operates with the MC-EP solution (see Figure 70):

- If PE3 does not have any forwarding path toward Domain1 mesh, it should block both PW1 and PW2 and inform PE3 so one of its pseudowires can be activated.
- To allow the use of block-on-mesh-failure for MC-EP, a block-on-mesh-failure parameter can be specified in the **config**>**service**>**vpls**>**endpoint** context with the following rules:
  - The default is **no block-on-mesh-failure** to allow for easy migration.
  - For a spoke-SDP to be added under an endpoint, the setting for its **block-on-mesh-failure** parameter must be in synchronization with the endpoint parameter.
  - After the spoke-SDP is added to an endpoint, the configuration of its **block-on-mesh-failure** parameter is disabled. A change in endpoint configuration for the **block-on-mesh-failure** parameter is propagated to the individual spoke-SDP configuration.
  - When a spoke-SDP is removed from the endpoint group, it will inherit the last configuration from the endpoint parameter.
  - Adding an MC-EP under the related endpoint configuration does not affect the above behavior.

  Before Release 7.0, the **block-on-mesh-failure** command could not be enabled under **config**>**service**>**vpls**>**endpoint** context. In order for a spoke-SDP to be added to an (single-chassis) endpoint, its **block-on-mesh-failure** had to be disabled (**config**>**service**>**vpls**>**spoke-sdp**>**no block-on-mesh-failure**). Then, the configuration of **block-on-mesh-failure** under a spoke-SDP is blocked.
- If **block-on-mesh-failure** is enabled on PE1 and PE2, these PEs will signal pseudowire standby status toward the MC-EP PE pair. PE3 and PE3 should consider the pseudowire status signaling from remote PE1 and PE2 when making the selection of the active pseudowire.

3.2.10.4.3 Support for Force Spoke SDP in MC-EP

In a regular (single chassis) endpoint scenario, the following command can be used to force a specific SDP binding (pseudowire) to become active:

```
tools perform service id service-id endpoint endpoint-name
force
```

In the MC-EP case, this command has a similar effect when there is a single forced SDP binding in an MC-EP. The forced SDP binding (pseudowire) will be elected as active.

However, when the command is run at the same time as both MC-EP PEs, when the endpoints belong to the same MC-EP, the regular MC-EP selection algorithm (for example, the operational status ⇒precedence value) will be applied to determine the winner.

3.2.10.4.4 Revertive Behavior for Primary Pseudowires in an MC-EP

For a single-chassis endpoint, a revert-time command is provided under the VPLS endpoint. Refer to the VPLS Service Configuration Command Reference for syntax and command usage information.

In a regular endpoint, the revert-time setting affects just the pseudowire defined as primary (precedence 0). For a failure of the primary pseudowire followed by restoration, the revert-timer is started. After it expires, the primary pseudowire takes the active role in the endpoint. This behavior does not apply for the case when both pseudowires are defined as secondary: that is, if the active secondary pseudowire fails and is restored, it will stay in standby until a configuration change or a force command occurs.

In the MC-EP case, the revertive behavior is supported for pseudowire defined as primary (precedence 0). The following rules apply:

- The revert-time setting under each individual endpoint control the behavior of the local primary pseudowire if one is configured under the local endpoint.
- The secondary pseudowires behave as in the regular endpoint case.

3.2.10.5 Using B-VPLS for Increased Scalability and Reduced Convergence Times

The PBB-VPLS solution can be used to improve scalability of the solution and to reduce convergence time. If PBB-VPLS is deployed starting at the edge PEs, the gateway PEs will contain only B-VPLS instances. The MC-EP procedures described for regular VPLS apply.

PBB-VPLS can also be enabled just on the gateway MC-EP PEs, as shown in Figure 71.

Figure 71 **MC-EP with B-VPLS**

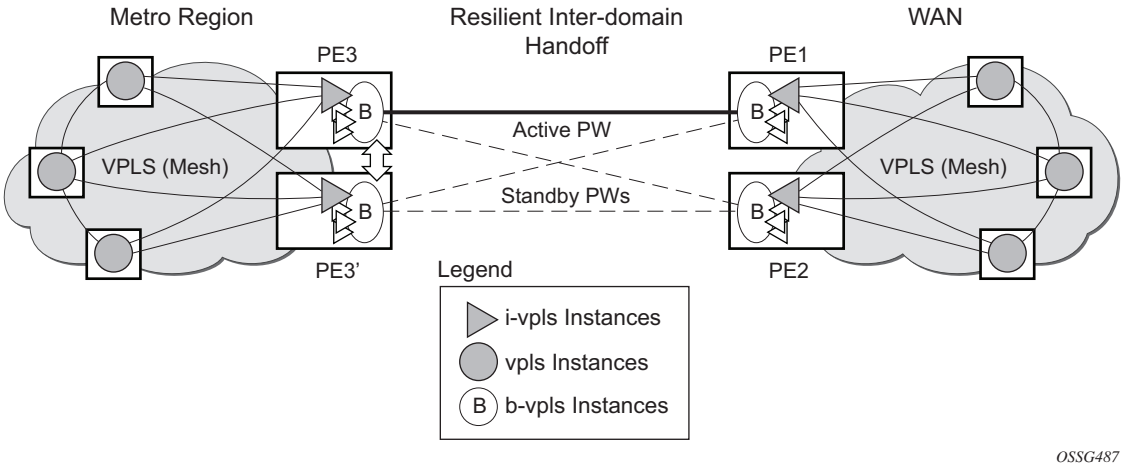

Multiple I-VPLS instances may be used to represent in the gateway PEs the customer VPLS instances using the PBB-VPLS M:1 model described in the PBB section. A backbone VPLS (B-VPLS) is used in this example to administer the resiliency for all customer VPLS instances at the domain borders. Just one MC-EP is required to be configured in the B-VPLS to address hundreds or even thousands of customer VPLS instances. If load balancing is required, multiple B-VPLS instances may be used to ensure even distribution of the customers across all the pseudowires interconnecting the two domains. In this example, four B-VPLS will be able to load share the customers across all four possible pseudowire paths.

The use of MC-EP with B-VPLS is strictly limited to cases where VPLS mesh exists on both sides of a B-VPLS. For example, active/standby pseudowires resiliency in the I-VPLS context where PE3, PE3' are PErs cannot be used because there is no way to synchronize the active/standby selection between the two domains.

For a similar reason, MC-LAG resiliency in the I-VPLS context on the gateway PEs participating in the MC-EP (PE3, PE3) should not be used.

For the PBB topology in Figure 71, block-on-mesh-failure in the I-VPLS domain will not have any effect on the B-VPLS MC-EP side. That is because mesh failure in one I-VPLS should not affect other I-VPLS sharing the same B-VPLS.

3.2.10.6 MAC Flush Additions for PBB VPLS

The scenario shown in Figure 72 is used to define the blackholing problem in PBB-VPLS using MC-EP.

Figure 72 **MC-EP with B-VPLS Failure Scenario**

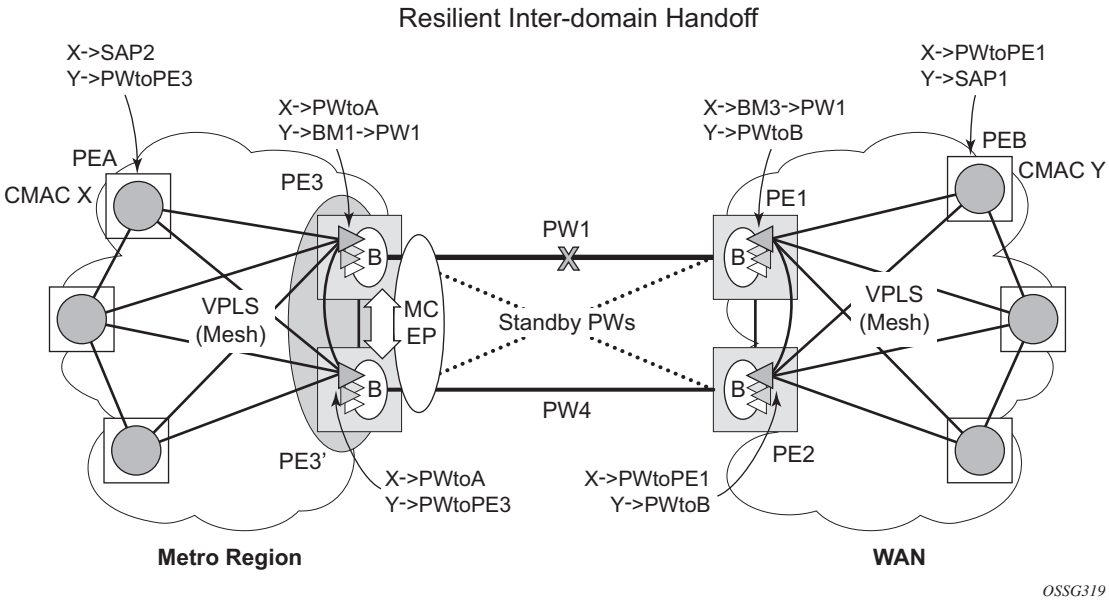

In the topology shown in Figure 72, PEA and PEB are regular VPLS PEs participating in the VPLS mesh deployed in the metro and WAN region, respectively. As the traffic flows between CEs with CMAC X and CMAC Y, the FDB entries in blue are installed. A failure of the active PW1 will result in the activation of PW4 between PE3 and PE2 in this example. An LDP flush-all-but-mine will be sent from PE3 to PE2 to clear the B-VPLS FDBs. The traffic between CMAC X and CMAC Y will be blackholed as long as the entries from the VPLS and I-VPLS FDBs along the path are not removed. This may take as long as 300 seconds, the usual aging timer used for MAC entries in a VPLS FDB.

A MAC flush is required in the I-VPLS space from PBB PEs to PEA and PEB to avoid blackholing in the regular VPLS space.

In the case of a regular VPLS, the following procedure is used:

- PE3 sends a flush-all-from-me toward its local blue I-VPLS mesh to PE3 and PEA when its MC-EP becomes disabled
- PE3 sends a flush-all-but-mine on the active PW4 to PE2, which is then propagated by PE2 (propagate-mac-flush must be on) to PEB in the WAN I-VPLS mesh.

For consistency, a similar procedure is used for the B-VPLS case as shown in Figure 73.

Figure 73 **MC-EP with B-VPLS MAC Flush Solution**

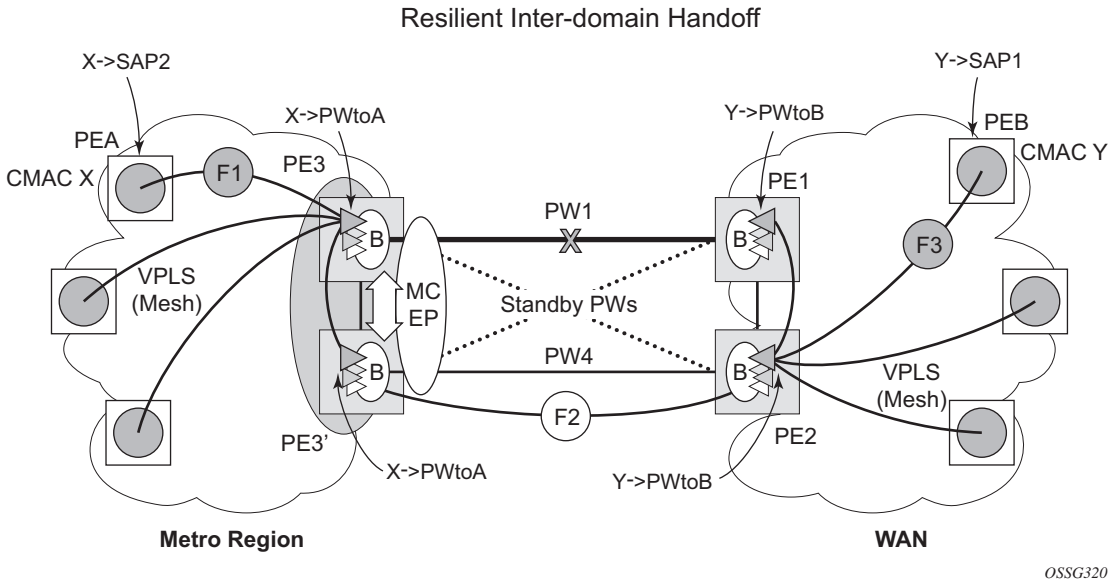

In this example, the MC-EP activates B-VPLS PW4 because of either a link/node failure or because of an MC-EP selection re-run that affected the previously active PW1. As a result, the endpoint on PE3 containing PW1 goes down.

The following steps apply:

- PE3 sends in the local I-VPLS context an LDP flush-all-from-me message (marked with F1) to PEA and to the other regular VPLS PEs, including PE3. The following command enables this behavior on a per I-VPLS basis: **config>service>vpls ivpls>send-flush-on-bvpls-failure.**
  - Result: PEA, PE3, and the other local VPLS PEs in the metro clear the VPLS FDB entries associated with PW to PE3.
- PE3 clears the entries associated with PW1 and sends in the B-VPLS context an LDP flush-all-but-mine (marked with F2) toward PE2 on the active PW4.
  - Result: PE2 clears the B-VPLS FDB entries not associated with PW4.
- PE2 propagates the MAC flush-all-but-mine (marked with F3) from B-VPLS in the related I-VPLS contexts toward all participating VPLS PEs; for example, in the blue I-VPLS to PEB, PE1. It also clears all the CMAC entries associated with I-VPLS pseudowires.

  The following command enables this behavior on a per I-VPLS basis:

  config>service>vpls ivpls>propagate-mac-flush-from-bvpls
  - Result: PEB, PE1, and the other local VPLS PEs in the WAN clear the VPLS FDB entries associated with PW to PE2.

  - This command does not control the propagation in the related I-VPLS of the B-VPLS LDP MAC flush containing a PBB TLV (BMAC and ISID list).
- Similar to regular VPLS, LDP signaling of the MAC flush will follow the active topology; for example, no MAC flush will be generated on standby pseudowires.

Other failure scenarios are addressed using the same or a subset of the above steps:

- If the pseudowire (PW2) in the same endpoint with PW1 becomes active instead of PW4, there will be no MAC flush of F1 type.
- If the pseudowire (PW3) in the same endpoint becomes active instead of PW4, the same procedure applies.

For an SC/MC endpoint configured in a B-VPLS, failure/deactivation of the active pseudowire member always generates a local MAC flush of all the BMAC associated with the pseudowire. It never generates a MAC move to the newly active pseudowire even if the endpoint stays up. That is because in SC-EP/MC-EP topology, the remote PE might be the terminating PBB PE and may not be able to reach the BMAC of the other remote PE. Therefore, connectivity between them exists only over the regular VPLS mesh.

For the same reasons, Nokia recommends that static BMAC not be used on SC/MC endpoints.

3.2.11 VPLS Access Redundancy

A second application of hierarchical VPLS is using MTUs that are not MPLS-enabled that must have Ethernet links to the closest PE node. To protect against failure of the PE node, an MTU can be dual-homed and have two SAPs on two PE nodes.

There are several mechanisms that can be used to resolve a loop in an access circuit; however, from operation perspective, they can be subdivided into two groups:

- STP-based access, with or without M-VPLS.
- Non-STP based access using mechanisms such as MC-LAG, MC-APS, MC-Ring.

3.2.11.1 STP-based Redundant Access to VPLS

Figure 74 **Dual-homed MTU-s in Two-Tier Hierarchy H-VPLS**

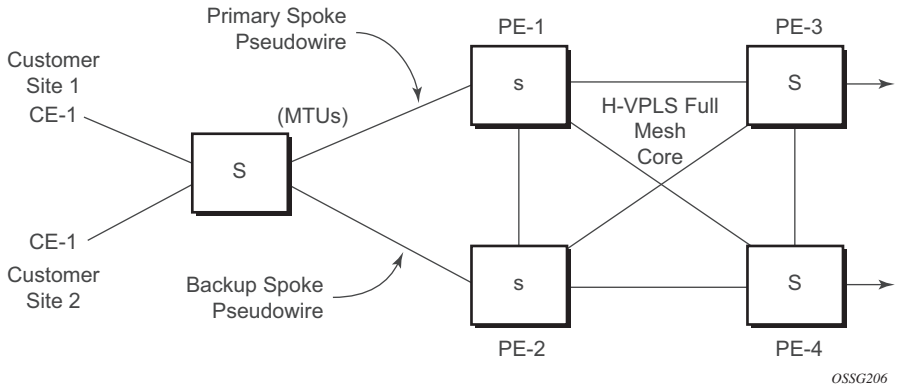

In the configuration shown in Figure 74, STP is activated on the MTU and two PEs in order to resolve a potential loop. STP only needs to run in a single VPLS instance, and the results of the STP calculations are applied to all VPLSs on the link.

In this configuration, the scope of the STP domain is limited to MTU and PEs, while any topology change needs to be propagated in the whole VPLS domain including mesh SDPs. This is done by using so called MAC-flush messages defined by RFC 4762. In the case of STP as an loop resolution mechanism, every TCN (Topology Change Notification) received in a context of STP instance is translated into LDP-MAC address withdrawal message (also referred to as MAC-flush message) requesting to clear all FDB entries except the ones learned from the originating PE. Such messages are sent to all PE peers connected through SDPs (mesh and spoke) in the context of VPLS services, which are managed by the specified STP instance.

3.2.11.2 Redundant Access to VPLS Without STP

The Nokia implementation also includes alternative methods for providing a redundant access to Layer 2 services, such as MC-LAG, MC-APS, or MC-Ring. Also in this case, the topology change event needs to be propagated into VPLS topology in order to provide fast convergence.

Figure 66 illustrates a dual-homed connection to VPLS service (PE-A, PE-B, PE-C, PE-D) and operation in case of link failure (between PE-C and L2-B). Upon detection of a link failure, PE-C will send MAC-Address-Withdraw messages, which will indicate to all LDP peers that they should flush all MAC addresses learned from PE-C. This will lead to a broadcasting of packets addressing affected hosts and relearning in case an alternative route exists.

The message described here is different than the message described in the previous section and in RFC 4762, *Virtual Private LAN Services Using LDP Signaling*. The difference is in the interpretation and action performed in the receiving PE. According to the standard definition, upon receipt of a MAC withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed.

This section specifies that all MAC addresses learned from the source are flushed. This message has been implemented as an LDP address message with vendor-specific type, length, value (TLV), and is called the flush-mine message.

The advantage of this approach (as compared to RSTP-based methods) is that only MAC-affected addresses are flushed and not the full forwarding database. While this method does not provide a mechanism to secure alternative loop-free topology, the convergence time is dependent on the speed that the specified CE device will open an alternative link (L2-B switch in Figure 57) as well as on the speed that PE routers will flush their FDB.

In addition, this mechanism is effective only if PE and CE are directly connected (no hub or bridge) as it reacts to the physical failure of the link.

3.2.12 Object Grouping and State Monitoring

This feature introduces a generic operational group object that associates different service endpoints (pseudowires, SAPs, IP interfaces) located in the same or in different service instances.

The operational group status is derived from the status of the individual components using certain rules specific to the application using the concept. A number of other service entities, the monitoring objects, can be configured to monitor the operational group status and to perform certain actions as a result of status transitions. For example, if the operational group goes down, the monitoring objects will be brought down.

3.2.12.1 VPLS Applicability — Block on VPLS a Failure

This concept is used in VPLS to enhance the existing BGP MH solution by providing a block-on-group failure function similar with the block-on-mesh failure feature implemented for LDP VPLS, mesh. On the PE selected as the Designated Forwarder (DF), if the rest of the VPLS endpoints fail (pseudowire spokes/pseudowire mesh and/or SAPs), there is no path forward for the frames sent to the MH site selected as DF. The status of the VPLS endpoints, other than the MH site, is reflected by bringing down/up the objects associated with the MH site.

Support for the feature is provided initially in VPLS and B-VPLS instance types for LDP VPLS, with or without BGP-AD and for BGP VPLS. The following objects may be placed as components of an operational group: BGP VPLS pseudowires, SAPs, spoke-pseudowire, BGP-AD pseudowires. The following objects are supported as monitoring objects: BGP MH site, individual SAP, spoke-pseudowire.

The following rules apply:

- An object can only belong to one group at a time.
- An object that is part of a group cannot monitor the status of a group.
- An object that monitors the status of a group cannot be part of a group.
- An operational group may contain any combination of member types: SAP, spoke-pseudowire, BGP-AD or BGP VPLS pseudowires.
- An operational group may contain members from different VPLS service instances.
- Objects from different services may monitor the operational group.
- The operational group feature may co-exist in parallel with the block-on-mesh failure feature as long as they are running in different VPLS instances.

There are two steps involved in enabling the block-on-mesh failure feature in a VPLS scenario:

1. Identify a set of objects whose forwarding state should be considered as a whole group then group them under an operational group using the **oper-group** CLI command.
2. Associate other existing objects (clients) with the **oper-group** command using the **monitor-group** CLI command; its forwarding state will be derived from the related operational group state.

The status of the operational group (oper-group) is dictated by the status of one or more members according to the following rule:

- The oper-group goes down if all the objects in the oper-group go down; the oper-group comes up if at least one of the components is up.
- An object in the group is considered down if it is not forwarding traffic in at least one direction. That could be because the operational state is down or the direction is blocked through some resiliency mechanisms.
- If a group is configured but no members are specified yet, its status is considered up. As soon as the first object is configured, the status of the operational group is dictated by the status of the provisioned members.

- For BGP-AD or BGP VPLS pseudowires associated with the oper-group (under the **config**>**service-vpls**>**bgp**>**pw-template-binding** context), the status of the oper-group is down as long as the pseudowire members are not instantiated (auto-discovered and signaled).

A simple configuration example is described for the case of a BGP VPLS mesh used to interconnect different customer locations. If we assume a customer edge (CE) device is dual-homed to two PEs using BGP MH, the following configuration steps apply:

- The **oper-group bgp-vpls-mesh** is created.
- The BGP VPLS mesh is added to the **bgp-vpls-mesh** group through the pseudowire template used to create the BGP VPLS mesh.
- The BGP MH site defined for the access endpoint is associated with the **bgp-vpls-mesh** group; its status from now on will be influenced by the status of the BGP VPLS mesh.

A simple configuration example follows:

```
service>oper-group bgp-vpls-mesh-1 create
service>vpls>bgp>pw-template-binding> oper-group bgp-vpls-mesh-1
service>vpls>site> monitor-group bgp-vpls-mesh-1
```

3.2.13 MAC Flush Message Processing

The previous sections described operating principles of several redundancy mechanisms available in the context of VPLS service. All of them rely on MAC flush message as a tool to propagate topology change in a context of the specified VPLS. This section summarizes basic rules for generation and processing of these messages.

As described on respective sections, the 7450 ESS, 7750 SR, and 7950 XRS support two types of MAC flush message: flush-all-but-mine and flush-mine. The main difference between these messages is the type of action they signal. Flush-all-but-mine messages request clearing of all FDB entries that were learned from all other LDP peers except the originating PE. This type is also defined by RFC 4762 as an LDP MAC address withdrawal with an empty MAC address list.

Flush-all-mine messages request clearing all FDB entries learned from the originating PE. This means that this message has the opposite effect of the flush-all-but-mine message. This type is not included in the RFC 4762 definition and it is implemented using vendor-specific TLV.

The advantages and disadvantages of the individual types should be apparent from examples in the previous section. The description here summarizes actions taken on reception and the conditions under which individual messages are generated.

Upon reception of MAC flush messages (regardless the type), an SR-series PE will take following actions:

- Clears FDB entries of all indicated VPLS services conforming to the definition.
- Propagates the message (preserving the type) to all LDP peers, if propagate-mac-flush flag is enabled at the corresponding VPLS level.

The flush-all-but-mine message is generated under the following conditions:

- The flush-all-but-mine message is received from the LDP peer and the propagate-mac-flush flag is enabled. The message is sent to all LDP peers in the context of the VPLS service it was received.
- TCN message in a context of STP instance is received. The flush-all-but-mine message is sent to all LDP peers connected with spoke and mesh SDPs in a context of VPLS service controlled by the specified STP instance (based on M-VPLS definition). If all LDP peers are in the STP domain, that is the M-VPLS and the uVPLS both have the same topology, the router will not send any flush-all-but-mine message. If the router has uVPLS LDP peers outside the STP domain, the router will send flush-all-but-mine messages to all its uVPLS peers.

Note: The 7750 SR will not send a withdrawal if the M-VPLS does not contain a mesh SDP. A mesh SDP must be configured in the M-VPLS to send withdrawals.

- Flush-all-but-mine message is generated when switchover between spoke-SDPs of the same endpoint occurs. The message is sent to the LDP peer connected through the newly active spoke-SDP.

The flush-mine message is generated under the following conditions:

- The flush-mine message is received from the LDP peer and the propagate-mac-flush flag is enabled. The message is sent to all LDP peers in the context of the VPLS service it was received.
- The flush-mine message is generated when a SAP or SDP transitions from operationally up to an operationally down state and the send-flush-on-failure flag is enabled in the context of the specified VPLS service. The message is sent to all LDP peers connected in the context of the specified VPLS service. The send-flush-on-failure flag is blocked in M-VPLS and is only allowed to be configured in a VPLS service managed by M-VPLS. This is to prevent both messages being sent at the same time.

- The flush-mine message is generated when an MC-LAG SAP or MC-APS SAP transitions from an operationally up state to an operationally down state. The message is sent to all LDP peers connected in the context of the specified VPLS service.
- The flush-mine message is generated when an MC-Ring SAP transitions from operationally up to an operationally down state or when MC-Ring SAP transitions to slave state. The message is sent to all LDP peers connected in the context of the specified VPLS service.

3.2.13.1 Dual Homing to a VPLS Service

Figure 75 **Dual-homed CE Connection to VPLS**

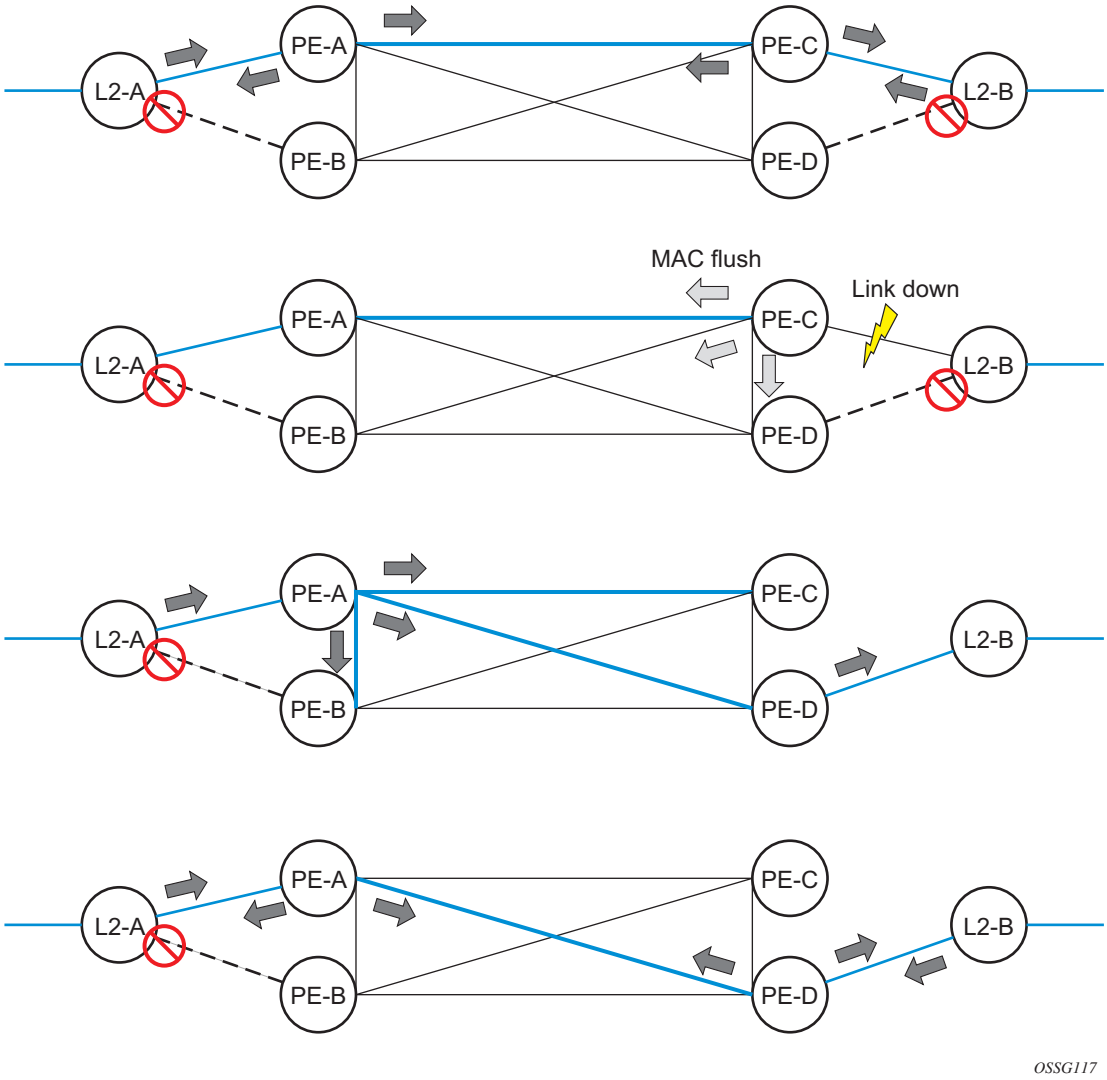

Figure 75 illustrates a dual-homed connection to VPLS service (PE-A, PE-B, PE-C, PE-D) and operation in case of link failure (between PE-C and L2-B). Upon detection of a link failure, PE-C will send MAC-Address-Withdraw messages, which will indicate to all LDP peers that they should flush all MAC addresses learned from PE-C. This will lead to a broadcasting of packets addressing affected hosts and relearning in case an alternative route exists.

The message described here is different than the message described in draft-ietf-l2vpn-vpls-ldp-xx.txt, *Virtual Private LAN Services over MPLS*. The difference is in the interpretation and action performed in the receiving PE. According to the draft definition, upon receipt of a MAC-withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed, This section specifies that all MAC addresses learned from the source are flushed. This message has been implemented as an LDP address message with vendor-specific type, length, and value (TLV), and is called the flush-all-from-me message.

The draft definition message is currently used in management VPLS which, is using RSTP for recovering from failures in Layer 2 topologies. The mechanism described in this document represents an alternative solution.

The advantage of this approach (as compared to RSTP-based methods) is that only MAC-affected addresses are flushed and not the full forwarding database. While this method does not provide a mechanism to secure alternative loop-free topology, the convergence time is dependent on the speed that the specified CE device will open an alternative link (L2-B switch in Figure 75) as well as on the speed that PE routers will flush their FDB.

In addition, this mechanism is effective only if PE and CE are directly connected (no hub or bridge) as it reacts to the physical failure of the link.

3.2.13.2 MC-Ring and VPLS

The use of multi-chassis ring control in a combination with the plain VPLS SAP is supported FDB in individual ring nodes, in case the link (or ring node) failure cannot be cleared on the 7750 SR or 7950 XRS.

This combination is not easily blocked in the CLI. If configured, the combination may be functional but the switchover times will be proportional to MAC aging in individual ring nodes and/or to relearning rate due to downstream traffic.

Redundant plain VPLS access in ring configurations, therefore, exclude corresponding SAPs from the multi-chassis ring operation. Configurations such as M-VPLS can be applied.

3.2.14 ACL Next-Hop for VPLS

The ACL next-hop for VPLS feature enables an ACL that has a forward to a SAP or SDP action specified to be used in a VPLS service to direct traffic with specific match criteria to a SAP or SDP. This allows traffic destined for the same gateway to be split and forwarded differently based on the ACL.

Figure 76 **Application 1 Diagram**

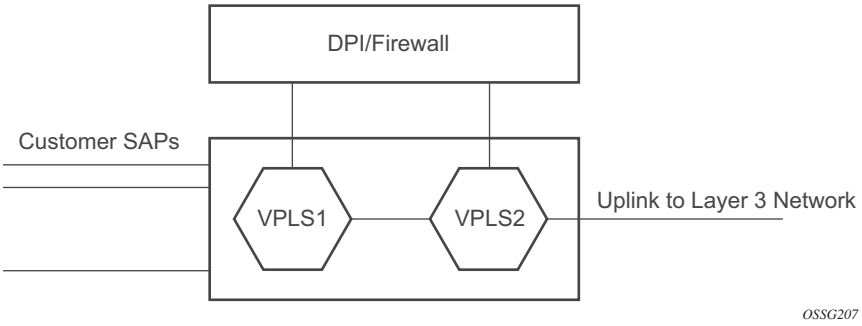

Policy routing is a popular tool used to direct traffic in Layer 3 networks. As Layer 2 VPNs become more popular, especially in network aggregation, policy forwarding is required. Many providers are using methods such as DPI servers, transparent firewalls, or Intrusion Detection/Prevention Systems (IDS/IPS). Since these devices are bandwidth limited, providers want to limit traffic forwarded through them. In the setup shown in Figure 76, a mechanism is required to direct some traffic coming from a SAP to the DPI without learning, and other traffic coming from the same SAP directly to the gateway uplink-based learning.

This feature will allow the provider to create a filter that will forward packets to a specific SAP or SDP. The packets are then forwarded to the destination SAP regardless of learned destination. The SAP can either terminate a Layer 2 firewall, perform deep packet inspection (DPI) directly, or may be configured to be part of a cross-connect bridge into another service. This will be useful when running the DPI remotely using VLLs. If an SDP is used, the provider can terminate it in a remote VPLS or VLL service where the firewall is connected. The filter can be configured under a SAP or SDP in a VPLS service. All packets (unicast, multicast, broadcast, and unknown) can be delivered to the destination SAP/SDP.

The filter may be associated with SAPs/SDPs belonging to a VPLS service only if all actions in the ACL forward to SAPs/SDPs that are within the context of that VPLS. Other services do not support this feature. An ACL that contains this feature is allowed, but the system will drop any packet that matches an entry with this action.

3.2.15 SDP Statistics for VPLS and VLL Services

The simple three-node network in Figure 77 shows two MPLS SDPs and one GRE SDP defined between the nodes. These SDPs connect VPLS1 and VPLS2 instances that are defined in the three nodes. With this feature, the operator will have local CLI-based as well as SNMP-based statistics collection for each VC used in the SDPs. This will allow for traffic management of tunnel usage by the different services and with aggregation of the total tunnel usage.

Figure 77 **SDP Statistics for VPLS and VLL Services**

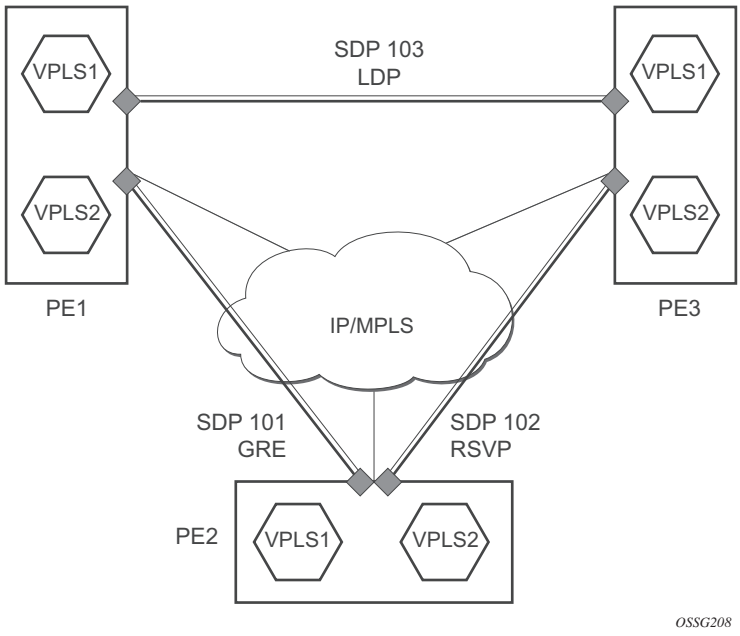

SDP statistics allow providers to bill customers on a per-SDP per-byte basis. This destination-based billing model can be used by providers with a variety of circuit types and have different costs associated with the circuits. An accounting file allows the collection of statistics in bulk.

3.2.16 BGP Auto-Discovery for LDP VPLS

BGP Auto-Discovery (BGP AD) for LDP VPLS is a framework for automatically discovering the endpoints of a Layer 2 VPN, offering an operational model similar to that of an IP VPN. This allows carriers to leverage existing network elements and functions, including but not limited to, route reflectors and BGP policies to control the VPLS topology.

BGP AD complements an already established and well-deployed Layer 2 VPN signaling mechanism target LDP, providing one-touch provisioning for LDP VPLS, where all the related PEs are discovered automatically. The service provider may make use of existing BGP policies to regulate the exchanges between PEs in the same, or in different, autonomous system (AS) domains. The addition of BGP AD procedures does not require carriers to uproot their existing VPLS deployments nor to change the signaling protocol.

3.2.16.1 BGP AD Overview

The BGP protocol establishes neighbor relationships between configured peers. An open message is sent after the completion of the three-way TCP handshake. This open message contains information about the BGP peer sending the message. This message contains Autonomous System Number (ASN), BGP version, timer information, and operational parameters, including capabilities. The capabilities of a peer are exchanged using two numerical values: the Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI). These numbers are allocated by the Internet Assigned Numbers Authority (IANA). BGP AD uses AFI 65 (L2VPN) and SAFI 25 (BGP VPLS). The complete list of allocations are at: *http://www.iana.org/assignments/address-family-numbers* and SAFI *http://www.iana.org/assignments/safi-namespace*.

3.2.16.2 Information Model

Following the establishment of the peer relationship, the discovery process begins as soon as a new VPLS service instance is provisioned on the PE.

Two VPLS identifiers are used to indicate the VPLS membership and the individual VPLS instance:

- VPLS-ID — Membership information, and unique network-wide identifier; same value is assigned for all VPLS switch instances (VSIs) belonging to the same VPLS. VPLS-ID is encodable and carried as a BGP extended community in one of the following formats:
  - A two-octet AS-specific extended community
  - An IPv4 address-specific extended community
- VSI-ID — The unique identifier for each individual VSI, built by concatenating a route distinguisher (RD) with a 4-byte identifier (usually the system IP of the VPLS PE), encoded and carried in the corresponding BGP NLRI.

To advertise this information, BGP AD employs a simplified version of the BGP VPLS NLRI where just the RD and the next 4 bytes are used to identify the VPLS instance. There is no need for Label Block and Label Size fields as T-LDP will signal the service labels later on.

The format of the BGP AD NLRI is very similar to the one used for IP VPN as shown in Figure 78. The system IP may be used for the last 4 bytes of the VSI ID, further simplifying the addressing and the provisioning process.

Figure 78 **BGP AD NLRI versus IP VPN NLRI**

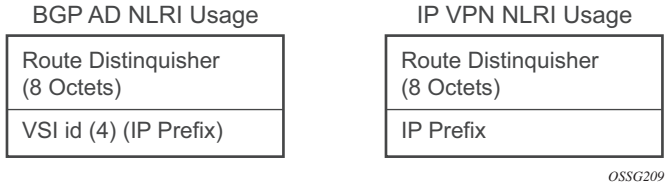

Network Layer Reachability Information (NLRI) is exchanged between BGP peers indicating how to reach prefixes. The NLRI is used in the Layer 2 VPN case to tell PE peers how to reach the VSI, rather than specific prefixes. The advertisement includes the BGP next hop and a route target (RT). The BGP next hop indicates the VSI location and is used in the next step to determine which signaling session is used for pseudowire signaling. The RT, also coded as an extended community, can be used to build a VPLS full mesh or an HVPLS hierarchy through the use of BGP import/ export policies.

BGP is only used to discover VPN endpoints and the corresponding far-end PEs. It is not used to signal the pseudowire labels. This task remains the responsibility of targeted-LDP (T-LDP).

3.2.16.3 FEC Element for T-LDP Signaling

Two LDP FEC elements are defined in RFC 4447, *PW Setup & Maintenance Using LDP*. The original pseudowire-ID FEC element 128 (0x80) employs a 32-bit field to identify the virtual circuit ID and was used extensively in the initial VPWS and VPLS deployments. The simple format is easy to understand, but does not provide the required information model for BGP auto-discovery function. To support BGP AD and other new applications, a new Layer 2 FEC element, the generalized FEC (0x81), is required.

The generalized pseudowire-ID FEC element has been designed for auto discovery applications. It provides a field, the address group identifier (AGI), that is used to signal the membership information from the VPLS-ID. Separate address fields are provided for the source and target address associated with the VPLS endpoints, called the Source Attachment Individual Identifier (SAII) and, Target Attachment Individual Identifier (TAII), respectively. These fields carry the VSI ID values for the two instances that are to be connected through the signaled pseudowire.

The detailed format for FEC 129 is shown in Figure 79.

Figure 79 **Generalized Pseudowire-ID FEC Element**

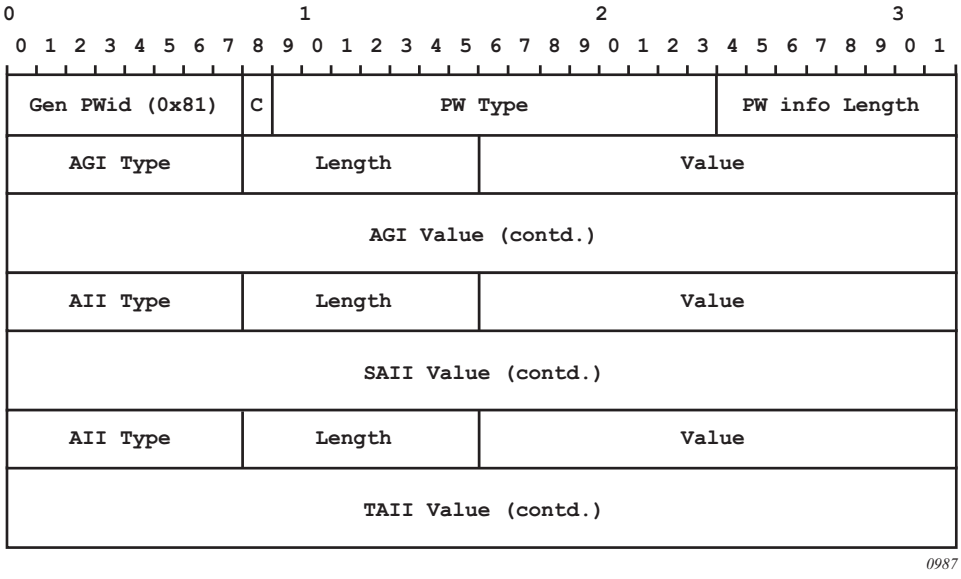

Each of the FEC fields are designed as a sub-TLV equipped with its own type and length, providing support for new applications. To accommodate the BGP AD information model, the following FEC formats are used:

- AGI (type 1) is identical in format and content to the BGP extended community attribute used to carry the VPLS-ID value.
- Source AII (type 1) is a 4-byte value to carry the local VSI-ID (outgoing NLRI minus the RD).
- Target AII (type 1) is a 4-byte value to carry the remote VSI-ID (incoming NLRI minus the RD).

3.2.16.4 BGP-AD and Target LDP (T-LDP) Interaction

BGP is responsible for discovering the location of VSIs that share the same VPLS membership. LDP protocol is responsible for setting up the pseudowire infrastructure between the related VSIs by exchanging service-specific labels between them.

Once the local VPLS information is provisioned in the local PE, the related PEs participating in the same VPLS are identified through BGP AD exchanges. A list of far-end PEs is generated and will trigger the creation, if required, of the necessary T-LDP sessions to these PEs and the exchange of the service-specific VPN labels. The steps for the BGP AD discovery process and LDP session establishment and label exchange are shown in Figure 80.

Figure 80 **BGP-AD and T-LDP Interaction**

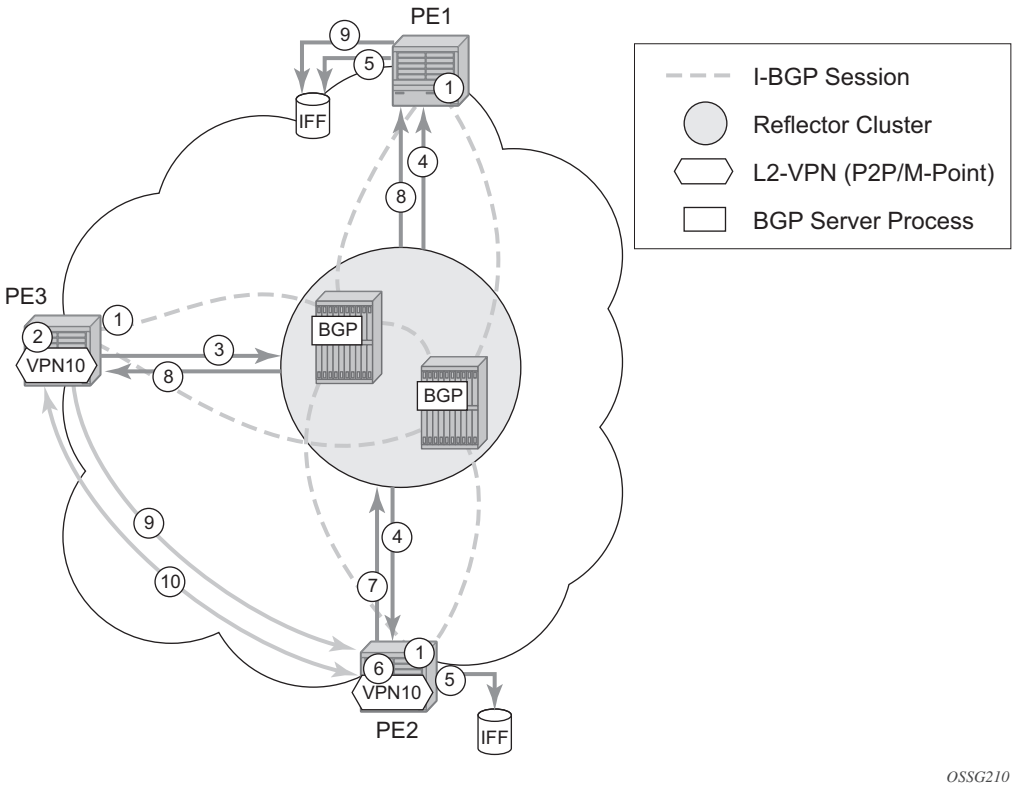

Key:

1. Establish I-BGP connectivity RR.
2. Configure VPN (10) on edge node (PE3).
3. Announce VPN to RR using BGP-AD.

4. Send membership update to each client of the cluster.
5. LDP exchange or inbound FEC filtering (IFF) of non-match or VPLS down.
6. Configure VPN (10) on edge node (PE2).
7. Announce VPN to RR using BGP-AD.
8. Send membership update to each client of the cluster.
9. LDP exchange or inbound FEC filtering (IFF) of non-match or VPLS down.
10. Complete LDP bidirectional pseudowire establishment FEC 129.

3.2.16.5 SDP Usage

Service Access Points (SAP) are linked to transport tunnels using Service Distribution Points (SDP). The service architecture allows services to be abstracted from the transport network.

MPLS transport tunnels are signaled using the Resource Reservation Protocol (RSVP-TE) or by the Label Distribution Protocol (LDP). The capability to automatically create an SDP only exists for LDP-based transport tunnels. Using a manually provisioned SDP is available for both RSVP-TE and LDP transport tunnels. Refer to the appropriate *7450 ESS, 7750 SR, and 7950 XRS MPLS Guide* for more information about MPLS, LDP, and RSVP.

3.2.16.6 Automatic Creation of SDPs

When BGP AD is used for LDP VPLS, and LDP is used as the transport tunnel, there is no requirement to manually create an SDP. The LDP SDP can be automatically instantiated using the information advertised by BGP AD. This simplifies the configuration on the service node.

Enabling LDP on the IP interfaces connecting all nodes between the ingress and the egress builds transport tunnels based on the best IGP path. LDP bindings are automatically built and stored in the hardware. These entries contain an MPLS label pointing to the best next hop along the best path toward the destination.

When two endpoints need to connect and no SDP exists, a new SDP will automatically be constructed. New services added between two endpoints that already have an automatically created SDP will be immediately used. No new SDP will be constructed. The far-end information is learned from the BGP next hop information in the NLRI. When services are withdrawn with a BGP_Unreach_NLRI, the automatically established SDP will remain up while at least one service is connected between those endpoints. An automatically created SDP will be removed and the resources released when the only or last service is removed.

The service provider has the option of associating the auto-discovered SDP with a split horizon group using the **pw-template-binding** option, control the forwarding between pseudowires and to prevent Layer 2 service loops.

An auto-discovered SDP using a **pw-template-binding** without a split horizon group configured will have similar traffic flooding behavior as a spoke-SDP.

3.2.16.7 Manually Provisioned SDP

The carrier is required to manually provision the SDP if they create transport tunnels using RSVP-TE. Operators have the option to choose a manually configured SDP if they use LDP as the tunnel signaling protocol. The functionality is the same regardless of the signaling protocol.

Creating a BGP AD-enabled VPLS service on an ingress node with the manually provisioned SDP option causes the Tunnel Manager to search for an existing SDP that connects to the far-end PE. The far-end IP information is learned from the BGP next hop information in the NLRI. If a single SDP exists to that PE, it is used. If no SDP is established between the two endpoints, the service will remain down until a manually configured SDP becomes active.

When multiple SDPs exist between two endpoints, the tunnel manager will select the appropriate SDP. The algorithm will prefer SDPs with the best (lower) metric. Should there be multiple SDPs with equal metrics, the operational state of the SDPs with the best metric will be considered. If the operational state is the same, the SDP with the higher SDP-ID will be used. If an SDP with a preferred metric is found with an operational state that is not active, the tunnel manager will flag it as ineligible and restart the algorithm.

3.2.16.8 Automatic Instantiation of Pseudowires (SDP Bindings)

The choice of manual or auto-provisioned SDPs has limited impact on the amount of required provisioning. Most of the savings are achieved through the automatic instantiation of the pseudowire infrastructure (SDP bindings). This is achieved for every auto-discovered VSI through the use of the pseudowire template concept. Each VPLS service that uses BGP AD contains the pw-template-binding option defining specific Layer 2 VPN parameters. This command references a pw-template, which defines the pseudowire parameters. The same pw-template may be referenced by multiple VPLS services. As a result, changes to these pseudowire templates have to be treated with caution as they may impact many customers at simultaneously.

The Nokia implementation provides for safe handling of pseudowire templates. Changes to the pseudowire templates are not automatically propagated. Tools are provided to evaluate and distribute the changes. The following command is used to distribute changes to a pw-template at the service level to one or all services that use that template.

PERs-4# tools perform service id 300 eval-pw-template 1 allow-service-impact

If the service ID is omitted, all services will be updated. The type of change made to the pw-template will influence how the service is impacted.

1. Adding or removing a split-horizon-group will cause the router to destroy the original object and recreate using the new value.
2. Changing parameters in the **vc-type** {**ether** | **vlan**} command requires LDP to re-signal the labels.

Both of these changes are service affecting. Other changes will not be service affecting.

3.2.16.9 Mixing Statically Configured and Auto-Discovered Pseudowires in a VPLS

The services implementation allows for manually provisioned and auto-discovered pseudowire (SDP bindings) to coexist in the same VPLS instance (for example, both FEC128 and FEC 129 are supported). This allows for gradual introduction of auto-discovery into an existing VPLS deployment.

As FEC 128 and 129 represent different addressing schemes, it is important to make sure that only one is used at any time between the same two VPLS instances. Otherwise, both pseudowires may become active causing a loop that might adversely impact the correct functioning of the service. It is recommended that FEC128 pseudowire be disabled as soon as the FEC129 addressing scheme is introduced in a portion of the network. Alternatively, RSTP may be used during the migration as a safety mechanism to provide additional protection against operational errors.

3.2.16.10 Resiliency Schemes

The use of BGP AD on the network side, or in the backbone, does not affect the different resiliency schemes Nokia has developed in the access network. This means that both Multi-Chassis Link Aggregation (MC-LAG) and Management-VPLS (M-VPLS) can still be used.

BGP AD may coexist with Hierarchical-VPLS (H-VPLS) resiliency schemes (for example, dual-homed MTU-s devices to different PE-rs nodes) using existing methods (M-VPLS and statically configured Active/Standby pseudowire endpoint).

If provisioned SDPs are used by BGP AD, M-VPLS may be employed to provide loop avoidance. However, it is currently not possible to auto-discover active/standby pseudowires and to instantiate the related endpoint.

3.2.17 BGP VPLS

The Nokia BGP VPLS solution, compliant with RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*, is described in this section.

Figure 81 **BGP VPLS Solution**

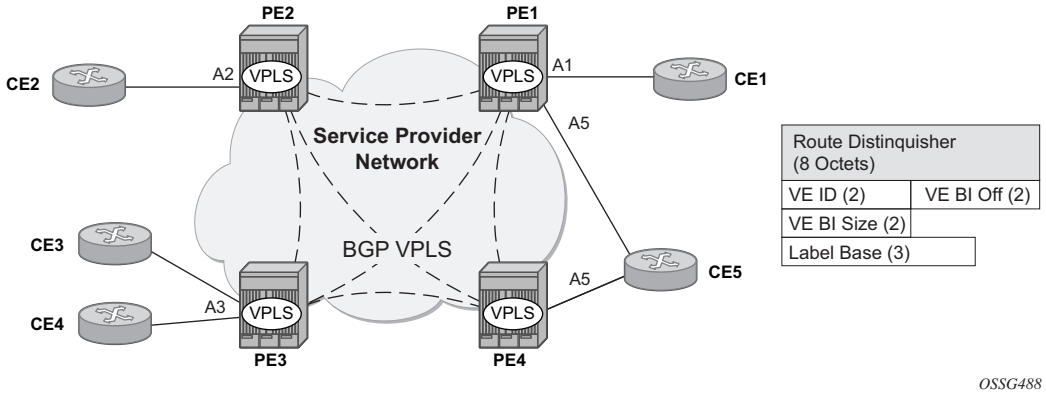

Figure 81 shows the service representation for BGP VPLS mesh. The major BGP VPLS components and the deltas from LDP VPLS with BGP AD are explained below:

- Data plane is identical with the LDP VPLS solution: for example, VPLS instances interconnected by pseudowire mesh. Split horizon groups may be used for loop avoidance between pseudowires.
- Addressing is based on a 2-byte VE-ID assigned to the VPLS instance.
  - BGP-AD for LDP VPLS: 4-byte VSI-ID (system IP) identifies the VPLS instance.
- The target VPLS instance is identified by the Route Target (RT) contained in the MP-BGP advertisement (extended community attribute).
  - BGP-AD: a new MP-BGP extended community is used to identify the VPLS. RT is used for topology control.
- Auto-discovery is MP-BGP based; the same AFI, SAFI is used as for LDP VPLS BGP-AD.
  - The BGP VPLS updates are distinguished from the BGP-AD ones based on the value of the NLRI prefix length: 17 bytes for BGP VPLS, 12 bytes for BGP-AD.
  - BGP-AD NLRI is shorter since there is no need to carry pseudowire label information as T-LDP does the pseudowire signaling for LDP VPLS.
- Pseudowire label signaling is MP-BGP based. Therefore, the BGP NLRI content also includes label related information; for example, block offset, block size, and label base.
  - LDP VPLS: target LDP (T-LDP) is used for signaling the pseudowire service label.
  - The Layer 2 extended community proposed in RFC 4761 is used to signal pseudowire characteristics; for example, VPLS status, control word, and sequencing.

3.2.17.1 Pseudowire Signaling Details

The pseudowire is set up using the following NLRI fields:

- VE Block offset (VBO): used to define for each VE-ID set the NLRI is targeted:
  - VBO = n*VBS+1; for VBS = 8, this results in 1, 9, 17, 25, …
  - Targeted Remote VE-IDs are from VBO to (VBO + VBS – 1)
- VE Block size (VBS): defines how many contiguous pseudowire labels are reserved, starting with the Label Base.
  - Nokia implementation uses always a value of eight (8).
- Label Base (LB): local allocated label base.
  - The next eight labels allocated for remote PEs.

This BGP update is telling the other PEs that accept the RT: reach me (VE-ID = x), use a pseudowire label of LB + VE-ID – VBO using the BGP NLRI for which VBO =< local VE-ID < VBO + VBS.

Following is an example of how this algorithm works assuming PE1 has VE-ID 7 configured:

- PE1 founds a Label Block of eight consecutive labels available, starting with LB = 1000
- PE1 starts sending a BGP update with pseudowire information of (VBO = 1, VBS = 8, LB = 1000) in the NLRI.
- This pseudowire information will be accepted by all participating PEs with VE-IDs from 1 to 8.
- Each of the receiving PEs will use the pseudowire label = LB + VE-ID – VBO to send traffic back to the originator PE. For example, VE-ID 2 will use pseudowire label 1001.

Assuming that VE-ID = 10 is configured in another PE4, the following procedure applies:

- PE4 sends a BGP update with the new VE-ID in the network that will be received by all the other participating PEs, including PE1.
- PE1 upon reception will generate another label block of 8 labels for the VBO = 9. For example, the initial PE will create new pseudowire signaling information of VBO = 9, VBS = 8, LB = 3000, and insert it in a new NLRI and BGP update that is sent in the network.
- This new NLRI will be used by the VE-IDs from 9 to 16 to establish pseudowires back to the originator PE1. For example, PE4 with VE-ID 10 will use pseudowire label 3001 to send VPLS traffic back to PE1.

- The PEs owning the set of VE-IDs from 1 to 8 will ignore this NLRI.

In addition to the pseudowire label information, the **Layer2 Info Extended Community** attribute must be included in the BGP update for BGP VPLS to signal the attributes of all the pseudowires that converge toward the originator VPLS PE.

The format is as follows:

```
+-----------------------------------+
| Extended community type (2 octets) |
+-----------------------------------+
|  Encaps Type (1 octet)            |
+-----------------------------------+
|  Control Flags (1 octet)          |
+-----------------------------------+
|  Layer-2 MTU (2 octet)            |
+-----------------------------------+
|  Reserved (2 octets)              |
+-----------------------------------+
```

The meaning of the fields are as follows:

- Extended community type – the value allocated by IANA for this attribute is 0x800A
- Encaps Type – Encapsulation type; identifies the type of pseudowire encapsulation. The only value used by BGP VPLS is 19 (13 in HEX). This value identifies the encapsulation to be used for pseudowire instantiated through BGP signaling which is the same as the one used for Ethernet pseudowire type in regular VPLS. There is no support for an equivalent Ethernet VLAN pseudowire in BGP VPLS in BGP signaling.
- Control Flags – control information regarding the pseudowires, as follows
- Layer 2 MTU – the Maximum Transmission Unit to be used on the pseudowires
- Reserved – this field is reserved and must be set to zero and ignored on reception except where it is used for VPLS preference.

  For inter-AS, the preference information must be propagated between autonomous systems. Consequently, as the VPLS preference in a BGP-VPLS or BGP multi-homing update extended community is zero, the local preference is copied by the egress ASBR into the VPLS preference field before sending the update to the eBGP peer. The adjacent ingress ASBR then copies the received VPLS preference into the local preference to prevent the update being considered malformed.

Figure 82 shows the detailed format for the control flags bit vector.

Figure 82 **Control Flag Bit Vector Format**

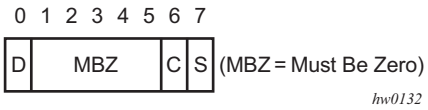

The following bits in the control flags are defined as follows:

- S – sequenced delivery of frames must or must not be used when sending VPLS packets to this PE, depending on whether S is 1 or 0, respectively.
- C – a Control word must or must not be present when sending VPLS packets to this PE, depending on whether C is 1 or 0, respectively. By default, Nokia implementation uses value 0.
- MBZ – Must Be Zero bits, set to zero when sending and ignored when receiving
- D – indicates the status of the whole VPLS instance (VSI); D = 0 if Admin and Operational status are up, D = 1 otherwise

Here are the events that set the D-bit to 1 to indicate VSI down status in BGP update message sent out from a PE:

- Local VSI is shutdown administratively using **config service vpls shutdown** command.
- All the related endpoints (SAPs or LDP pseudowires) are down.
- There are no related endpoints (SAPs or LDP pseudowires) configured yet in the VSI.
  - The intent is to save the core bandwidth by not establishing the BGP pseudowires to an empty VSI.
- Upon reception of a BGP update message with D-bit set to 1, all the receiving VPLS PEs must mark related pseudowires as down.

The following events do not set the D-bit to 1:

- The local VSI is delete; a BGP update with unreach-NLRI is sent out. Upon reception, all remote VPLS PEs must remove the related pseudowires and BGP routes.
- If the local SDP goes down, only the BGP pseudowires mapped to that SDP go down. There is no BGP update sent.

3.2.17.2 Supported VPLS Features

BGP VPLS added support for a new type of pseudowire signaling based on MP-BGP, based on the existing VPLS instance; therefore, it inherits all the existing Ethernet switching functions. Following are some of the most important VPLS features ported to BGP VPLS:

- VPLS data plane features: for example, FDB management, SAPs, LAG access, and BUM rate limiting
- MPLS tunneling: LDP, LDP over RSVP-TE, RSVP-TE, GRE, and MP-BGP based on RFC3107 (Inter-AS option C solution)

Note: Pre-provisioned SDPs must be configured when GRE- or RSVP-signaled transport tunnels are used.

- HVPLS topologies, hub and spoke traffic distribution
- Coexists with LDP VPLS (with or without BGP-AD) in the same VPLS instance.
  - LDP and BGP-signaling should operate in disjoint domains to simplify loop avoidance
- Coexists with BGP-based multi-homing.
- BGP VPLS is supported as the control plane for B-VPLS.
- Supports IGMP/PIM snooping for IPv4
- Support for High Availability is provided
- Ethernet Service OAM toolset is supported: IEEE 802.1ag, Y.1731.
  - Not supported OAM features: CPE Ping, MAC trace/ping/populate/purge
- Support for RSVP and LSP P2MP LSP for VPLS/B-VPLS BUM

3.2.18 VCCV BFD Support for VPLS Services

The SR OS supports RFC 5885, which specifies a method for carrying BFD in a pseudowire associated channel. For general information about VCCV BFD, limitations, and configuring, see the VLL Services chapter.

VCCV BFD is supported on the following VPLS services:

- T-LDP spoke-SDP termination on VPLS (including I-VPLS, B-VPLS, and R-VPLS)
- H-VPLS spoke-SDP

- BGP VPLS
- VPLS with BGP auto-discovery

To configure VCCV BFD for H-VPLS (where the pseudowire template does not apply), configure the BFD template using the **config**>**service**>**vpls**>**spoke-sdp**>**bfd-template** *name* command and then enable it using the **config**>**service**>**vpls**>**spoke-sdp**>**bfd-enable** command.

For BGP VPLS, a BFD template is referenced from the pseudowire template binding context. To configure VCCV BFD for BGP VPLS, use the **config**>**service**>**vpls**>**bgp**>**pw-template-binding**>**bfd-template** *name* command and enable it using the **config**>**service**>**vpls**>**bgp**>**pw-template-binding**>**bfd-enable** command.

For BGP-AD VPLS, a BFD template is referenced from the pseudowire template context. To configure VCCV BFD for BGP-AD, use the command **config**>**service**>**vpls**>**bgp-ad**>**pw-template-binding**>**bfd-template** *name* and enable it using the command **config**>**service**>**vpls**>**bgp-ad**>**pw-template-binding**>**bfd-enable**.

3.2.19 BGP Multi-Homing for VPLS

This section describes BGP-based procedures for electing a designated forwarder among the set of PEs that are multi-homed to a customer site. Only the local PEs are actively participating in the selection algorithm. The PEs remote from the dual-homed CE are not required to participate in the designated forwarding election for a remote dual-homed CE.

The main components of the BGP-based multi-homing solution for VPLS are:

- Provisioning model
- MP-BGP procedures
- Designated Forwarder Election
- Blackhole avoidance – indicating the designated forwarder change toward the core PEs and access PEs or CEs
- The interaction with pseudowire signaling (BGP/LDP)

Figure 83 **BGP Multi-Homing for VPLS**

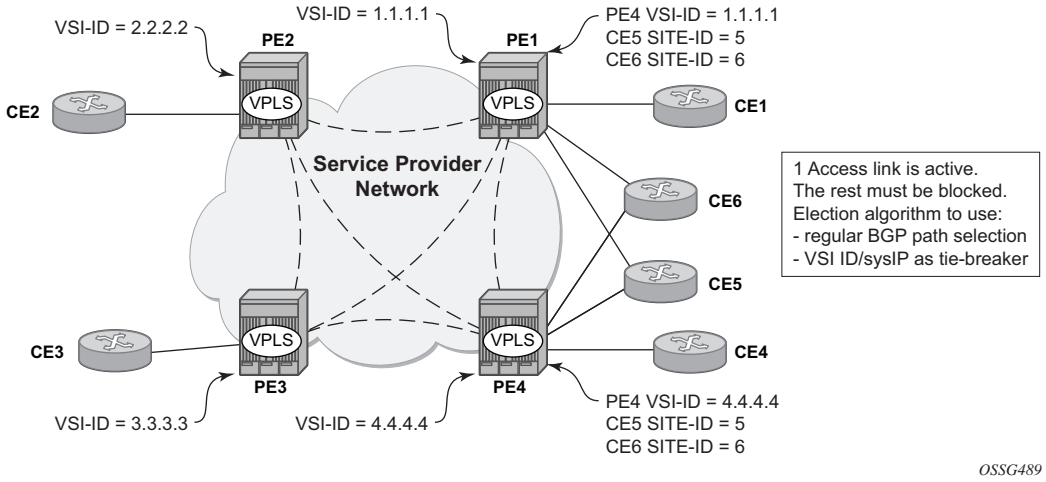

Figure 83 shows the VPLS using BGP multi-homing for the case of multi-homed CEs. Although the figure shows the case of a pseudowire infrastructure signaled with LDP for an LDP VPLS using BGP-AD for discovery, the procedures are identical for BGP VPLS or for a mix of BGP and LDP signaled pseudowires.

3.2.19.1 Information Model and Required Extensions to L2VPN NLRI

VPLS Multi-homing using BGP-MP expands on the BGP AD and BGP VPLS provisioning model. The addressing for the multi-homed site is still independent from the addressing for the base VSI (VSI-ID or, respectively, VE-ID). Every multi-homed CE is represented in the VPLS context through a site-ID, which is the same on the local PEs. The site-ID is unique within the scope of a VPLS. It serves to differentiate between the multi-homed CEs connected to the same VPLS Instance (VSI). For example, in Figure 84, CE5 will be assigned the same site-ID on both PE1 and PE4. For the same VPLS instance, different site-IDs are assigned for multi-homed CE5 and CE6; for example, site-ID 5 is assigned for CE5 and site-ID 6 is assigned for CE6. The single-homed CEs (CE1, 2, 3, and 4) do not require allocation of a multi-homed site-ID. They are associated with the addressing for the base VSI, either VSI-ID or VE-ID.

The new information model required changes to the BGP usage of the NLRI for VPLS. The extended MH NLRI for Multi-Homed VPLS is compared with the BGP AD and BGP VPLS NLRIs in Figure 84.

Figure 84 **BGP MH-NLRI for VPLS Multi-Homing**

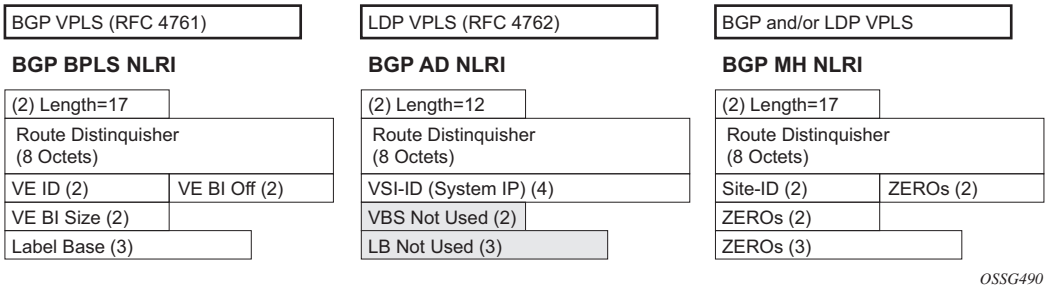

The BGP VPLS NLRI described in RFC 4761 is used to carry a 2-byte site-ID that identifies the MH site. The last seven bytes of the BGP VPLS NLRI used to instantiate the pseudowire are not used for BGP-MH and are zeroed out. This NLRI format translates into the following processing path in the receiving VPLS PE:

- BGP VPLS PE: no label information means there is no need to set up a BGP pseudowire.
- BGP AD for LDP VPLS: length =17 indicates a BGP VPLS NLRI that does not require any pseudowire LDP signaling.

The processing procedures described in this section start from the above identification of the BGP Update as not destined for pseudowire signaling.

The RD ensures the NLRIs associated with a certain site-ID on different PEs are seen as different by any of the intermediate BGP nodes (RRs) on the path between the multi-homed PEs. In other words, different RDs must be used on the MH PEs every time an RR or an ASBR is involved to guarantee the MH NLRIs reach the PEs involved in VPLS MH.

The L2-Info extended community from RFC 4761 is used in the BGP update for MH NLRI to initiate a MAC flush for blackhole avoidance to indicate the operational and admin status for the MH Site or the DF election status.

After the pseudowire infrastructure between VSIs is built using either RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, or RFC 4761 procedures or a mix of pseudowire Signaling procedure, on activation of a multi- homed site, an election algorithm must be run on the local and remote PEs to determine which site will be the designated forwarder (DF). The end result is that all the related MH sites in a VPLS will be placed in standby except for the site selected as DF. Nokia BGP-based multi-homing solution uses the DF election procedure described in the IETF working group document *draft-ietf-bess-vpls-multihoming-01*. The implementation allows the use of BGP Local Preference and the received VPLS preference but does not support setting the VPLS preference to a non-zero value.

3.2.19.2 Supported Services and Multi-Homing Objects

This feature is supported for the following services:

- LDP VPLS with or without BGP-AD
- BGP VPLS (BGP multi-homing for inter-AS BGP-VPLS services is not supported)
- mix of the above
- PBB B-VPLS on BCB
- PBB I-VPLS (refer to IEEE 802.1ah Provider Backbone Bridging for more information)

The following access objects can be associated with MH SITE:

- SAPs
- SDP bindings (pseudowire object), both mesh-SDP and spoke-SDP
- Split Horizon Group
  - Under the SHG we can associate either one or multiple of the following objects: SAPs, pseudowires (BGP VPLS, BGP-AD, provisioned and LDP signaled spoke-SDP and mesh-SDP)

3.2.19.3 Blackhole Avoidance

Blackholing refers to the forwarding of frames to a PE that is no longer carrying the designated forwarder. This could happen for traffic from:

- Core PE participating in the main VPLS
- Customer Edge devices (CEs)
- Access PEs - pseudowires between them and the MH PEs are associated with MH Sites

Changes in DF election results or MH site status must be detected by all of the above network elements to provide for Blackhole Avoidance.

3.2.19.3.1 MAC Flush to the Core PEs

Assuming there is a transition of the existing DF to non-DF status. The PE that owns the MH site experiencing this transition will generate a MAC flush-all-from-me (negative MAC flush) toward the related core PEs. Upon reception, the remote PEs will flush all the MACs learned from the MH PE.

MAC flush-all-from-me indication is sent using the following core mechanisms:

- For LDP VPLS running between core PEs, existing LDP MAC flush will be used.
- For pseudowire signaled with BGP VPLS, MAC flush will be provided implicitly using the L2-Info Extended community to indicate a transition of the active MH-site: for example the attached objects going down or more generically, the entire site going from Designated Forwarder (DF) to non-DF.
- Double flushing will not happen as it is expected that between any pair of PEs it will exist only one type of pseudowires – either BGP or LDP pseudowire but not both.

3.2.19.3.2 Indicating non-DF status toward the access PE or CE

For the CEs or access PEs support is provided for indicating the blocking of the MH site using the following procedures:

- For MH Access PE running LDP pseudowires the LDP standby-status is sent to all LDP pseudowires.
- For MH CEs site deactivation is linked to a CCM failure on a SAP that has a down MEP configured.

3.2.19.4 BGP Multi-Homing for VPLS Inter-Domain Resiliency

BGP MH for VPLS can be used to provide resiliency between different VPLS domains. An example of a multi-homing topology is shown in Figure 85.

Figure 85 **BGP MH Used in an HVPLS Topology**

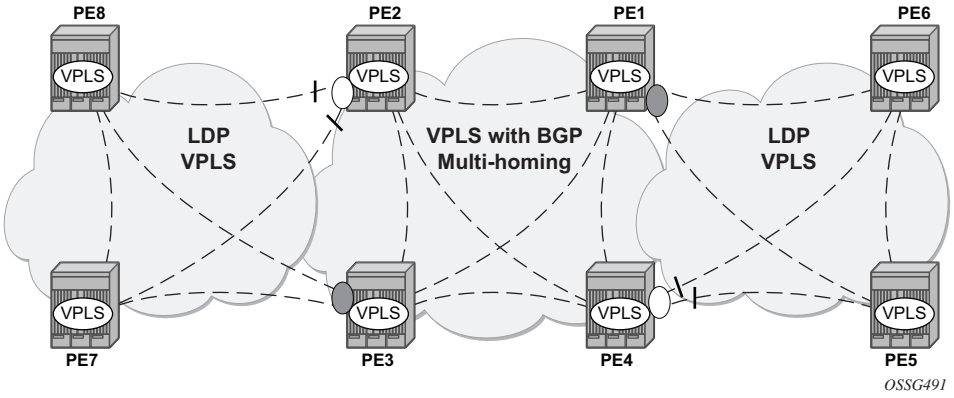

LDP VPLS domains are interconnected using a core VPLS domain either BGP VPLS or LDP VPLS. The gateway PEs, for example PE2 and PE3, are running BGP multi-homing where one MH site is assigned to each of the pseudowires connecting the access PE, PE7, and PE8 in this example.

Alternatively, one may choose to associate the MH site to multiple access pseudowires using an access SHG. The **config**>**service**>**vpls**>**site**>**failed-threshold** command can be used to indicates the number of pseudowire failures that are required for the MH site to be declared down.

3.2.20 Multicast-Aware VPLS

VPLS is a Layer 2 service; hence multicast and broadcast frames are normally flooded in a VPLS. Broadcast frames are targeted to all receivers. However, for IP multicast, normally for a multicast group, only some receivers in the VPLS are interested. Flooding to all sites can cause wasted network bandwidth and unnecessary replication on the ingress PE router.

To avoid this condition, VPLS is IP multicast-aware; therefore, it forwards IP multicast traffic based on multicast states to the object on which the IP multicast traffic is requested. This is achieved by enabling the following related IP multicast protocol snooping:

- IGMP snooping
- MLD snooping
- PIM snooping

3.2.20.1 IGMP Snooping for VPLS

When IGMP snooping is enabled in a VPLS service, IGMP messages received on SAPs and SDPs are snooped in order to determine the scope of the flooding for a specified stream or (S,G). IGMP snooping operates in a proxy mode, where the system summarizes upstream IGMP reports and responds to downstream queries. See “IGMP Snooping” in the *7450 ESS and 7750 SR Triple Play Guide* for a description of IGMP snooping.

Streams are sent to all SAPs/SDPs on which there is a multicast router (either discovered dynamically from received query messages or configured statically using the **mrouter-port** command) and on which an active join for that stream has been received. The mrouter port configuration adds a (*,*) entry into the MFIB, which causes all groups (and IGMP messages) to be sent out of the respective object and causes IGMP messages received on that object to be discarded.

Directly connected multicast sources are supported when IGMP snooping is enabled.

IGMP snooping is enabled at the service level and is not supported in the following services:

- B-VPLS
- Routed I-VPLS
- BGP EVPN in routed VPLS services
- PBB B-VPLS services

IGMP snooping is not supported on a router configured with **enable-inter-as-vpn** or **enable-rr-vpn-forwarding**.

IGMP snooping is not supported under the following forms of default SAP:

- *
- *.null
- *.*

3.2.20.2 MLD Snooping for VPLS

MLD snooping is an IPv6 version of IGMP snooping. The guidelines and procedures are similar to IGMP snooping as described above. However, MLD snooping uses MAC-based forwarding. See MAC-Based IPv6 Multicast Forwarding for more information. Directly connected multicast sources are supported when MLD snooping is enabled.

MLD snooping is enabled at the service level and is not supported in the following services:

- B-VPLS
- Routed I-VPLS
- PBB-EVPN services

MLD snooping is not supported under the following forms of default SAP:

- *
- *.null
- *.*

3.2.20.3 PIM Snooping for VPLS

PIM snooping for VPLS allows a VPLS PE router to build multicast states by snooping PIM protocol packets that are sent over the VPLS. The VPLS PE then forwards multicast traffic based on the multicast states. When all receivers in a VPLS are IP multicast routers running PIM, multicast forwarding in the VPLS is efficient when PIM snooping for VPLS is enabled.

Because of PIM join/prune suppression, in order to make PIM snooping operate over VPLS pseudowires, two options are available: plain PIM snooping and PIM proxy. PIM proxy is the default behavior when PIM snooping is enabled for a VPLS.

PIM snooping is supported for both IPv4 and IPv6 multicast by default and can be configured to use SG based forwarding (see IPv6 Multicast Forwarding for more information).

Directly connected multicast sources are supported when PIM snooping is enabled.

The following restrictions apply to PIM snooping:

- PIM snooping for IPv4 and IPv6 is not supported:
  - in the following services:
    - PBB B-VPLS
    - Routed VPLS (including I-VPLS and BGP EVPN)
    - PBB-EVPN B-VPLS
    - EVPN-VXLAN
  - on a router configured with **enable-inter-as-vpn** or **enable-rr-vpn-forwarding**

- under the following forms of default SAP:
  - *
  - *.null
  - *.*
- with connected SR OSs configured with **improved-assert**
- with subscriber management in the VPLS service
- as a mechanism to drive MCAC

• PIM snooping for IPv6 is not supported:
- in the following services:
  - PBB I-VPLS
  - BGP-VPLS
  - BGP EVPN (including PBB-EVPN)
  - VPLS E-Tree
  - Management VPLS
- with the configuration of MLD snooping

3.2.20.3.1 Plain PIM Snooping

In a plain PIM snooping configuration, VPLS PE routers only snoop; PIM messages are generated on their own. Join/prune suppression must be disabled on CE routers.

When plain PIM snooping is configured, if a VPLS PE router detects a condition where join/prune suppression is not disabled on one or more CE routers, the PE router will put PIM snooping into the PIM proxy state. A trap is generated that reports the condition to the operator and is logged to the syslog. If the condition changes, for example, join/prune suppression is disabled on CE routers, the PE reverts to the plain PIM snooping state. A trap is generated and is logged to the syslog.

3.2.20.3.2 PIM Proxy

For PIM proxy configurations, VPLS PE routers perform the following:

- Snoop hellos and flood hellos in the fast data path.
- Consume join/prune messages from CE routers.
- Generate join/prune messages upstream using the IP address of one of the downstream CE routers.

- Run an upstream PIM state machine to determine whether a join/prune message should be sent upstream.

Join/prune suppression is not required to be disabled on CE routers, but it requires all PEs in the VPLS to have PIM proxy enabled. Otherwise, CEs behind the PEs that do not have PIM proxy enabled may not be able to get multicast traffic that they are interested in if they have join/prune suppression enabled.

When PIM proxy is enabled, if a VPLS PE router detects a condition where join/prune suppression is disabled on all CE routers, the PE router put PIM proxy into a plain PIM snooping state to improve efficiency. A trap is generated to report the scenario to the operator and is logged to the syslog. If the condition changes, for example, join/prune suppression is enabled on a CE router, PIM proxy is placed back into the operational state. Again, a trap is generated to report the condition to the operator and is logged to the syslog.

3.2.20.4 IPv6 Multicast Forwarding

When MLD snooping or PIM snooping for IPv6 is enabled, the forwarding of IPv6 multicast traffic is MAC-based; see MAC-Based IPv6 Multicast Forwarding for more information.

The operation with PIM snooping for IPv6 can be changed to SG-based forwarding; see SG-Based IPv6 Multicast Forwarding for more information.

The following command configures the IPv6 multicast forwarding mode with the default being **mac-based**:

```
configure service vpls mcast-ipv6-snooping-scope {sg-
 based | mac-based}
```

The forwarding mode can only be changed when PIM snooping for IPv6 is disabled.

3.2.20.4.1 MAC-Based IPv6 Multicast Forwarding

This section describes IPv6 multicast address to MAC address mapping and IPv6 multicast forwarding entries.

For IPv6 multicast address to MAC address mapping, Ethernet MAC addresses in the range of 33-33-00-00-00-00 to 33-33-FF-FF-FF-FF are reserved for IPv6 multicast. To map an IPv6 multicast address to a MAC-layer multicast address, the low-order 32 bits of the IPv6 multicast address are mapped directly to the low-order 32 bits in the MAC-layer multicast address.

For IPv6 multicast forwarding entries, IPv6 multicast snooping forwarding entries are based on MAC addresses, while native IPv6 multicast forwarding entries are based on IPv6 addresses. When both MLD snooping or PIM snooping for IPv6 and native IPv6 multicast are enabled on the same device, both types of forwarding entries are supported on the same forward plane, although they are used for different services.

The output below shows a service with PIM snooping for IPv6 that has received joins for two multicast groups from different sources. As the forwarding mode is MAC-based, there is a single MFIB entry created to forward these two groups.

```
*A:PE# show service id 1 pim-snooping group ipv6
===============================================================================
PIM Snooping Groups ipv6
===============================================================================
Group Address           Source Address          Type     Incoming          Num
                                                         Intf              Oifs
-------------------------------------------------------------------------------
ff0e:db8:1000::1        2001:db8:1000::1        (S,G)    SAP:1/1/1         2
ff0e:db8:1001::1        2001:db8:1001::1        (S,G)    SAP:1/1/1         2
-------------------------------------------------------------------------------
Groups : 2
===============================================================================
*A:PE#

*A:PE# show service id 1 all | match "Mcast IPv6 scope"
Mcast IPv6 scope  : mac-based
*A:PE#

*A:PE# show service id 1 mfib
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                            Blk
-------------------------------------------------------------------------------
*               33:33:00:00:00:01     sap:1/1/1                    Local    Fwd
                                      sap:1/1/2                    Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
*A:PE#
```

3.2.20.4.2 SG-Based IPv6 Multicast Forwarding

When PIM snooping for IPv6 is configured, SG-based forwarding can be enabled, which causes the IPv6 multicast forwarding to be based on both the source (if specified) and destination IPv6 address in the received join.

Enabling SG-based forwarding increases the MFIB usage if the source IPv6 address or higher 96 bits of the destination IPv6 address varies in the received joins compared to using MAC-based forwarding.

The output below shows a service with PIM snooping for IPv6 that has received joins for two multicast groups from different sources. As the forwarding mode is SG-based, there are two MFIB entries, one for each of the two groups.

```
*A:PE# show service id 1 pim-snooping group ipv6
===============================================================================
PIM Snooping Groups ipv6
===============================================================================
Group Address           Source Address          Type     Incoming          Num
                                                         Intf              Oifs
-------------------------------------------------------------------------------
ff0e:db8:1000::1        2001:db8:1000::1        (S,G)    SAP:1/1/1         2
ff0e:db8:1001::1        2001:db8:1001::1        (S,G)    SAP:1/1/1         2
-------------------------------------------------------------------------------
Groups : 2
===============================================================================
*A:PE#

*A:PE# show service id 1 all | match "Mcast IPv6 scope"
Mcast IPv6 scope  : sg-based
*A:PE#

*A:PE# show service id 1 mfib
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                            Blk
-------------------------------------------------------------------------------
2001:db8:1000:* ff0e:db8:1000::1      sap:1/1/1                    Local    Fwd
                                      sap:1/1/2                    Local    Fwd
2001:db8:1001:* ff0e:db8:1001::1      sap:1/1/1                    Local    Fwd
                                      sap:1/1/2                    Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
*A:PE#
```

SG-based IPv6 multicast forwarding is supported when both plain PIM snooping and PIM proxy are supported.

SG-based forwarding is only supported on FP3 line cards. It is supported in all services in which PIM snooping for IPv6 is supported, with the same restrictions.

It is not supported in the following services:

- PBB B-VPLS
- PBB I-VPLS
- Routed-VPLS (including with I-VPLS and BGP-EVPN)
- BGP-EVPN (including PBB-EVPN)
- VPLS E-Tree
- Management VPLS

In any specific service, SG-based forwarding and MLD snooping are mutually exclusive. Consequently, MLD snooping uses MAC-based forwarding.

It is not supported in services with:

- subscriber management
- multicast VLAN Registration
- video interface

It is not supported on connected SR OS routers configured with **improved-assert**.

It is not supported with the following forms of default SAP:

- *
- *.null
- *.*

3.2.20.5 PIM and IGMP/MLD Snooping Interaction

When both PIM snooping for IPv4 and IGMP snooping are enabled in the same VPLS service, multicast traffic is forwarded based on the combined multicast forwarding table.

There is no interaction between PIM snooping for IPv6 and PIM snooping for IPv4/ IGMP snooping when all are enabled within the same VPLS service. The configurations of PIM snooping for IPv6 and MLD snooping are mutually exclusive.

When PIM snooping is enabled within a VPLS service, all IP multicast traffic and flooded PIM messages (these include all PIM snooped messages when not in PIM proxy mode and PIM hellos when in PIM proxy mode) will be sent to any SAP or SDP binding configured with an IGMP-snooping mrouter port. This will occur even without IGMP-snooping enabled, but is not supported in a BGP-VPLS or M-VPLS service.

3.2.20.6 Multi-Chassis Synchronization for Layer 2 Snooping States

To achieve a faster failover in scenarios with redundant active/standby routers performing Layer 2 multicast snooping, it is possible to synchronize the snooping state from the active router to the standby router, so that if a failure occurs the standby router has the Layer 2 multicast snooped states and is able to forward the multicast traffic immediately. Without this capability, there would be a longer delay in re-establishing the multicast traffic path due to having to wait for the Layer 2 states to be snooped.

Multi-chassis synchronization (MCS) is enabled per peer router and uses a **sync-tag**, which is configured on the objects requiring synchronization on both of the routers. This allows MCS to map the state of a set of objects on one router to a set of objects on the other router. Specifically, objects relating to a **sync-tag** on one router are backed up by, or are backing up, the objects using the same **sync-tag** on the other router (the state is synchronized from the active object on one router to its backup objects on the standby router).

The object type must be the same on both routers; otherwise, a mismatch error is reported. The same **sync-tag** value can be reused for multiple peer/object combinations, where each combination represents a different set of synchronized objects; however, a **sync-tag** cannot be configured on the same object to more than one peer.

The **sync-tag** is configured per port and can relate to a specific set of dot1q or QinQ VLANs on that port, as follows:

CLI Syntax:
```
configure
     redundancy
          multi-chassis
               peer ip-address [create]
                    sync
                         port port-id [sync-tag sync-tag]
                            [create]
                              range encap-range sync-tag
                                 sync-tag
```

In order for IGMP snooping and PIM snooping for IPv4 to work correctly with MCS on QinQ ports using x.* SAPs, one of the following must be true:

- MCS is configured with a **sync-tag** for the entire port
- The IGMP snooping SAP and the MCS **sync-tag** must be provisioned with the same Q-tag values when using the range parameter

3.2.20.6.1 IGMP Snooping Synchronization

MCS for IGMP snooping synchronizes the join/prune state information from IGMP messages received on the related port/VLANs corresponding to their associated **sync-tag**. It is enabled as follows:

CLI Syntax:
```
configure
    redundancy
        multi-chassis
            peer ip-address [create]
                sync
                    igmp-snooping
```

IGMP snooping synchronization is supported wherever IGMP snooping is supported (except in EVPN for VXLAN). See IGMP Snooping for VPLS for more information. IGMP snooping synchronization is also only supported for the following active/standby redundancy mechanisms:

- multi-chassis LAG
- multi-chassis ring
- Single-Active Multihoming (EVPN-MPLS and PBB-EVPN I-VPLS)

Configuring an mrouter port under an object that has the synchronizing of IGMP snooping states enabled is not recommended. The mrouter port configuration adds a (*,*) entry into the MFIB, which causes all groups (and IGMP messages) to be sent out of the respective object. In addition, the **mrouter port** command causes all IGMP messages on that object to be discarded. However, the (*,*) entry is not synchronized by MCS. Consequently, the mrouter port could cause the two MCS peers to be forwarding different sets of multicast streams out of the related object when each is active.

3.2.20.6.2 MLD Snooping Synchronization

MCS for MLD snooping is not supported. The command is not blocked for backward-compatibility reasons, but has no effect on the system if configured.

3.2.20.6.3 PIM Snooping for IPv4 Synchronization

MCS for PIM snooping for IPv4 synchronizes the neighbor information from PIM hellos and join/prune state information from PIM for IPv4 messages received on the related SAPs and spoke-SDPs corresponding to the **sync-tag** associated with the related ports and SDPs, respectively. Use the following CLI syntax to enable MCS for PIM snooping for IPv4 synchronization:

CLI Syntax:

```
configure
     redundancy
          multi-chassis
               peer ip-address [create]
                    sync
                         pim-snooping [saps] [spoke-sdps]
```

Any PIM hello state information received over the MCS connection from the peer router takes precedence over locally snooped hello information. This ensures that any PIM hello messages received on the active router that are then flooded, for example through the network backbone, and received over a local SAP or SDP on the standby router are not inadvertently used in the standby router's VPLS service.

The synchronization of PIM snooping state is only supported for manually-configured spoke-SDPs. It is not supported for spoke-SDPs configured within an endpoint.

When synchronizing the PIM state between two spoke-SDPs, if both spoke-SDPs go down, the PIM state is maintained on both until one becomes active in order to ensure that the PIM state is preserved when a spoke-SDP recovers.

Appropriate actions based on the expiration of PIM related-timers on the standby router are only taken after it has become the active peer for the related object (after a failover).

PIM snooping for IPv4 synchronization is supported wherever PIM snooping for IPv4 is supported, excluding the following services:

- BGP-VPLS
- VPLS E-Tree
- management VPLS

See PIM Snooping for VPLS for more details.

PIM snooping for IPv4 synchronization is also only supported for the following active/standby redundancy mechanisms on dual-homed systems:

- multi-chassis LAG
- BGP Multi-homing

- active/standby pseudowires
- Single-Active Multi-homing (EVPN-MPLS and PBB-EVPN I-VPLS)

Configuring an mrouter port under an object that has the synchronizing of PIM snooping for IPv4 states enabled is not recommended. The mrouter port configuration adds a (*,*) entry into the MFIB, which causes all groups (and PIM messages) to be sent out of the respective object. In addition, the **mrouter port** command causes all PIM messages on that object to be discarded. However, the (*,*) entry is not synchronized by MCS. Consequently, the mrouter port could cause the two MCS peers to be forwarding different sets of multicast streams out of the related object when each is active.

3.2.20.7 VPLS Multicast-Aware High Availability Features

The following features are HA capable:

- Configuration redundancy — All the VPLS multicast-aware configurations can be synchronized to the standby CPM.
- Local snooping states as well as states distributed by LDP can be synchronized to the standby CPM.
- Operational states can also be synchronized, for example, the operational state of PIM proxy.

3.2.21 RSVP and LDP P2MP LSP for Forwarding VPLS/B-VPLS BUM and IP Multicast Packets

This feature enables the use of a P2MP LSP as the default tree for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLS instance. The P2MP LSP is referred to in this case as the Inclusive Provider Multicast Service Interface (I-PMSI).

When enabled, this feature relies on BGP Auto-Discovery (BGP-AD) or BGP-VPLS to discover the PE nodes participating in a specified VPLS/B-VPLS instance. The BGP route contains the information required to signal both the point-to-point (P2P) PWs used for forwarding unicast known Ethernet frames and the RSVP P2MP LSP used to forward the BUM frames. The root node signals the P2MP LSP based on an LSP template associated with the I-PMSI at configuration time. The leaf node will join automatically the P2MP LSP which matches the I-PMSI tunnel information discovered via BGP.

If IGMP or PIM snooping are configured on the VPLS instance, multicast packets matching a L2 multicast Forwarding Information Base (FIB) record will also be forwarded over the P2MP LSP.

The user enables the use of an RSVP P2MP LSP as the I-PMSI for forwarding Ethernet BUM and IP multicast packets in a VPLS/B-VPLS instance using the following commands:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>rsvp>lsp-template *p2mp-lsp-template-name*

The user enables the use of an LDP P2MP LSP as the I-PMSI for forwarding Ethernet BUM and IP multicast packets in a VPLS instance using the following command:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>mldp

After the user performs a ‘no shutdown’ under the context of the inclusive node and the expiration of a delay timer, BUM packets will be forwarded over an automatically signaled mLDP P2MP LSP or over an automatically signaled instance of the RSVP P2MP LSP specified in the LSP template.

The user can specify if the node is both root and leaf in the VPLS instance:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>root-and-leaf

The **root-and-leaf** command is required; otherwise, this node will behave as a leaf only node by default. When the node is leaf only for the I-PMSI of type P2MP RSVP LSP, no PMSI Tunnel Attribute is included in BGP-AD route update messages and thus no RSVP P2MP LSP is signaled but the node can join RSVP P2MP LSP rooted at other PE nodes participating in this VPLS/B-VPLS service. The user must still configure a LSP template even if the node is a leaf only. For the I-PMSI of type mLDP, the leaf-only node will join I-PMSI rooted at other nodes it discovered but will not include a PMSI Tunnel Attribute in BGP route update messages. This way, a leaf only node will forward packets to other nodes in the VPLS/B-VPLS using the point-to-point spoke-SDPs.

BGP-AD (or BGP-VPLS) must have been enabled in this VPLS/B-VPLS instance or the execution of the ‘no shutdown” command under the context of the inclusive node is failed and the I-PMSI will not come up.

Any change to the parameters of the I-PMSI, such as disabling the P2MP LSP type or changing the LSP template requires that the inclusive node be first shutdown. The LSP template is configured in MPLS.

If the P2MP LSP instance goes down, VPLS/B-VPLS immediately reverts the forwarding of BUM packets to the P2P PWs. The user can, however, restore at any time the forwarding of BUM packets over the P2P PWs by performing a ‘shutdown’ under the context of the inclusive node.

This feature is supported with VPLS, H-VPLS, B-VPLS and BGP-VPLS. It is not supported with I-VPLS and Routed VPLS.

3.2.22 MPLS Entropy Label and Hash Label

The router supports the MPLS entropy label (RFC 6790) and the Flow Aware Transport label (known as the hash label) (RFC 6391). These labels allow LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. See the *7450 ESS, 7750 SR, and 7950 XRS MPLS Guide* for further information.

3.3 Routed VPLS and I-VPLS

This section provides information about Routed VPLS and I-VPLS. Routed VPLS and I-VPLS applies to the 7450 ESS and 7750 SR.

3.3.1 IES or VPRN IP Interface Binding

A standard IP interface within an existing IES or VPRN service context may be bound to a service name. Subscriber and group IP interfaces are not allowed to bind to a VPLS or I-VPLS service context or I-VPLS. **For the remainder of this section Routed VPLS and Routed I-VPLS will both be described as a VPLS service and differences will be pointed out where applicable.** A VPLS service only supports binding for a single IP interface.

While an IP interface may only be bound to a single VPLS service, the routing context containing the IP interface (IES or VPRN) may have other IP interfaces bound to other VPLS service contexts of the same type (all VPLS or all I-VPLS). In other words, Routed VPLS allows the binding of IP interfaces in IES or VPRN services to be bound to VPLS services and Routed I-VPLS allows of IP interfaces in IES or VPRN services to be bound to I-VPLS services.

3.3.1.1 Assigning a Service Name to a VPLS Service

When a service name is applied to any service context, the name and service ID association is registered with the system. A service name cannot be assigned to more than one service ID.

Special consideration is given to a service name that is assigned to a VPLS service that has the **config>service>vpls>allow-ip-int-bind** command enabled. If a name is applied to the VPLS service while the flag is set, the system will scan the existing IES and VPRN services for an IP interface that is bound to the specified service name. If an IP interface is found, the IP interface will be attached to the VPLS service associated with the name. Only one interface can be bound to the specified name.

If the **allow-ip-int-bind** command is not enabled on the VPLS service, the system will not attempt to resolve the VPLS service name to an IP interface. As soon as the **allow-ip-int-bind** flag is configured on the VPLS, the corresponding IP interface will be bound and become operational up. There is no need to toggle the **shutdown/no shutdown** command.

If an IP interface is not currently bound to the service name used by the VPLS service, no action is taken at the time of the service name assignment.

3.3.1.2 Service Binding Requirements

In the event that the defined service ID is created on the system, the system will check to ensure that the service type is VPLS. If the service type is not VPLS or I-VPLS, service creation will not be allowed and the service ID will remain undefined within the system.

If the created service type is VPLS, the IP interface will be eligible to enter the operationally up state.

3.3.1.3 Bound Service Name Assignment

In the event that a bound service name is assigned to a service within the system, the system will first check to ensure the service type is VPLS or I-VPLS. Secondly the system will ensure that the service is not already bound to another IP interface via the service ID. If the service type is not VPLS or I-VPLS or the service is already bound to another IP interface via the service ID, the service name assignment will fail.

In the event that a single VPLS Service ID and service name is assigned to two separate IP interfaces, the VPLS service will not be allowed to enter and be operational/up state.

3.3.1.4 Binding a Service Name to an IP Interface

An IP interface within an IES or VPRN service context may be bound to a service name at anytime. Only one interface can be bound to a service.

When an IP interface is bound to a service name and the IP interface is administratively up, the system will scan for a VPLS service context using the name and take the following actions:

- If the name is not currently in use by a service, the IP interface will be placed in an operationally down: Non-existent service name or inappropriate service type state.

- If the name is currently in use by a non-VPLS service or the wrong type of VPLS service, the IP interface will be placed in the operationally down: Non-existent service name or inappropriate service type state.
- If the name is currently in use by a VPLS service without the **allow-ip-int-bind** flag set, the IP interface will be placed in the operationally down: VPLS service **allow-ip-int-bind** flag not set state. There is no need to toggle the **shutdown/ no shutdown** command.
- If the name is currently in use by a valid VPLS service and the **allow-ip-int-bind** flag is set, the IP interface will be eligible to be placed in the operationally up state depending on other operational criteria being met.

3.3.1.5 Bound Service Deletion or Service Name Removal

In the event that a VPLS service is deleted while bound to an IP interface, the IP interface will enter the ‘Down: Non-existent svc-ID’ operational state. If the IP interface was bound to the VPLS service name, the IP interface will enter the ‘Down: Non-existent svc-name’ operational state. No console warning is generated.

If the created service type is VPLS, the IP interface will be eligible to enter the operationally up state.

3.3.1.6 IP Interface Attached VPLS Service Constraints

Once a VPLS service has been bound to an IP interface through its service name, the service name assigned to the service cannot be removed or changed unless the IP interface is first unbound from the VPLS service name.

A VPLS service that is currently attached to an IP interface cannot be deleted from the system unless the IP interface is unbound from the VPLS service name.

The **allow-ip-int-bind** flag within an IP interface attached VPLS service cannot be reset. The IP interface must first be unbound from the VPLS service name to reset the flag.

3.3.1.7 IP Interface and VPLS Operational State Coordination

When the IP interface is successfully attached to a VPLS service, the operational state of the IP interface will be dependent upon the operational state of the VPLS service.

The VPLS service itself remains down until at least one virtual port (SAP, spoke-SDP or mesh SDP) is operational.

3.3.2 IP Interface MTU and Fragmentation

The VPLS service is affected by two MTU values; port MTUs and the VPLS service MTU. The MTU on each physical port defines the largest Layer 2packet (including all DLC headers) that may be transmitted out a port. The VPLS itself has a service level MTU that defines the largest packet supported by the service. This MTU does not include the local encapsulation overhead for each port (QinQ, Dot1Q, TopQ or SDP service delineation fields and headers) but does include the remainder of the packet. As virtual ports are created in the system, the virtual port cannot become operational unless the configured port MTU minus the virtual port service delineation overhead is greater than or equal to the configured VPLS service MTU. Thus, an operational virtual port is ensured to support the largest packet traversing the VPLS service. The service delineation overhead on each Layer 2 packet is removed before forwarding into a VPLS service. VPLS services do not support fragmentation and must discard any Layer 2 packet larger than the service MTU after the service delineation overhead is removed.

When an IP interface is associated with a VPLS service, the IP-MTU is based on either the administrative value configured for the IP interface or an operational value derived from VPLS service MTU. The operational IP-MTU cannot be greater than the VPLS service MTU minus 14 bytes.

- If the configured (administrative) IP-MTU is configured for a value greater than the normalized IP-MTU, based on the VPLS service-MTU, then the operational IP-MTU is reset to equal the normalized IP-MTU value (VPLS service MTU – 14 bytes).
- If the configured (administrative) IP-MTU is configured for a value less than or equal to the normalized IP-MTU, based on the VPLS service-MTU, then the operational IP-MTU is set to equal the configured (administrative) IP-MTU value.

3.3.2.1 Unicast IP Routing into a VPLS Service

The VPLS service MTU and the IP interface MTU parameters may be changed at anytime.

3.3.3 ARP and VPLS FDB Interactions

Two address-oriented table entries are used when routing into a VPLS service. On the routing side, an ARP entry is used to determine the destination MAC address used by an IP next-hop. In the case where the destination IP address in the routed packet is a host on the local subnet represented by the VPLS instance, the destination IP address itself is used as the next-hop IP address in the ARP cache lookup. If the destination IP address is in a remote subnet that is reached by another router attached to the VPLS service, the routing lookup will return the local IP address on the VPLS service of the remote router will be returned. If the next-hop is not currently in the ARP cache, the system will generate an ARP request to determine the destination MAC address associated with the next-hop IP address. IP routing to all destination hosts associated with the next-hop IP address stops until the ARP cache is populated with an entry for the next-hop. The ARP cache may be populated with a static ARP entry for the next-hop IP address. While dynamically populated ARP entries will age out according to the ARP aging timer, static ARP entries never age out.

The second address table entry that affects VPLS routed packets is the MAC destination lookup in the VPLS service context. The MAC associated with the ARP table entry for the IP next-hop may or may not currently be populated in the VPLS Layer 2 FDB table. While the destination MAC is unknown (not populated in the VPLS FDB), the system will flood all packets destined for that MAC (routed or bridged) to all virtual ports within the VPLS service context. Once the MAC is known (populated in the VPLS FDB), all packets destined for the MAC (routed or bridged) will be targeted to the specific virtual port where the MAC has been learned. As with ARP entries, static MAC entries may be created in the VPLS FDB. Dynamically learned MAC addresses are allowed to age out or be flushed from the VPLS FDB while static MAC entries always remain associated with a specific virtual port. Dynamic MACs may also be relearned on another VPLS virtual port than the current virtual port in the FDB. In this case, the system will automatically move the MAC FDB entry to the new VPLS virtual port.

The MAC address associated with the routed VPLS IP interface is protected within its VPLS service such that frames received with this MAC address as the source address are discarded. VRRP MAC addresses are not protected in this way.

3.3.3.1 Routed VPLS Specific ARP Cache Behavior

In typical routing behavior, the system uses the IP route table to select the egress interface and then at the egress forwarding engine, an ARP entry is used forward the packet to the appropriate Ethernet MAC. With routed VPLS, the egress IP interface may be represented by multiple egress forwarding engine (wherever the VPLS service virtual ports exists).

To optimize routing performance, the ingress forwarding engine processing has been augmented to perform an ingress ARP lookup in order to resolve which VPLS MAC address the IP frame must be routed toward. This MAC address may be currently known or unknown within the VPLS FDB. If the MAC is unknown, the packet is flooded by the ingress forwarding engine to all egress forwarding engines where the VPLS service exists. When the MAC is known on a virtual port, the ingress forwarding engine forwards the packet to the correct egress forwarding engine. Table 36 describes how the ARP cache and MAC FDB entry states interact at ingress and Table 37 describes the corresponding egress behavior.

Table 36 **Ingress Routed to VPLS Next-Hop Behavior**

Next-Hop ARP Cache Entry	Next-Hop MAC FDB Entry	Ingress Behavior
ARP Cache Miss (No Entry)	Known or Unknown	Flood to all egress forwarding engines associated with the VPLS or I-VPLS context.
	Unknown	Flood to all egress forwarding engines associated with the VPLS or I-VPLS context.
	Unknown	Flood to all egress forwarding engines associated with the VPLS for forwarding out all VPLS or I-VPLS virtual ports.

Table 37 **Egress Routed VPLS Next-Hop Behavior**

Next-Hop ARP Cache Entry	Next-Hop MAC FDB Entry	Egress Behavior
ARP Cache Miss (No Entry)2	Known	No ARP entry. The MAC address is unknown and the ARP request is flooded out of all virtual ports of the VPLS or I-VPLS instance.
	Unknown	Request control engine ARP processing ARP request transmitted out all virtual port associated with the VPLS or I-VPLS service. Only the first egress forwarding engine ARP processing request triggers egress ARP request.

Table 37 **Egress Routed VPLS Next-Hop Behavior (Continued)**

Next-Hop ARP Cache Entry	Next-Hop MAC FDB Entry	Egress Behavior
ARP Cache Hit	Known	Forward out specific egress VPLS or I-VPLS virtual port where MAC has been learned.
	Unknown	Flood to all egress VPLS or I-VPLS virtual ports on forwarding engine.

3.3.4 The allow-ip-int-bind VPLS Flag

The **allow-ip-int-bind** flag on a VPLS service context is used to inform the system that the VPLS service is enabled for routing support. The system uses the setting of the flag as a key to determine what type of ports and which type of forwarding planes the VPLS service may span.

The system also uses the flag state to define which VPLS features are configurable on the VPLS service to prevent enabling a feature that is not supported when routing support is enabled.

3.3.4.1 Routed VPLS SAPs Only Supported on Standard Ethernet Ports

The **allow-ip-int-bind** flag is set (routing support enabled) on a VPLS/I-VPLS service. SAPs within the service can be created on standard Ethernet, HSMDA, and CCAG ports. ATM and POS are not supported.

3.3.4.2 LAG Port Membership Constraints

If a LAG has a non-supported port type as a member, a SAP for the routing-enabled VPLS service cannot be created on the LAG. Once one or more routing enabled VPLS SAPs are associated with a LAG, a non-supported Ethernet port type cannot be added to the LAG membership.

3.3.4.3 Routed VPLS Feature Restrictions

When the **allow-ip-int-bind** flag is set on a VPLS service, the following restrictions apply. The flag also cannot be enabled while any of these features are applied to the VPLS service:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined.
- SDPs used in spoke or mesh SDP bindings cannot be configured as GRE.
- The VPLS service type cannot be B-VPLS or M-VPLS.
- MVR from Routed VPLS and to another SAP is not supported.
- Enhanced and Basic Subscriber Management (BSM) features cannot be enabled.
- Network domain on SDP bindings cannot be enabled.
- Per-service hashing is not supported.
- BGP-VPLS is not supported.
- Ingress queuing for split horizon groups is not supported.
- Multiple virtual routers are not supported.

3.3.4.4 Routed I-VPLS Feature Restrictions

The following restrictions apply to routed I-VPLS:

- Multicast is not supported.
- VC-VLANs are not supported on SDPs.
- **force-qtag-forwarding** is not supported.
- Control words are not supported on B-VPLS SDPs.
- Hash Label is not supported on B-VPLS SDPs.

3.3.5 IPv4 and IPv6 Multicast Routing Support

IPv4 and IPv6 multicast routing is supported in a routed VPLS service through its IP interface when the source of the multicast stream is on one side of its IP interface and the receivers are on either side of the IP interface. For example, the source for multicast stream G1 could be on the IP side sending to receivers on both other regular IP interfaces and the VPLS of the routed VPLS service, while the source for group G2 could be on the VPLS side sending to receivers on both the VPLS and IP side of the routed VPLS service.

IPv4 and IPv6 multicast routing is not supported with Multicast VLAN Registration functions or the configuration of a video interface within the associated VPLS service. It is also not supported in a routed I-VPLS service or in BGP EVPN services. Forwarding IPv4 or IPv6 multicast traffic from the routed VPLS IP interface into its VPLS service on a P2MP LSP is not supported.

The IP interface of a routed VPLS supports the configuration of both PIM and IGMP for IPv4 multicast and for both PIM and MLD for IPv6 multicast.

To forward IPv4/IPv6 multicast traffic from the VPLS side of the routed VPLS service to the IP side, the **forward-ipv4-multicast-to-ip-int** and/or **forward-ipv6-multicast-to-ip-int** parameters must be configured as shown below:

```
configure
    service
        vpls <service-id>
            allow-ip-int-bind
                forward-ipv4-multicast-to-ip-int
                forward-ipv6-multicast-to-ip-int
            exit
        exit
    exit
exit
```

Enabling IGMP snooping or MLD snooping in the VPLS service is optional. If IGMP/MLD snooping is enabled, IGMP/MLD must be enabled on the routed VPLS IP interface in order for multicast traffic to be sent into, or received from, the VPLS service. IPv6 multicast uses MAC-based forwarding, see MAC-Based IPv6 Multicast Forwarding for more information.

If both IGMP/MLD and PIM for IPv4/IPv6 are configured on the routed VPLS IP interface in a redundant PE topology, the associated IP interface on one of the PEs must be configured as both the PIM designated router and the IGMP/MLD querier in order that the multicast traffic is sent into the VPLS service, as IGMP/MLD joins are only propagated to the IP interface if it is the IGMP/MLD querier. An alternative to this is to configure the routed VPLS IP interface in the VPLS service as an mrouter port as follows:

```
configure
    service
        vpls <service-id>
            allow-ip-int-bind
                igmp-snooping
                    mrouter-port
                mld-snooping
                    mrouter-port
            exit
        exit
    exit
exit
```

This configuration achieves a faster failover in scenarios with redundant routers where multicast traffic is sent to systems on the VPLS side of their routed VPLS services and IGMP/MLD snooping is enabled in the VPLS service. If the active router fails, the remaining router does not have to wait until it sends an IGMP/MLD query into the VPLS service before it starts receiving IGMP/MLD joins, and starts sending the multicast traffic into the VPLS service. When the mrouter port is configured as above, all IGMP/MLD joins (and multicast traffic) are sent to the VPLS service IP interface.

IGMP/MLD snooping should only be enabled when systems, as opposed to PIM routers, are connected to the VPLS service. If IGMP/MLD snooping is enabled when the VPLS service is used for transit traffic for connected PIM routers, the IGMP/MLD snooping would prevent multicast traffic being forwarded between the PIM routers (as PIM snooping is not supported). A workaround would be to configure the VPLS SAPs and spoke-SDPs (and the routed VPLS IP interface) to which the PIM routers are connected as mrouter ports.

If IMPM is enabled on an FP on which there is a routed VPLS service with **forward-ipv4-multicast-to-ip-int** or **forward-ipv6-multicast-to-ip-int** configured, the IPv4/IPv6 multicast traffic received in the VPLS service that is forwarded through the IP interface will be IMPM-managed even without IGMP/MLD snooping being enabled. This does not apply to traffic that is only flooded within the VPLS service.

When IPv4/IPv6 multicast traffic is forwarded from a VPLS SAP through the routed VPLS IP interface, the packet count is doubled in the following statistics to represent both the VPLS and IP replication (this reflects the capacity used for this traffic on the ingress queues, which is subject to any configured rates and IMPM capacity management):

- Offered queue statistics
- IMPM managed statistics
- IMPM unmanaged statistics for policed traffic

IPv4 or IPv6 multicast traffic entering the IP side of the routed VPLS service and exiting over a multi-port LAG on the VPLS side of the service is sent on a single link of that egress LAG, specifically the link used for all broadcast, unknown and multicast traffic.

An example of IPv4/IPv6 multicast in a routed VPLS service is shown in Figure 86. There are two routed VPLS IP interfaces connected to an IES service with the upper interface connected to a VPLS service in which there is a PIM router and the lower interface connected to a VPLS service in which there is a system using IGMP/MLD.

Figure 86 **IPv4/IPv6 Multicast with a Router VPLS Service**

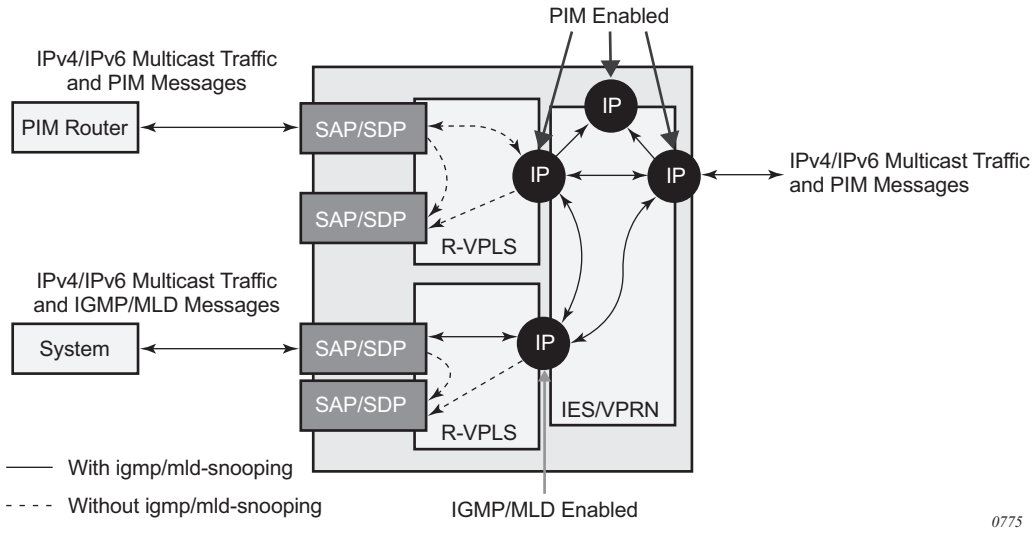

The IPv4/IPv6 multicast traffic entering the IES/VPRN service through the regular IP interface is replicated to both the other regular IP interface and the two routed VPLS interfaces if PIM/IGMP/MLD joins have been received on the respective IP interfaces. This traffic will be flooded into both VPLS services unless IGMP/MLD snooping is enabled in the lower VPLS service, in which case it is only sent to the system originating the IGMP/MLD join.

The IPv4/IPv6 multicast traffic entering the upper VPLS service from the connected PIM router will be flooded in that VPLS service and, if related joins have been received, forwarded to the regular IP interfaces in the IES/VPRN. It will also be forwarded to the lower VPLS service if an IGMP/MLD join is received on its IP interface, and will be flooded in that VPLS service unless IGMP/MLD snooping is enabled.

The IPv4/IPv6 multicast traffic entering the lower VPLS service from the connected system will be flooded in that VPLS service, unless IGMP/MLD snooping is enabled, in which case it will only be forwarded to SAPs, spoke-SDPs, or the routed VPLS IP interface if joins have been received on them. It will be forwarded to the regular IP interfaces in the IES/VPRN service if related joins have been received on those interfaces, and it will also be forwarded to the upper VPLS service if a PIM IPv4/IPv6 join is received on its IP interface, this being flooded in that VPLS service.

3.3.6 BGP Auto-Discovery (BGP-AD) for Routed VPLS Support

BGP Auto-Discovery (BGP-AD) for Routed VPLS is supported. BGP-AD for LDP VPLS is an already supported framework for automatically discovering the endpoints of a Layer 2 VPN offering an operational model similar to that of an IP VPN.

3.3.7 Routed VPLS Caveats

3.3.7.1 VPLS SAP Ingress IP Filter Override

When an IP Interface is attached to a VPLS or an I-VPLS service context, the VPLS SAP provisioned IP filter for ingress routed packets may be optionally overridden in order to provide special ingress filtering for routed packets. This allows different filtering for routed packets and non-routed packets. The filter override is defined on the IP interface bound to the VPLS service name. A separate override filter may be specified for IPv4 and IPv6 packet types.

If a filter for a specified packet type (IPv4 or IPv6) is not overridden, the SAP specified filter is applied to the packet (if defined).

3.3.7.2 IP Interface Defined Egress QoS Reclassification

The SAP egress QoS policy defined forwarding class and profile reclassification rules are not applied to egress routed packets. To allow for egress reclassification, a SAP egress QoS policy ID may be optionally defined on the IP interface which will be applied to routed packets that egress the SAPs on the VPLS or I-VPLS service associated with the IP interface. Both unicast directed and MAC unknown flooded traffic apply to this rule. Only the reclassification portion of the QoS policy is applied which includes IP precedence or DSCP classification rules and any defined IP match criteria and their associated actions.

The policers and queues defined within the QoS policy applied to the IP interface are not created on the egress SAPs of the VPLS service. Instead, the actual QoS policy applied to the egress SAPs defines the egress policers and queues that will be used by both routed and non-routed egress packets. The forwarding class mappings defined in the egress SAP's QoS policy will also define which policer or queue will handle each forwarding class for both routed and non-routed packets.

3.3.7.3 Remarking for VPLS and Routed Packets

The remarking of packets to and from an IP interface in an R-VPLS service corresponds to that supported on IP interface, even though the packets ingress or egress a SAP in the VPLS service bound to the IP service. Specifically, this results in the ability to remark the DSCP/prec for these packets.

Packets ingressing and egressing SAPs in the VPLS service (not routed through the IP interface) support the regular VPLS QoS and therefore the DSCP/prec cannot be remarked.

3.3.7.4 7450 Mixed Mode Chassis

The mixed mode on the 7450 ESS that allows 7750 SR-based IOM3s to be populated and operational in a 7450 ESS chassis supports routed VPLS as long as all the forwarding plane and port type restrictions are observed.

3.3.7.5 IPv4 Multicast Routing

When using IPv4 Multicast routing, the following are not supported:

- Multicast VLAN registration functions within the associated VPLS service.
- The configuration of a video ISA within the associated VPLS service.
- The configuration of MFIB-allowed MDA destinations under spoke/mesh SDPs within the associated VPLS service.
- IPv4 multicast routing is not supported in Routed I-VPLS.
- RFC 6037 multicast tunnel termination (including when the system is a bud node) is not supported on the routed VPLS IP interface for multicast traffic received in the VPLS service.
- Forwarding of multicast traffic from the VPLS side of the service to the IP interface side of the service is not supported for routed VPLS services in which VXLAN is enabled.

3.3.7.6 Routed VPLS Supported Routing Related Protocols

The following protocols are supported on IP interfaces bound to a VPLS service:

- BGP
- OSPF
- ISIS
- PIM
- IGMP
- BFD
- VRRP
- ARP
- DHCP Relay

3.3.7.7 Spanning Tree and Split Horizon

A routed VPLS context supports all spanning tree and split horizon capabilities that a non-routed VPLS service supports.

3.4 VPLS Service Considerations

This section describes the 7450 ESS, 7750 SR, and 7950 XRS service features and any special capabilities or considerations as they relate to VPLS services.

3.4.1 SAP Encapsulations

VPLS services are designed to carry Ethernet frame payloads, so it can provide connectivity between any SAPs and SDPs that pass Ethernet frames. The following SAP encapsulations are supported on the 7450 ESS, 7750 SR, and 7950 XRS VPLS service:

- Ethernet null
- Ethernet dot1q
- Ethernet QinQ
- SONET/SDH BCP-null
- SONET/SDH BCP-dot1q
- ATM VC with RFC 2684 Ethernet bridged encapsulation (See ATM/Frame Relay PVC Access and Termination on a VPLS Service.)
- FR VC with RFC 2427 Ethernet bridged encapsulation (See ATM/Frame Relay PVC Access and Termination on a VPLS Service.)

3.4.2 VLAN Processing

The SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs:

1. Null encapsulation defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP.
2. dot1q encapsulation defined on ingress — Only first label is considered.
3. QinQ encapsulation defined on ingress— Both labels are considered. The SAP can be defined with a wildcard for the inner label (for example, “100:100.*”). In this situation all packets with an outer label of 100 will be treated as belonging to the SAP. If, on the same physical link, there is also a SAP defined with a QinQ encapsulation of 100:100.1, then traffic with 100:1 will go to that SAP and all other traffic with 100 as the first label will go to the SAP with the 100:100.* definition.

In situations 2 and 3 above, traffic encapsulated with tags for which there is no definition are discarded.

3.4.3 Ingress VLAN Swapping

This feature is supported on VPLS and VLL service where the end-to-end solution is built using two node solutions (requiring SDP connections between the nodes).

In VLAN swapping, only the VLAN-id value will be copied to the inner VLAN position. Ethertype of the inner tag will be preserved and all consecutive nodes will work with that value. Similarly, the dot1p bits value of outer-tag will not be preserved.

Figure 87 **Ingress VLAN Swapping**

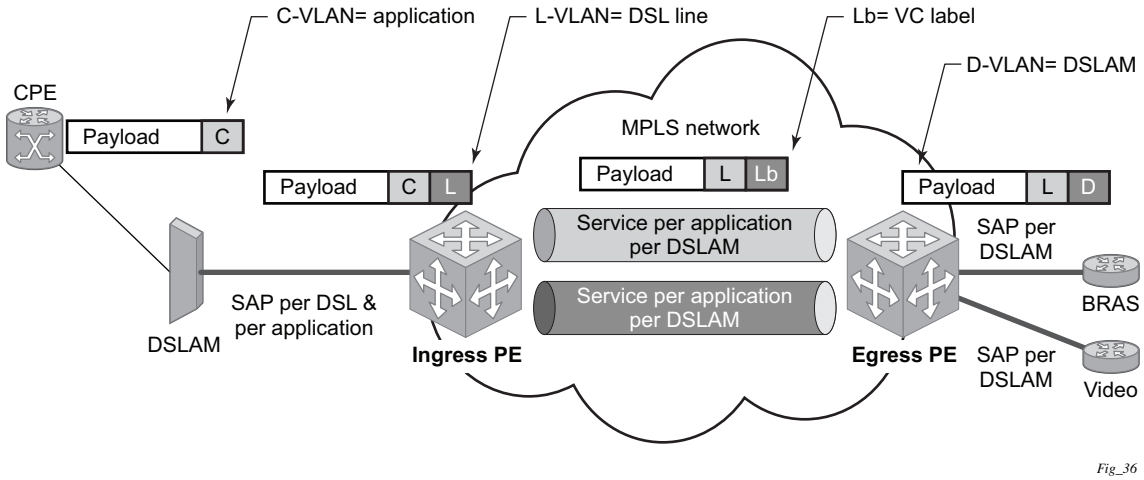

Figure 87 describes the network where at user access side (DSLAM facing SAPs) every subscriber is represented by several QinQ SAPs with inner-tag encoding service and outer-tag encoding subscriber (DSL line). The aggregation side (BRAS or PE facing SAPs) the is represented by DSL line number (inner VLAN tag) and DSLAM (outer VLAN tag). The effective operation on VLAN tag is to drop inner tag at access side and push another tag at the aggregation side.

3.4.4 Service Auto-Discovery using Multiple VLAN Registration Protocol (MVRP)

IEEE 802.1ak Multiple VLAN Registration Protocol (MVRP) is used to advertise throughout a native Ethernet switching domain one or multiple VLAN IDs to build automatically native Ethernet connectivity for multiple services. These VLAN IDs can be either Customer VLAN IDs (CVID) in an enterprise switching environment, Stacked VLAN IDs (SVID) in a Provider Bridging, QinQ Domain (refer to IEEE 802.1ad) or Backbone VLAN IDs (BVID) in a Provider Backbone Bridging (PBB) domain (refer to IEEE 802.1ah).

The initial focus of Nokia MVRP implementation is a Service Provider QinQ domain with or without a PBB core. The QinQ access into a PBB core example is used throughout this section to describe the MVRP implementation. With the exception of end-station components, a similar solution can be used to address a QinQ only or enterprise environments.

The components involved in the MVRP control plane are shown in Figure 88.

Figure 88 **Infrastructure for MVRP Exchanges**

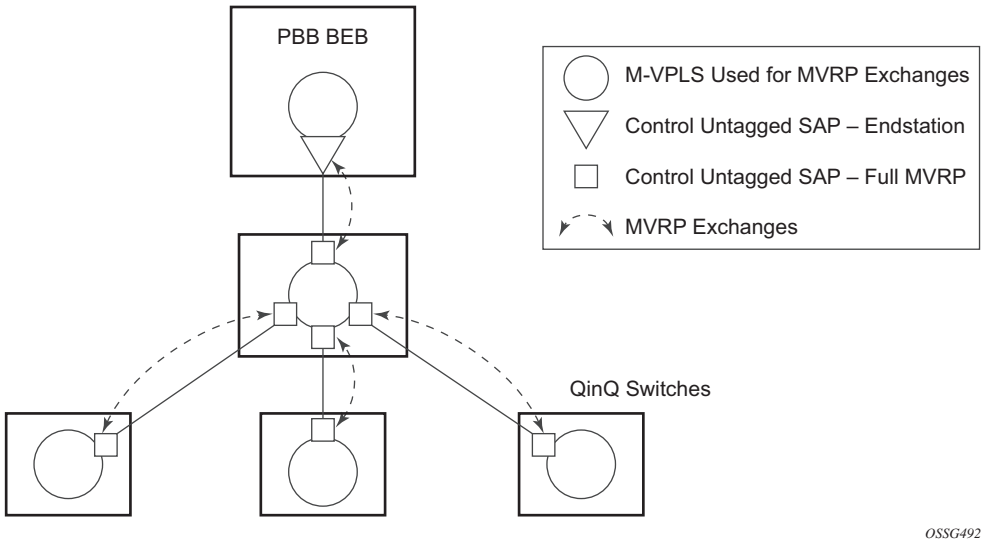

All the devices involved are QinQ switches with the exception of the PBB BEB which delimits the QinQ domain and ensures the transition to the PBB core. The red circles represent Management VPLS instances interconnected by SAPs to build a native Ethernet switching domain used for MVRP control plane exchanges.

The following high level steps are involved in auto-discovery of VLAN connectivity in a native Ethernet domain using MVRP:

- Configure the MVRP infrastructure
  - This involves the configuration of a Management VPLS (M-VPLS) context
  - MSTP may be used in M-VPLS to provide the loop-free topology over which the MVRP exchanges take place.
- Instantiate related VLAN FDB, trunks in the MVRP, M-VPLS scope
  - The VLAN FDBs (VPLS instances) and associated trunks (SAPs) are instantiated in the same Ethernet switches and on the same “trunk ports” as the M-VPLS
  - There is no need to instantiate data VPLS instances in the BEB. I-VPLS instances and related downward facing SAPs will be provisioned manually because the ISID to VLAN association must be configured.
- MVRP activation of service connectivity
  - When the first two customer UNI and/or PBB end-station SAPs are configured on different Ethernet switches in a certain service context the MVRP exchanges will activate service connectivity

3.4.4.1 Configure the MVRP Infrastructure using an M-VPLS Context

The following provisioning steps apply:

- Configure M-VPLS instances in the switches that will participate in MVRP control plane
- Configure under the M-VPLS the untagged SAPs to be used for MVRP exchanges; only dot1q or qinq ports are accepted for MVRP enabled M-VPLS
- Configure MVRP parameters at M-VPLS instance or SAP level

3.4.4.2 Instantiate Related VLAN FDBs and Trunks in MVRP Scope

This involves the configuration in the M-VPLS, under vpls-group of the following attributes: VLAN ranges, vpls-template and vpls-sap-template bindings. As soon as the VPLS group is enabled the configured attributes are used to auto-instantiate on a per VLAN basis a VPLS FDB and related SAPs in the switches and on the “trunk ports” specified in the M-VPLS context. The trunk ports are ports associated with an M-VPLS SAP not configured as an end-station.

The following procedure is used:

- The vpls-template binding is used to instantiate the VPLS instance where the service ID is derived from the VLAN value as per service-range configuration
- The vpls-sap-template binding is used to create dot1q SAPs by deriving from the VLAN value the service delimiter as per service-range configuration

The above procedure may be used outside of the MVRP context to pre-provision a large number of VPLS contexts that share the same infrastructure and attributes.

The MVRP control of the auto-instantiated services can be enabled using the **mvrp-contrl** command under vpls-group:

- If mvrp-control is disabled the auto-created VPLS instances and related SAPs are ready to forward.
- If mvrp-control is enabled the auto-created VPLS instances will be instantiated initially with an empty flooding domain. The MVRP exchanges will gradually enabled service connectivity according to the operator configuration – between configured SAPs in the data VPLS context
  - This provides also protection against operational mistakes that may generate flooding throughout the auto-instantiated VLAN FDBs.

From an MVRP perspective these SAPs can be either “full MVRP” or “end-stations” interfaces.

A full MVRP interface is a full participant in the local M-VPLS scope:

- VLAN attributes received in an MVRP registration on this MVRP interface are declared on all the other full MVRP SAPs in the control VPLS.
- VLAN attributes received in an MVRP registration on other full MVRP interfaces in the local M-VPLS context are declared on this MVRP interface.

In an MVRP end-station the attributes registered on that interface have local significance:

- VLAN attributes received in an MVRP registration on this interface are not declared on any other MVRP SAPs in the control VPLS. The attributes are registered only on the local port.
- Only locally active VLAN attributes are declared on the end-station interface; VLAN attributes registered on any other MVRP interfaces are not declared on end-station interfaces
- Also defining an M-VPLS SAP as end-station does not instantiate any objects on the local switch; the command is used just to define which SAP needs to be monitored by MVRP to declare the related VLAN value.

The following example describes the M-VPLS configuration required to auto-instantiate the VLAN FDBs and related trunks in non-PBB switches:

Example:

```
mrp
     no shutdown
     mmrp
          shutdown
     mvrp
          no shutdown
sap 1/1/1:0
     mrp mvrp
          no shutdown
sap 2/1/2:0
     mrp mvrp
          no shutdown
sap 3/1/10:0
     mrp mvrp
          no shutdown
vpls-group 1
     service-range 100-2000
     vpls-template-binding Autovpls1
     sap-template-binding Autosap1
          mvrp-control
     no shutdown
```

A similar M-VPLS configuration may be used to auto-instantiate the VLAN FDBs and related trunks in PBB switches. The vpls-group command is replaced by the end-station command under the downwards SAPs as in the following example:

Example:

```
config>service>vpls control-mvrp m-vpls create customer
 1
     [..]
     sap 1/1/1:0
          mvrp mvrp
               endstation-vid-group 1 vlan-id 100-2000
               no shutdown
```

3.4.4.3 MVRP Activation of Service Connectivity

As new Ethernet services are activated, UNI SAPs need to be configured and associated with the VLAN IDs (VPLS instances) auto-created using the procedures described in the previous sections. These UNI SAPs may be located in the same VLAN domain or over a PBB backbone. When UNI SAPs are located in different VLAN domains, an intermediate service translation point must be used at the PBB BEB which maps the local VLAN ID through an I-VPLS SAP to a PBB ISID. This BEB SAP will be playing the role of an end-station from an MVRP perspective for the local VLAN domain. This section will discuss how MVRP is used to activate service

connectivity between a BEB SAP and a UNI SAP located on one of the switches in the local domain. Similar procedure is used for the case of UNI SAPs configured on two switches located in the same access domain. No end-station configuration is required on the PBB BEB if all the UNI SAPs in a service are located in the same VLAN domain.

The service connectivity instantiation through MVRP is shown in Figure 89.

Figure 89 **Service Instantiation with MVRP - QinQ to PBB Example**

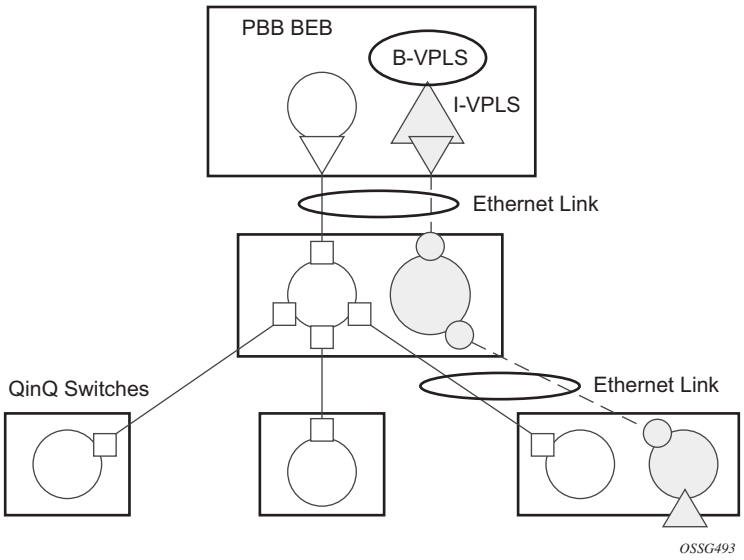

In this example the UNI and service translation SAPs are configured in the data VPLS represented by the yellow circle. This instance and associated trunk SAPs were instantiated using the procedures described in the previous sections. The following configuration steps are involved:

- on the BEB an I-VPLS SAP must be configured toward the local switching domain – see yellow triangle facing downwards
- on the UNI facing the customer a “customer” SAP is configured on the bottom left switch – see yellow triangle facing upwards

As soon as the first UNI SAP becomes “active” in the data VPLS on the ES, the associated VLAN value is advertised by MVRP throughout the related M-VPLS context. As soon as the second UNI SAP becomes available on a different switch or in our example on the PBB BEB the MVRP proceeds to advertise the associated VLAN value throughout the same M-VPLS. The trunks that experience MVRP declaration and registration in both directions will become active instantiating service connectivity as represented by the big and small yellow circles shown in the picture.

A hold-time parameter (**config>service>vpls>mrp>mvrp>hold-time**) is provided in the M-VPLS configuration to control when the end-station or last UNI SAP is considered active from an MVRP perspective. The hold-time controls the amount of MVRP advertisements generated on fast transitions of the end-station or UNI SAPs.

If the **no hold-time** setting is used:

- MVRP will stop declaring the VLAN only when the last provisioned UNI SAP associated locally with the service is deleted.
- MVRP will start declare the VLAN as soon as the first provisioned SAP is created in the associated VPLS instance, regardless of the operational state of the SAP.

If a non-zero “hold-time” setting is used:

- When a SAP in down state is added, MVRP does not declare the associated VLAN attribute. The attribute is declared immediately when the SAP comes up.
- When the SAP goes down, MVRP will wait until “hold-time” expiry before withdrawing the declaration.

For QinQ endstation SAPs only “no hold-time” setting is allowed

Only the following PBB Epipe and I-VPLS SAP types are eligible to activate MVRP declarations:

- dot1q: for example 1/1/2:100
- qinq or qinq default: for example, 1/1/1:100.1 and respectively 1/1/1:100.*; the outer VLAN 100 will be used as MVRP attribute as long as it belongs to the MVRP range configured for the port
- null port and dot1q default cannot be used

An example of steps required to activate service connectivity for VLAN 100 using MVRP follows.

In the data VPLS instance (VLAN 100) controlled by MVRP, on the QinQ switch:

Example:
```
config>service>vpls 100
     sap 9/1/1:10 //UNI sap using CVID 10 as service
       delimiter
          no shutdown
```

In I-VPLS on PBB BEB:

Example:
```
config>service>vpls 1000 i-vpls
     sap 8/1/2:100 //sap (using MVRP VLAN 100 on
       endstation port in M-VPLS
          no shutdown
```

3.4.4.4 MVRP Control Plane

MVRP is based on the IEEE 802.1ak MRP specification where STP is the supported method to be used for loop avoidance in a native Ethernet environment. M-VPLS and associated MSTP (or P-MSTP) control plane provides the loop avoidance component in Nokia implementation. Nokia MVRP may be used also in a non-MSTP, loop free topology.

3.4.4.5 STP-MVRP Interaction

Table 38 captures the expected interaction between STP (MSTP or P-MSTP) and MVRP:

Table 38 **MSTP and MVRP Interaction Table**

Item	M-VPLS Service xSTP	M-VPLS SAP STP	Register/Declare Data VPLS VLAN on M-VPLS SAP	DSFS (Data SAP Forwarding State) controlled by	Data Path Forwarding with MVRP enabled controlled by
1	(p)MSTP	Enabled	Based on M-VPLS SAP's MSTP forwarding state	MSTP only	DSFS and MVRP
2	(p)MSTP	Disabled	Based on M-VPLS SAP's oper state	None	MVRP
3	Disabled	Enabled or Disabled	Based on M-VPLS SAP's oper state	None	MVRP

Notes:

- Running STP in data VPLS instances controlled by MVRP is not allowed.
- Running STP on MVRP-controlled end-station SAPs is not allowed.

3.4.4.5.1 Interaction Between MVRP and Instantiated SAP Status

This section describes how MVRP reacts to changes in the instantiated SAP status.

There are a number of mechanisms that may generate operational or admin down status for the SAPs and VPLS instances controlled by MVRP:

1. Port down
2. MAC Move
3. Port MTU too small
4. Service MTU too small

The shutdown of the whole instantiated VPLS or instantiated SAPs is disabled in both VPLS and VPLS SAP templates. The **no shutdown** option is automatically configured.

In the **port down** case MVRP will also be operationally down on the port so no VLAN declaration will take place.

When MAC move is enabled in a data VPLS controlled by MVRP, in case a MAC move hit happens, one of the instantiated SAPs controlled by MVRP may be blocked. The SAP blocking by MAC Move is not reported though to the MVRP control plane. As a result MVRP keeps declaring and registering the related VLAN value on the control SAPs including the one which shares the same port with the instantiate SAP blocked by MAC move as long as MVRP conditions are met. For MVRP, an active control SAP is one that has MVRP enabled and MSTP is not blocking it for the VLAN value on the port. Also in the related data VPLS one of the two conditions must be met for the declaration of the VLAN value: there must be either a local user SAP or at least one MVRP registration received on one of the control SAPs for that VLAN.

In the last two cases VLAN attributes get declared or registered even when the instantiated SAP is operationally down, similarly with the MAC move case.

3.4.4.5.2 Using Temporary Flooding to Optimize Failover Times

MVRP advertisements use the active topology which may be controlled through loop avoidance mechanisms like MSTP. When the active topology changes as a result of network failures, the time it takes for MVRP to bring up the optimal service connectivity may be added on top of the regular MSTP convergence time. Full connectivity also depends on the time it takes for the system to complete flushing of bad MAC entries.

To minimize the effects of MAC Flushing and MVRP convergence, a temporary flooding behavior is implemented. When enabled the temporary flooding eliminates the time it takes to flush the MAC tables. In the initial implementation the temporary flooding is initiated only on reception of an STP TCN.

While temporary flooding is active all the frames received in the extended data VPLS context are flooded while the MAC flush and MVRP convergence takes place. The extended data VPLS context comprises all instantiated trunk SAPs regardless of MVRP activation status. A timer option is also available to configure a fixed amount of time, in seconds, during which all traffic is flooded (BUM or known unicast). Once the flood-time expires, traffic will be delivered according to the regular FDB content. The timer value should be configured to allow auxiliary processes like MAC Flush and MVRP to converge. The temporary flooding behavior applies to all VPLS types. MAC learning continues during temporary flooding. Temporary flooding behavior is enabled using the temp-flooding command under **config> service>vpls** or **config> service>template>vpls-template** contexts and is supported in VPLS regardless of whether MVRP is enabled or not.

The following rules apply for temporary flooding in VPLS:

- If discard-unknown is enabled then there is no temporary flooding
- Temporary flooding while active applies also to static MAC entries; after the MAC FDB is flushed it reverts back to the static MAC entries
- If MAC learning is disabled fast or temporary flooding is still enabled
- Temporary flooding is not supported in B-VPLS context when MMRP is enabled. The use of flood-time procedure provides a better procedure for this kind of environment.

3.4.5 VPLS E-Tree Services

This section describes VPLS E-Tree services.

3.4.5.1 VPLS E-Tree Services Overview

The VPLS E-Tree service offers a VPLS service with Root and Leaf designated access SAPs and SDP bindings, which prevent any traffic flow from leaf to leaf directly. With a VPLS E-Tree the split horizon group capability is inherent for leaf SAPs (or SDP bindings) and extends to all the remote PEs part of the same VPLS E-Tree service. This feature is based on IETF Draft *draft-ietf-l2vpn-vpls-pe-etree*.

A VPLS E-Tree service may support an arbitrary number of leaf access (leaf-ac) interfaces, root access (root-ac) interfaces and root-leaf tagged (root-leaf-tag) interfaces. Leaf-ac interfaces are supported on SAPs and SDP binds and can only communicate with root-ac interfaces (also supported on SAPs and SDP binds). Leaf-ac to leaf-ac communication is not allowed. Root-leaf-tag interfaces (supported on SAPs and SDP bindings) are tagged with root and leaf VIDs to allow remote VPLS instances to enforce the E-Tree forwarding.

Figure 90 shows a network with two root-ac interfaces and several leaf-ac SAPs (also could be SDPs). The diagram indicates two VIDs in use to each service within the service with no restrictions on the AC interfaces. The service guarantees no leaf-ac to leaf-ac traffic.

Figure 90 **E-Tree Service**

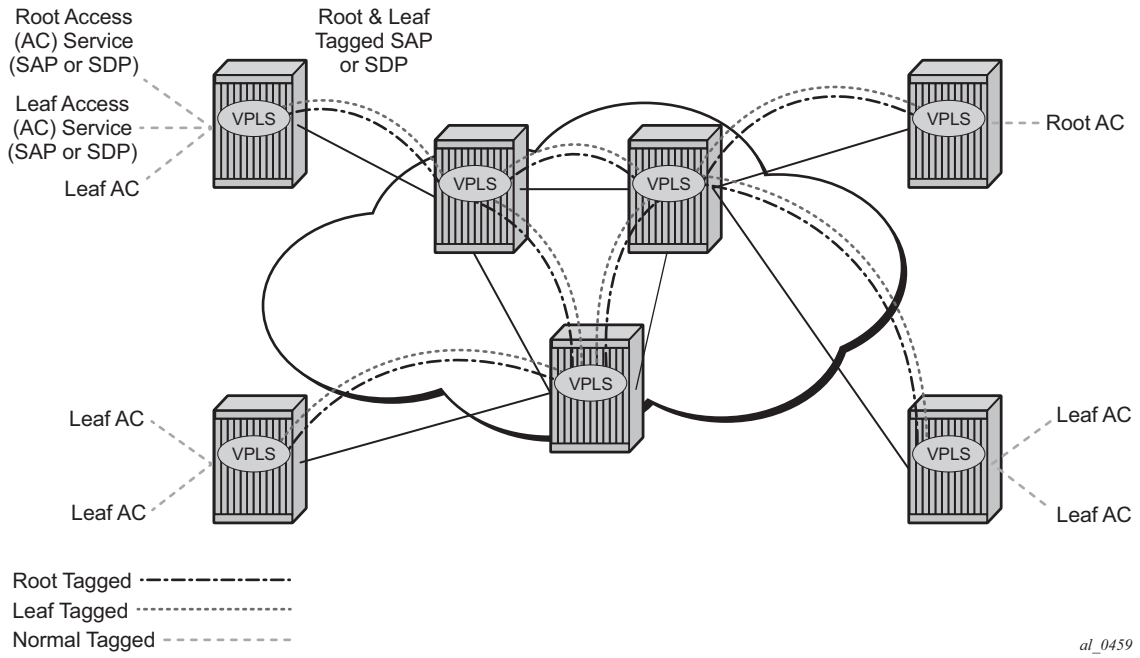

3.4.5.2 Leaf-ac and Root-ac SAPs

Figure 91 illustrates the terminology used for E-Tree in IETF Draft *draft-ietf-l2vpn-vpls-pe-etree* and a mapping to SR OS terms.

An Ethernet service access SAP is characterized as either a leaf-ac or a root-ac for a VPLS E-Tree service. As far as SR OS is concerned, these are normal SAPs with either no tag (Null)/ priority tag or dot1q or QinQ encapsulation on the frame. Functionally, a root-ac is a normal SAP and does not need to be differentiated from the regular SAPs except that it will be associated with a root behavior in a VPLS E-Tree.

Leaf-ac SAPs have restrictions; for example, a SAP is configured for a leaf-ac can never send frames to other leaf-ac directly (local) or through a remote node. Leaf-ac SAPs on the same VPLS instance behave as if they are part of a split horizon group (SHG) locally. Leaf-ac SAPs that are on other nodes need to have the traffic marked as originating “from a Leaf” in the context of the VPLS service when carried on PWs and SAPs with tags (VLANs).

Root-ac SAPs on the same VPLS can talk to any root-ac or leaf-ac.

Figure 91 **Mapping PE Model to VPLS Service**

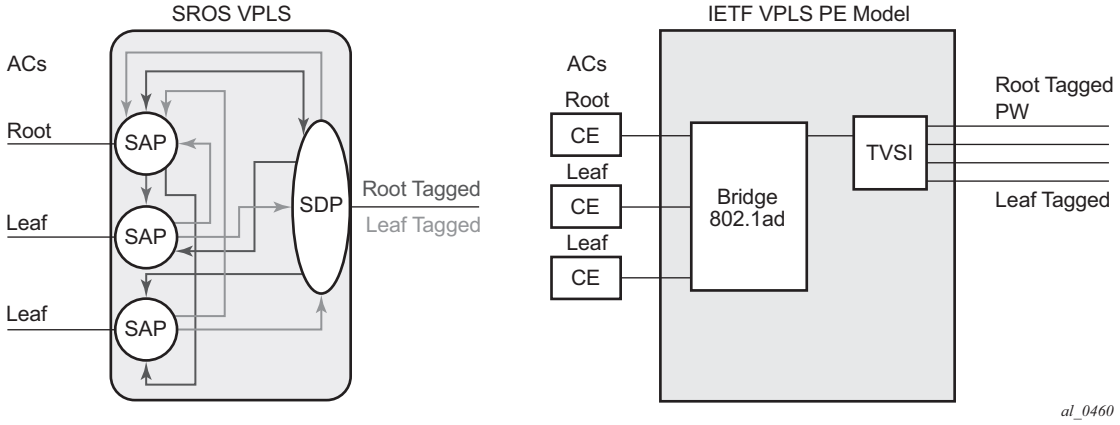

3.4.5.3 Leaf-ac and Root-ac SDP Binds

Untagged SDP binds for access can also be designated as root-ac or leaf-ac. This type of E-Tree interface is required for devices that do not support E-Tree, such as the 7210 SAS, to enable them to be connected with pseudowires. Such devices are root or leaf only and do not require having a tagged frame with a root or leaf indication.

3.4.5.4 Root-leaf-tag SAPs

Support on root-leaf-tag SAPs requires that the outer VID is overloaded to indicate root and leaf. To support the SR service model for a SAP the ability to send and receive 2 different tags on a single SAP has been added. Figure 92 illustrates the behavior when a root-ac and leaf-ac exchange traffic over a root-leaf-tag SAP. Although the figure shows two SAPs connecting VPLS instances 1 and 2, the CLI will show a single SAP with the format:

```
sap 2/1/1:25 root-leaf-tag leaf-tag 26 create
```

Figure 92 **Leaf and Root Tagging dot1q**

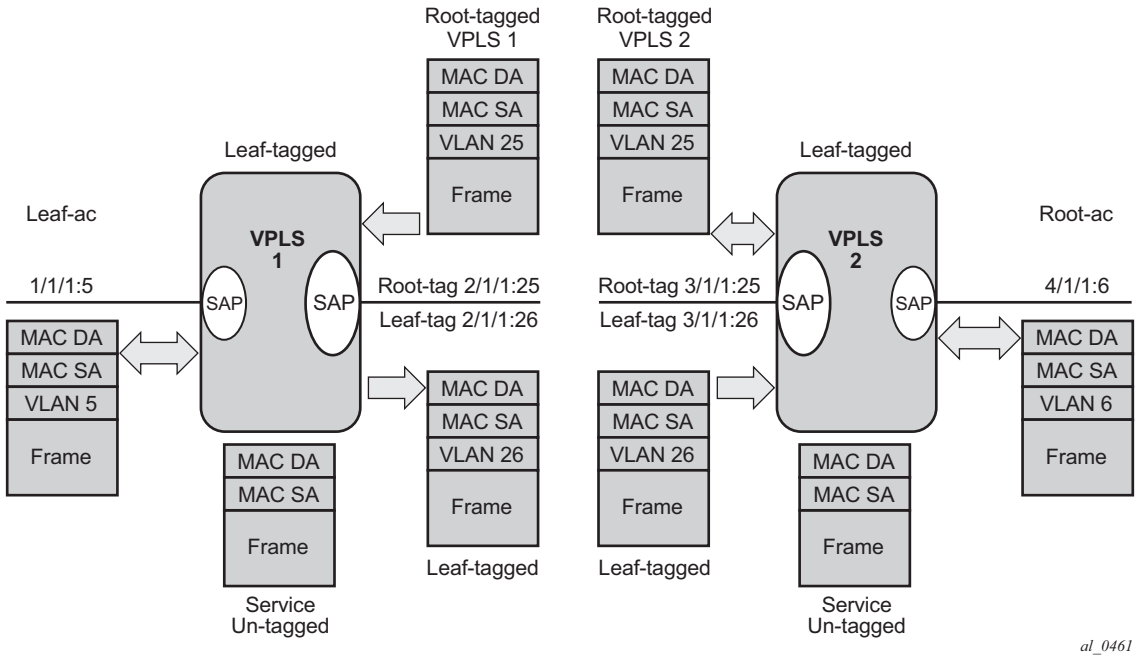

The root-leaf-tag SAP performs all of the operations for egress and ingress traffic for both tags (root and leaf):

- When receiving a frame, the outer tag VID will be compared against the configured root or leaf VIDs and the frame forwarded accordingly.
- When transmitting, the system will add a root VLAN (in the outer tag) on frames with an internal indication of Root, and a leaf VLAN on frames with an internal indication of Leaf.

3.4.5.5 Root-leaf-tag SDP Binds

Typically, in a VPLS environment over MPLS, mesh and spoke-SDP binds interconnect the local VPLS instances to remote PEs. To support VPLS E-Tree the root and leaf traffic is sent over the SDP bind using a fixed VLAN tag value. The SR OS implementation uses a fixed VLAN ID 1 for root and fixed VLAN ID 2 for leaf. The root and leaf tags are a considered a global value and signaling is not supported. The vc-type on root-leaf-tag SDP binds must be VLAN. The vlan-vc-tag command will be blocked in root-leaf-tag SDP-binds.

Figure 93 illustrates the behavior when leaf-ac or root-ac interfaces exchange traffic over a root-leaf-tag SDP-binding.

Figure 93 **Leaf and Root Tagging PW**

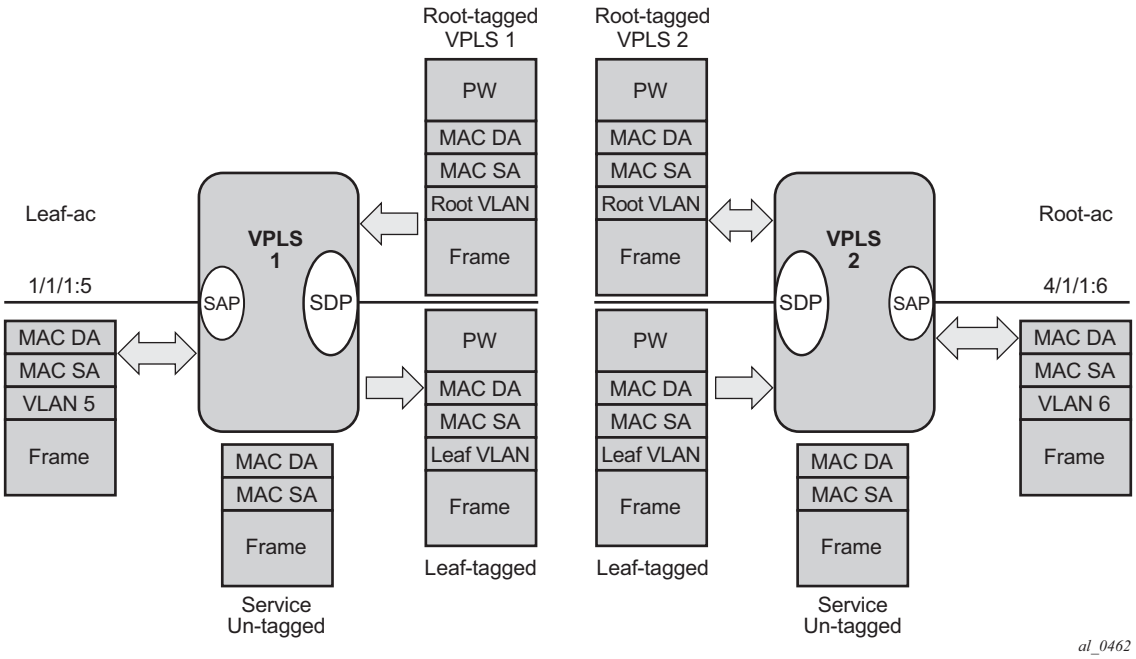

3.4.5.6 Interaction between VPLS E-Tree Services and Other Features

As a general rule, any CPM-generated traffic is always root traffic (STP, OAM, and so on) and any received control plane frame is marked with a root/leaf indication based on which E-Tree interface it arrived at. Some other particular feature interactions are described below:

- ETH-CFM and E-Tree — ETH-CFM allows the operator to verify connectivity between the various endpoint of the service as well as execute troubleshooting and performance gathering functions. Continuity Checking, ETH-CC, is a method by which endpoints are configured and messages are passed between them at regular configured intervals. When CCM Enabled MEPs are configured all MEPs in the same maintenance association, the grouping typically along the service lines, must know about every other endpoint in the service. This is the main principle behind continuity verification (all endpoints in communication). Although the maintenance points configured within the E-Tree service adhere to the forwarding rules of the Leaf and the Root, local population of the MEP database used by the ETH-CFM function may make it appear as the forwarding plane is broken when it is not. All MEPs that are locally configured within a service will automatically be added to the local MEP database. However, because of the Leaf and Root forwarding rules not all of these MEPs can receive the required peer CCM message to avoid CCM Defect conditions. It is suggested, when deploying CCM enabled MEPs in an E-Tree configuration, these CCM-enabled MEPs are configured on Root entities. If Leaf access requires CCM verification then down MEPs in separate maintenance associations should be configured. This consideration is only for operators who wish to deploy CCM in E-Tree environments. No other ETH-CFM tools query or utilize this database.
- Legacy OAM commands (cpe-ping, mac-ping, mac-trace, mac-populate and mac-purge) are not supported in E-Tree service contexts. Although some configuration may result in normal behavior for some commands not all commands or configurations will yield the expected results. Standards based ETH-CFM tools should be used in place of the proprietary legacy OAM command set.
- IGMP and PIM snooping for IPv4 work on VPLS E-Tree services. Routers should use root-ac interfaces so that the multicast traffic can be delivered properly.
- xSTP is supported in VPLS E-Tree services; however, when configuring STP in VPLS E-Tree services the following considerations apply:
  - STP must be carefully used so that STP does not block undesired objects.
  - xSTP is not aware of the leaf-to-leaf topology, e.g. for leaf-to-leaf traffic, even if there is no loop in the forwarding plane, xSTP may block leaf-ac SAPs or SDP binds.
  - Since xSTP is not aware of the root-leaf topology either, root ports might end up blocked before leaf interfaces.
  - When xSTP is used as a access redundancy mechanism, Nokia recommends that the dual-homed device is connected to the same type of E-Tree AC, to avoid unexpected forwarding behaviors when xSTP converges.

- Redundancy mechanisms such as MC-LAG, SDP bind end-points, or BGP-MH are fully supported on VPLS E-Tree services. However, eth-tunnel SAPs or eth-ring control SAPs are not supported on VPLS E-Tree services.

3.5 Configuring a VPLS Service with CLI

This section provides information to configure VPLS services using the command line interface.

3.5.1 Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPLS service:

- Customer ID (refer to the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide* for more information)
- For a local service, configure two SAPs, specifying local access ports and encapsulation values.
- For a distributed service, configure a SAP and an SDP for each far-end node.

The following example shows a sample configuration of a local VPLS service on ALA-1.

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
...
        vpls 9001 customer 6 create
            description "Local VPLS"
            stp
                shutdown
            exit
            sap 1/2/2:0 create
                description "SAP for local service"
            exit
            sap 1/1/5:0 create
                description "SAP for local service"
            exit
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

The following example shows a sample configuration of a distributed VPLS service between ALA-1, ALA-2, and ALA-3.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            shutdown
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
```

```
                shutdown
            exit
            sap 1/1/5:16 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 7:750 create
            exit
        exit
...
-----------------------------------------------
*A:ALA-1>config>service#

*A:ALA-2>config>service# info
-----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/5:16 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 8:750 create
            exit
            no shutdown
        exit
...
-----------------------------------------------
*A:ALA-2>config>service#

*A:ALA-3>config>service# info
-----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/3:33 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 8:750 create
            exit
            no shutdown
        exit
...
-----------------------------------------------
*A:ALA-3>config>service#
```

3.5.2 Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure both local and distributed VPLS services and provides the CLI commands.

For VPLS services:

Step 1. Associate VPLS service with a customer ID

Step 2. Define SAPs:

- Select node(s) and port(s)
- Optional — Select QoS policies other than the default (configured in config>qos context)
- Optional — Select filter policies (configured in config>filter context)
- Optional — Select accounting policy (configured in config>log context)

Step 3. Associate SDPs for (distributed services)

Step 4. Modify STP default parameters (optional) (see VPLS and Spanning Tree Protocol)

Step 5. Enable service

3.5.3 Configuring VPLS Components

Use the CLI syntax displayed in the following sections to configure VPLS components.

3.5.3.1 Creating a VPLS Service

Use the following CLI syntax to create a VPLS service:

CLI Syntax:

```
config>service# vpls service-id [customer customer-id]
  [vpn vpn-id] [m-vpls] [b-vpls | i-vpls] [create]
description description-string
no shutdown
```

The following example shows a VPLS configuration:

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
```

```
            def-mesh-vc-id 750
            stp
                shutdown
            exit
        exit
...
-----------------------------------------------
*A:ALA-1>config>service>vpls#
```

3.5.3.2 Enabling Multiple MAC Registration Protocol (MMRP)

Once MMRP is enabled in the B-VPLS, it advertises the presence of the I-VPLS instances associated with this B-VPLS.

The following example shows a configuration with MMRP enabled.

```
*A:PE-B>config>service# info
----------------------------------------------
        vpls 11 customer 1 vpn 11 i-vpls create
            backbone-vpls 100:11
            exit
            stp
                shutdown
            exit
            sap 1/5/1:11 create
            exit
            sap 1/5/1:12 create
            exit
            no shutdown
        exit
        vpls 100 customer 1 vpn 100 b-vpls create
            service-mtu 2000
            stp
                shutdown
            exit
            mrp
                flood-time 10
                no shutdown
            exit
            sap 1/5/1:100 create
            exit
            spoke-sdp 3101:100 create
            exit
            spoke-sdp 3201:100 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:PE-B>config>service#
```

Since I-VPLS 11 is associated with B-VPLS 100, MMRP advertises the group B-MAC 01:1e:83:00:00:0b) associated with I-VPLS 11 through a declaration on all the B-SAPs and B-SDPs. If the remote node also declares an I-VPLS 11 associated to its B-VPLS 10, then this results in a registration for the group B-MAC. This also creates the MMRP multicast tree (MFIB entries). In this case, sdp 3201:100 is connected to a remote node that declares the group B-MAC.

The following show commands display the current MMRP information for this scenario:

```
*A:PE-C# show service id 100 mrp
-------------------------------------------------------------------------------
MRP Information
-------------------------------------------------------------------------------
Admin State        : Up                 Failed Register Cnt: 0
Max Attributes     : 1023               Attribute Count    : 1
Attr High Watermark: 95%                Attr Low Watermark : 90%
Flood Time         : 10
-------------------------------------------------------------------------------
*A:PE-C# show service id 100 mmrp mac
-------------------------------------------------------------------------------
SAP/SDP                                 MAC Address       Registered  Declared
-------------------------------------------------------------------------------
sap:1/5/1:100                           01:1e:83:00:00:0b No          Yes
sdp:3101:100                            01:1e:83:00:00:0b No          Yes
sdp:3201:100                            01:1e:83:00:00:0b Yes         Yes
-------------------------------------------------------------------------------


*A:PE-C# show service id 100 sdp 3201:100 mrp
-------------------------------------------------------------------------------
Sdp Id 3201:100 MRP Information
-------------------------------------------------------------------------------
Join Time          : 0.2 secs                Leave Time        : 3.0 secs
Leave All Time     : 10.0 secs               Periodic Time     : 1.0 secs
Periodic Enabled   : false
Rx Pdus            : 7                       Tx Pdus           : 23
Dropped Pdus       : 0
Rx New Event       : 0                       Rx Join-In Event  : 6
Rx In Event        : 0                       Rx Join Empty Evt : 1
Rx Empty Event     : 0                       Rx Leave Event    : 0
Tx New Event       : 0                       Tx Join-In Event  : 4
Tx In Event        : 0                       Tx Join Empty Evt : 19
Tx Empty Event     : 0                       Tx Leave Event    : 0
-------------------------------------------------------------------------------
SDP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
01:1e:83:00:00:0b Yes               Yes
-------------------------------------------------------------------------------
Number of MACs=1 Registered=1 Declared=1
-------------------------------------------------------------------------------
*A:PE-C#


*A:PE-C#  show service id 100 mfib
```

```
===============================================================================
Multicast FIB, Service 100
===============================================================================
Source Address  Group Address         Sap/Sdp Id              Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               01:1E:83:00:00:0B     sdp:3201:100            Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
*A:PE-C#
```

3.5.3.2.1 Enabling MAC Move

The **mac-move** feature is useful to protect against undetected loops in your VPLS topology as well as the presence of duplicate MACs in a VPLS service. For example, if two clients in the VPLS have the same MAC address, the VPLS will experience a high re-learn rate for the MAC and will shut down the SAP or spoke-SDP when the threshold is exceeded.

Use the following CLI syntax to configure **mac-move** parameters.

CLI Syntax:

```
config>service# vpls service-id [customer customer-id]
  [vpn vpn-id] [m-vpls]
mac-move
     primary-ports
          spoke-sdp
          cumulative-factor
     exit
     secondary-ports
          spoke-sdp
          sap
     exit
     move-frequency frequency
     retry-timeout timeout
     no shutdown
```

The following example shows a **mac-move** configuration:

```
*A:ALA-2009>config>service>vpls>mac-move# show service id 500 mac-move
===============================================================================
Service Mac Move Information
===============================================================================
Service Id        : 500                 Mac Move          : Enabled
Primary Factor    : 4                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
-------------------------------------------------------------------------------
SAP Mac Move Information: 2/1/3:501
-------------------------------------------------------------------------------
Admin State       : Up                  Oper State        : Down
```

```
Flags              : RelearnLimitExceeded
Time to come up    : 1 seconds           Retries Left      : 1
Mac Move           : Blockable           Blockable Level   : Tertiary
-------------------------------------------------------------------------------
SAP Mac Move Information: 2/1/3:502
-------------------------------------------------------------------------------
Admin State        : Up                  Oper State        : Up
Flags              : None
Time to RetryReset: 267 seconds          Retries Left      : none
Mac Move           : Blockable           Blockable Level   : Tertiary
-------------------------------------------------------------------------------
SDP Mac Move Information: 21:501
-------------------------------------------------------------------------------
Admin State        : Up                  Oper State        : Up
Flags              : None
Time to RetryReset: never                Retries Left      : 3
Mac Move           : Blockable           Blockable Level   : Secondary
-------------------------------------------------------------------------------
SDP Mac Move Information: 21:502
-------------------------------------------------------------------------------
Admin State        : Up                  Oper State        : Down
Flags              : RelearnLimitExceeded
Time to come up    : never               Retries Left      : none
Mac Move           : Blockable           Blockable Level   : Tertiary
===============================================================================
*A:*A:ALA-2009>config>service>vpls>mac-move#
```

3.5.3.2.2 Configuring STP Bridge Parameters in a VPLS

Modifying some of the Spanning Tree Protocol parameters allows the operator to balance STP between resiliency and speed of convergence extremes. Modifying particular parameters, mentioned below, must be done in the constraints of the following two formulae:

2 x (Bridge_Forward_Delay - 1.0 seconds) >= Bridge_Max_Age
Bridge_Max_Age >= 2 x (Bridge_Hello0_Time + 1.0 seconds)

The following STP parameters can be modified at VPLS level:

- Bridge STP Admin State
- Mode
- Bridge Priority
- Max Age
- Forward Delay
- Hello Time
- MST Instances
- MST Max Hops
- MST Name

- MST Revision

STP always uses the locally configured values for the first three parameters (Admin State, Mode and Priority).

For the parameters Max Age, Forward Delay, Hello Time and Hold Count, the locally configured values are only used when this bridge has been elected root bridge in the STP domain, otherwise the values received from the root bridge are used. The exception to this rule is: when STP is running in RSTP mode, the Hello Time is always taken from the locally configured parameter. The other parameters are only used when running mode MSTP.

Bridge STP Admin State

The administrative state of STP at the VPLS level is controlled by the shutdown command.

When STP on the VPLS is administratively disabled, any BPDUs are forwarded transparently through the 7450 ESS, 7750 SR, or 7950 XRS. When STP on the VPLS is administratively enabled, but the administrative state of a SAP or spoke-SDP is down, BPDUs received on such a SAP or spoke-SDP are discarded.

CLI Syntax:
```
config>service>vpls service-id# stp
no shutdown
```

Mode

To be compatible with the different iterations of the IEEE 802.1D standard, the 7450 ESS, 7750 SR, and 7950 XRS support several variants of the Spanning Tree protocol:

- **rstp** — Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode.
- **dot1w** — Compliant with IEEE 802.1w.
- **comp-dot1w** — Operation as in RSTP but backwards compatible with IEEE 802.1w
  (this mode was introduced for interoperability with some MTU types).
- **mstp** — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS.

- **pmstp** — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D3.0-04/2005 but with some changes to make it backwards compatible to 802.1Q 2003 edition and IEEE 802.1w.

See section Spanning Tree Operating Modes for details on these modes.

CLI Syntax: `config>service>vpls service-id# stp`
`mode {rstp | comp-dot1w | dot1w | mstp}`

Default: rstp

Bridge Priority

The **bridge-priority** command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. When running MSTP, this is the bridge priority used for the CIST.

All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

CLI Syntax: `config>service>vpls service-id# stp`
`priority bridge-priority`

Range: 1 to 65535

Default: 32768

Restore Default: no priority

Max Age

The **max-age** command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge will take the message_age value from BPDUs received on their root port and increment this value by 1. The message_age thus reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges by the BPDUs.
The default value of **max-age** is 20. This parameter can be modified within a range of 6 to 40, limited by the standard STP parameter interaction formulae.

CLI Syntax: `config>service>vpls service-id# stp`
`max-age max-info-age`

Range: 6 to 40 seconds

Default: 20 seconds

Restore Default: no max-age

Forward Delay

RSTP, as defined in the IEEE 802.1D-2004 standards, will normally transition to the forwarding state by a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (e.g. on shared links, see below), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two Ethernet bridges (for example, a shared 10/100BaseT segment). The port-type command is used to configure a link as point-to-point or shared (see section SAP Link Type).

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP or spoke-SDP spends in the discarding and learning states when transitioning to the forwarding state. The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in rstp mode, but only when the SAP or spoke-SDP has not fallen back to legacy STP operation, the value configured by the **hello-time** command is used;
- in all other situations, the value configured by the **forward-delay** command is used.

CLI Syntax: `config>service>vpls service-id# stp`
`forward-delay seconds`

Range: 4 to 30 seconds

Default: 15 seconds

Restore Default: no forward-delay

Hello Time

The **hello-time** command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The *seconds* parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time value can also be used to calculate the bridge forward delay, see Forward Delay.

CLI Syntax: `config>service>vpls service-id# stp`
`hello-time hello-time`

Range: 1 to 10 seconds

Default: 2 seconds

Restore Default: no hello-time

Hold Count

The **hold-count** command configures the peak number of BPDUs that can be transmitted in a period of one second.

CLI Syntax: `config>service>vpls service-id# stp`
`hold-count count-value`

Range: 1 to 10

Default: 6

Restore Default: no hold-count

MST Instances

You can create up to 15 MST-instances. They can range from 1 to 4094. By changing path-cost and priorities, you can make sure that each instance will form it's own tree within the region, therefore making sure different VLANs follow different paths.

You can assign non-overlapping VLAN ranges to each instance. VLANs that are not assigned to an instance are implicitly assumed to be in instance 0, which is also called the CIST. This CIST cannot be deleted or created.

The parameter that can be defined per instance are mst-priority and vlan-range.

- mst-priority — The bridge-priority for this specific mst-instance. It follows the same rules as bridge-priority. For the CIST, the bridge-priority is used.
- vlan-range — The VLANs are mapped to this specific mst-instance. If no VLAN-ranges are defined in any mst-instances, then all VLANs are mapped to the CIST.

MST Max Hops

The mst-max-hops command defines the maximum number of hops the BPDU can traverse inside the region. Outside the region max-age is used.

MST Name

The MST name defines the name that the operator gives to a region. Together with MST revision and the VLAN to MST-instance mapping, it forms the MST configuration identifier. Two bridges that have the same MST configuration identifier form a region if they exchange BPDUs.

MST Revision

The MST revision together with MST-name and VLAN to MST-instance mapping define the MST configuration identifier. Two bridges that have the same MST configuration identifier form a region if they exchange BPDUs.

3.5.3.3 Configuring GSMP Parameters

The following parameters must be configured in order for GSMP to function:

- One or more GSMP sessions
- One or more ANCP policies
- For basic subscriber management only, ANCP static maps

- For enhanced subscriber management only, associate subscriber profiles with ANCP policies.

Use the following CLI syntax to configure GSMP parameters.

CLI Syntax:

```
config>service>vpls# gsmp
group name [create]
     ancp
          dynamic-topology-discover
          oam
     description description-string
     hold-multiplier multiplier
     keepalive seconds
     neighbor ip-address [create]
          description v
          local-address ip-address
          priority-marking dscp dscp-name
          priority-marking prec ip-prec-value
          [no] shutdown
     [no] shutdown
[no] shutdown
```

This example shows a GSMP group configuration.

```
A:ALA-48>config>service>vpls>gsmp# info
----------------------------------------------
               group "group1" create
                   description "test group config"
                   neighbor 10.10.10.104 create
                       description "neighbor1 config"
                       local-address 10.10.10.103
                       no shutdown
                   exit
                   no shutdown
               exit
               no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls>gsmp#
```

3.5.3.4 Configuring a VPLS SAP

A default QoS policy is applied to each ingress and egress SAP. Additional QoS policies can be configured in the **config>qos** context. There are no default filter policies. Filter policies are configured in the **config>filter** context and must be explicitly applied to a SAP.
Use the following CLI syntax to create:

- Local VPLS SAPs

- Distributed VPLS SAPs

3.5.3.4.1 Local VPLS SAPs

To configure a local VPLS service, enter the **sap** *sap-id* command twice with different port IDs in the same service configuration.

The following example shows a local VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vpls 90001 customer 6 create
            description "Local VPLS"
            stp
                shutdown
            exit
            sap 1/2/2:0 create
                description "SAP for local service"
            exit
            sap 1/1/5:0 create
                description "SAP for local service"
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#
*A:ALA-1>config>service# info
----------------------------------------------
       vpls 1150 customer 1 create
            fdb-table-size 1000
            fdb-table-low-wmark 5
            fdb-table-high-wmark 80
            local-age 60
            stp
                shutdown
            exit
            sap 1/1/1:1155 create
            exit
            sap 1/1/2:1150 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#
```

3.5.3.4.2 Distributed VPLS SAPs

To configure a distributed VPLS service, you must configure service entities on originating and far-end nodes. You must use the same service ID on all ends (for example, create a VPLS service ID 9000 on ALA-1, ALA-2, and ALA-3). A distributed VPLS consists of a SAP on each participating node and an SDP bound to each participating node.

For SDP configuration information, refer to the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide*. For SDP binding information, see Configuring SDP Bindings.

The following example shows a configuration of VPLS SAPs configured for ALA-1, ALA-2, and ALA-3.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 vpn 750 create
            description "Distributed VPLS services."
def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/2/5:0 create
                description "VPLS SAP"
                multi-service-site "West"
            exit
        exit
...
----------------------------------------------
*A:ALA-1>config>service#


*A:ALA-2>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 vpn 750 create
            description "Distributed VPLS services."
def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/2:22 create
                description "VPLS SAP"
                multi-service-site "West"
            exit
        exit
...
----------------------------------------------
*A:ALA-2>config>service#


*A:ALA-3>config>service# info
----------------------------------------------
```

```
...
        vpls 9000 customer 6 vpn 750 create
            description "Distributed VPLS services."
def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/3:33 create
                description "VPLS SAP"
                multi-service-site "West"
            exit
        exit
...
----------------------------------------------
*A:ALA-3>config>service#
```

3.5.3.4.3 Configuring SAP-Specific STP Parameters

When a VPLS has STP enabled, each SAP within the VPLS has STP enabled by default. The operation of STP on each SAP is governed by:

- SAP STP Administrative State
- SAP Virtual Port Number
- SAP Priority
- SAP Path Cost
- SAP Edge Port
- SAP Auto Edge
- SAP Link Type

SAP STP Administrative State

The administrative state of STP within a SAP controls how BPDUs are transmitted and handled when received. The allowable states are:

- SAP Admin Up

  The default administrative state is up for STP on a SAP. BPDUs are handled in the normal STP manner on a SAP that is administratively up.
- SAP Admin Down

  An administratively down state allows a service provider to prevent a SAP from becoming operationally blocked. BPDUs will not originate out the SAP toward the customer.

If STP is enabled on VPLS level, but disabled on the SAP, received BPDUs are discarded. Discarding the incoming BPDUs allows STP to continue to operate normally within the VPLS service while ignoring the down SAP. The specified SAP will always be in an operationally forwarding state.

Note: The administratively down state allows a loop to form within the VPLS.

CLI Syntax: `config>service>vpls>sap>stp#`
`[no] shutdown`

Range: shutdown or no shutdown

Default: no shutdown (SAP admin up)

SAP Virtual Port Number

The virtual port number uniquely identifies a SAP within configuration BPDUs. The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with it's own virtual port number that is unique to every other SAP defined on the VPLS. The virtual port number is assigned at the time that the SAP is added to the VPLS.

Since the order in which SAPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

CLI Syntax: `config>service>vpls>sap# stp`
`port-num` *`number`*

Range: 1 to 2047

Default: (automatically generated)

Restore Default: no port-num

SAP Priority

SAP priority allows a configurable “tie breaking” parameter to be associated with a SAP. When configuration BPDUs are being received, the configured SAP priority will be used in some circumstances to determine whether a SAP will be designated or blocked. These are the values used for CIST when running MSTP for the 7450 ESS or 7750 SR.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP within the STP instance. See SAP Virtual Port Number for details on the virtual port number.

STP computes the actual SAP priority by taking the configured priority value and masking out the lower four bits. The result is the value that is stored in the SAP priority parameter. For example, if a value of 0 was entered, masking out the lower 4 bits would result in a parameter value of 0. If a value of 255 was entered, the result would be 240.

The default value for SAP priority is 128. This parameter can be modified within a range of 0 to 255, 0 being the highest priority. Masking causes the values actually stored and displayed to be 0 to 240, in increments of 16.

CLI Syntax:
```
config>service>vpls>sap>stp#
priority stp-priority
```

Range: 0 to 255 (240 largest value, in increments of 16)

Default: 128

Restore Default: no priority

SAP Path Cost

The SAP path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP. When BPDUs are sent out other egress SAPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs are controlled by complex queuing dynamics, in the 7450 ESS, 7750 SR, and 7950 XRS the STP path cost is a purely static configuration.

The default value for SAP path cost is 10. This parameter can be modified within a range of 1 to 65535, 1 being the lowest cost.

CLI Syntax: `config>service>vpls>sap>stp#`
`path-cost sap-path-cost`

Range: 1 to 200000000

Default: 10

Restore Default: no path-cost

SAP Edge Port

The SAP edge-port command is used to reduce the time it takes a SAP to reach the forwarding state when the SAP is on the edge of the network, and therefore has no further STP bridge to handshake with.

The edge-port command is used to initialize the internal OPER_EDGE variable. At any time, when OPER_EDGE is false on a SAP, the normal mechanisms are used to transition to the forwarding state (see Forward Delay). When OPER_EDGE is true, STP assumes that the remote end agrees to transition to the forwarding state without actually receiving a BPDU with an agreement flag set.

The OPER_EDGE variable will dynamically be set to false if the SAP receives BPDUs (the configured edge-port value does not change). The OPER_EDGE variable will dynamically be set to true if auto-edge is enabled and STP concludes there is no bridge behind the SAP.

When STP on the SAP is administratively disabled and re-enabled, the OPER_EDGE is re-initialized to the value configured for edge-port.

Valid values for SAP edge-port are enabled and disabled with disabled being the default.

CLI Syntax: `config>service>vpls>sap>stp#`
`[no] edge-port`

Default: no edge-port

SAP Auto Edge

The SAP **edge-port** command is used to instruct STP to dynamically decide whether the SAP is connected to another bridge.

If auto-edge is enabled, and STP concludes there is no bridge behind the SAP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see SAP Edge Port).

Valid values for SAP auto-edge are enabled and disabled with enabled being the default.

CLI Syntax:
```
config>service>vpls>sap>stp#
[no] auto-edge
```

Default: auto-edge

SAP Link Type

The SAP **link-type** parameter instructs STP on the maximum number of bridges behind this SAP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected by a shared media, their SAPs should all be configured as shared, and timer-based transitions are used.

Valid values for SAP link-type are shared and pt-pt with pt-pt being the default.

CLI Syntax:
```
config>service>vpls>sap>stp#
link-type {pt-pt | shared}
```

Default: link-type pt-pt

Restore Default: no link-type

3.5.3.4.4 STP SAP Operational States

The operational state of STP within a SAP controls how BPDUs are transmitted and handled when received. Defined states are:

- Operationally Disabled
- Operationally Discarding
- Operationally Learning

- Operationally Forwarding

Operationally Disabled

Operationally disabled is the normal operational state for STP on a SAP in a VPLS that has any of the following conditions:

- VPLS state administratively down
- SAP state administratively down
- SAP state operationally down

If the SAP enters the operationally up state with the STP administratively up and the SAP STP state is up, the SAP will transition to the STP SAP discarding state.

When, during normal operation, the router detects a downstream loop behind a SAP or spoke-SDP, BPDUs can be received at a very high rate. To recover from this situation, STP will transition the SAP to disabled state for the configured forward-delay duration.

Operationally Discarding

A SAP in the discarding state only receives and sends BPDUs, building the local correct STP state for each SAP while not forwarding actual user traffic. The duration of the discarding state is explained in section Forward Delay.

Note: In previous versions of the STP standard, the discarding state was called a blocked state.

Operationally Learning

The learning state allows population of the MAC forwarding table before entering the forwarding state. In this state, no user traffic is forwarded.

Operationally Forwarding

Configuration BPDUs are sent out a SAP in the forwarding state. Layer 2 frames received on the SAP are source learned and destination forwarded according to the FDB. Layer 2 frames received on other forwarding interfaces and destined for the SAP are also forwarded.

SAP BPDU Encapsulation State

IEEE 802.1d (referred as Dot1d) and Cisco's per VLAN Spanning Tree (PVST) BPDU encapsulations are supported on a per SAP basis for the 7450 ESS and 7750 SR. STP is associated with a VPLS service like PVST is associated per VLAN. The main difference resides in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDU.

The following table shows differences between Dot1d and PVST Ethernet BPDU encapsulations based on the interface encap-type field:

Each SAP has a Read-Only operational state that shows which BPDU encapsulation is currently active on the SAP. The states are:

- Dot1d — This state specifies that the switch is currently sending IEEE 802.1d standard BPDUs. The BPDUs are tagged or non-tagged based on the encapsulation type of the egress interface and the encapsulation value defined in the SAP. A SAP defined on an interface with encapsulation type dot1q continues in the dot1d BPDU encapsulation state until a PVST encapsulated BPDU is received in which case, the SAP will convert to the PVST encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged if the interface encapsulation type is defined as dot1q. PVST BPDUs will be silently discarded if received when the SAP is on an interface defined with encapsulation type null.
- PVST — This state specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The SAP continues in the PVST BPDU encapsulation state until a dot1d encapsulated BPDU is received, in which case, the SAP reverts to the dot1d encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged with the encapsulation value defined for the SAP. PVST BPDUs are silently discarded if received when the SAP is on an interface defined with a null encapsulation type.

Dot1d is the initial and only SAP BPDU encapsulation state for SAPs defined on Ethernet interface with encapsulation type set to null.

Each transition between encapsulation types optionally generates an alarm that can be logged and optionally transmitted as an SNMP trap on the 7450 ESS or 7750 SR.

3.5.3.4.5 Configuring VPLS SAPs with Split Horizon

To configure a VPLS service with a split horizon group, add the **split-horizon-group** parameter when creating the SAP. Traffic arriving on a SAP within a split horizon group will not be copied to other SAPs in the same split horizon group.

The following example shows a VPLS configuration with split horizon enabled:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
    vpls 800 customer 6001 vpn 700 create
        description "VPLS with split horizon for DSL"
        stp
            shutdown
        exit
        sap 1/1/3:100 split-horizon-group DSL-group1 create
            description "SAP for residential bridging"
        exit
        sap 1/1/3:200 split-horizon-group DSL-group1 create
            description "SAP for residential bridging"
        exit
        split-horizon-group DSL-group1
            description "Split horizon group for DSL"
        exit
        no shutdown
    exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

3.5.3.4.6 Configuring MAC Learning Protection

To configure MAC learning protection, configure split horizon, MAC protection, and SAP parameters on the 7450 ESS or 7750 SR.

The following example shows a VPLS configuration with split horizon enabled:

```
A:ALA-48>config>service>vpls# info
----------------------------------------------
            description "IMA VPLS"
            split-horizon-group "DSL-group1" create
                restrict-protected-src
                restrict-unprotected-dst
            exit
            mac-protect
                 mac ff:ff:ff:ff:ff:ff
            exit
            sap 1/1/9:0 create
                ingress
                    scheduler-policy "SLA1"
                    qos 100 shared-queuing
```

```
                exit
                egress
                    scheduler-policy "SLA1"
                    filter ip 10
                exit
                restrict-protected-src
                arp-reply-agent
                host-connectivity-verify source-ip 143.144.145.1
            exit
...
-------------------------------------------
A:ALA-48>config>service>vpls#
```

3.5.3.5 Configuring SAP Subscriber Management Parameters

Use the following CLI syntax to configure subscriber management parameters on a VPLS service SAP on the 7450 ESS and 7750 SR. The policies and profiles that are referenced in the **def-sla-profile**, **def-sub-profile**, **non-sub-traffic**, and **sub-ident-policy** commands must already be configured in the **config>subscr-mgmt** context.

CLI Syntax:

```
config>service>vpls service-id
sap sap-id [split-horizon-group group-name]
     sub-sla-mgmt
          def-sla-profile default-sla-profile-name
          def-sub-profile default-subscriber-profile-
             name
          mac-da-hashing
          multi-sub-sap [number-of-sub]
          no shutdown
          single-sub-parameters
               non-sub-traffic sub-profile sub-profile-
                  name sla-profile sla-profile-name
                  [subscriber sub-ident-string]
               profiled-traffic-only
          sub-ident-policy sub-ident-policy-name
```

The following example shows a subscriber management configuration:

```
A:ALA-48>config>service>vpls#
-------------------------------------------
            description "Local VPLS"
            stp
                shutdown
            exit
            sap 1/2/2:0 create
                description "SAP for local service"
                sub-sla-mgmt
                    def-sla-profile "sla-profile1"
                    sub-ident-policy "SubIdent1"
                exit
```

```
            exit
            sap 1/1/5:0 create
                description "SAP for local service"
            exit
            no shutdown
-----------------------------------------------
A:ALA-48>config>service>vpls#
```

3.5.3.6 MSTP Control over Ethernet Tunnels

When MSTP is used to control VLANs, a range of VLAN IDs is normally used to specify the VLANs to be controlled on the 7450 ESS and 7750 SR.

If an Ethernet tunnel SAP is to be controlled by MSTP, the Ethernet tunnel SAP ID needs to be within the VLAN range specified under the mst-instance.

```
vpls 400 customer 1 m-vpls create
            stp
                mode mstp
                mst-instance 111 create
                    vlan-range 1-100
                exit
                mst-name "abc"
                mst-revision 1
                no shutdown
            exit
            sap 1/1/1:0 create // untagged
            exit
            sap eth-tunnel-1 create
            exit
            no shutdown
        exit
        vpls 401 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/1:12 create
            exit
            sap eth-tunnel-1:12 create
                // Ethernet tunnel SAP ID 12 falls within the VLAN
                // range for mst-instance 111
                eth-tunnel
                    path 1 tag 1000
                    path 8 tag 2000
                exit
            exit
            no shutdown
        exit
```

3.5.3.7 Configuring SDP Bindings

VPLS provides scaling and operational advantages. A hierarchical configuration eliminates the need for a full mesh of VCs between participating devices. Hierarchy is achieved by enhancing the base VPLS core mesh of VCs with access VCs (spoke) to form two tiers. Spoke SDPs are generally created between Layer 2 switches and placed at the Multi-Tenant Unit (MTU). The PE routers are placed at the service provider's Point of Presence (POP). Signaling and replication overhead on all devices is considerably reduced.

A spoke-SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received (unless a split horizon group was defined on the spoke-SDP, see section Configuring VPLS Spoke SDPs with Split Horizon).

A spoke-SDP connects a VPLS service between two sites and, in its simplest form, could be a single tunnel LSP. A set of ingress and egress VC labels are exchanged for each VPLS service instance to be transported over this LSP. The PE routers at each end treat this as a virtual spoke connection for the VPLS service in the same way as the PE-MTU connections. This architecture minimizes the signaling overhead and avoids a full mesh of VCs and LSPs between the two metro networks.

A mesh SDP bound to a service is logically treated like a single bridge “port” for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke-SDPs and SAPs) and not transmitted on any mesh SDPs.

A VC-ID can be specified with the SDP-ID. The VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer SRs on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

Figure 94 shows an example of a distributed VPLS service configuration of spoke and mesh SDPs (uni-directional tunnels) between routers and MTUs.

3.5.3.8 Configuring Overrides on Service SAPs

The following output shows a service SAP queue override configuration example.

```
*A:ALA-48>config>service>vpls>sap# info
----------------------------------------------
...
exit
ingress
```

```
        scheduler-policy "SLA1"
        scheduler-override
            scheduler "sched1" create
                parent weight 3 cir-weight 3
            exit
        exit
        policer-control-policy "SLA1-p"
        policer-control-override create
            max-rate 50000
        exit
        qos 100 multipoint-shared
        queue-override
            queue 1 create
                rate 1500000 cir 2000
            exit
        exit
        policer-override
            policer 1 create
                rate 10000
            exit
        exit
    exit
    egress
        scheduler-policy "SLA1"
        policer-control-policy "SLA1-p"
        policer-control-override create
            max-rate 60000
        exit
        qos 100
        queue-override
            queue 1 create
                adaptation-rule pir max cir max
            exit
        exit
        policer-override
            policer 1 create
                mbs 2000 kilobytes
            exit
        exit
        filter ip 10
    exit
    ------------------------------------------------
    *A:ALA-48>config>service>vpls>sap#
```

Figure 94 **SDPs — Uni-Directional Tunnels**

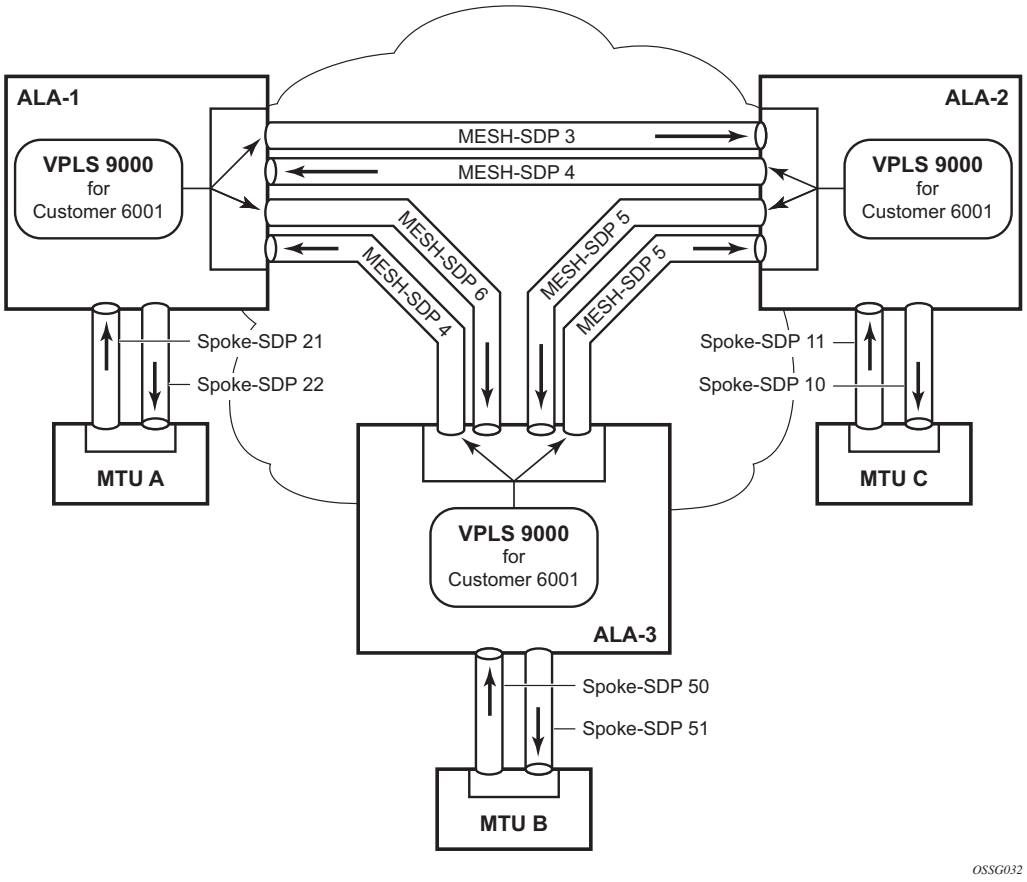

Use the following CLI syntax to create a mesh or spoke-SDP bindings with a distributed VPLS service. SDPs must be configured prior to binding. Refer to the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide* for information about creating SDPs.

Use the following CLI syntax to configure mesh SDP bindings:

CLI Syntax:

```
config>service# vpls service-id
mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
     egress
          filter {ip ip-filter-id|mac mac-filter-id}
          mfib-allowed-mda-destinations
                    mda mda-id
          vc-label egress-vc-label
     ingress
          filter {ip ip-filter-id|mac mac-filter-id}
          vc-label ingress-vc-label
```

```
        no shutdown
        static-mac ieee-address
        vlan-vc-tag 0..4094
```

Use the following CLI syntax to configure spoke-SDP bindings:

CLI Syntax:

```
config>service# vpls service-id
spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}] [split-
 horizon-group group-name]
    egress
        filter {ip ip-filter-id|mac mac-filter-id}
        vc-label egress-vc-label
    ingress
        filter {ip ip-filter-id|mac mac-filter-id}
        vc-label ingress-vc-label
    limit-mac-move[non-blockable]
    vlan-vc-tag 0..4094
    no shutdown
    static-mac ieee-address
    stp
        path-cost stp-path-cost
        priority stp-priority
        no shutdown
    vlan-vc-tag [0..4094]
```

The following examples show SDP binding configurations for ALA-1, ALA-2, and ALA-3 for VPLS service ID 9000 for customer 6:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/2/5:0 create
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 5:750 create
            exit
            mesh-sdp 7:750 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#


*A:ALA-2>config>service# info
----------------------------------------------
```

```
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/2:22 create
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 5:750 create
            exit
            mesh-sdp 7:750 create
            exit
            no shutdown
        exit
-------------------------------------------


*A:ALA-3>config>service# info
-------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/3:33 create
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 5:750 create
            exit
            mesh-sdp 7:750 create
            exit
            no shutdown
        exit
-------------------------------------------
*A:ALA-3>config>service#
```

3.5.3.8.1 Configuring Spoke SDP Specific STP Parameters

When a VPLS has STP enabled, each spoke-SDP within the VPLS has STP enabled by default. The operation of STP on each spoke-SDP is governed by:

- Spoke SDP STP Administrative State
- Spoke SDP Virtual Port Number
- Spoke SDP Priority
- Spoke SDP Path Cost
- Spoke SDP Edge Port

- Spoke SDP Auto Edge
- Spoke SDP Link Type

Spoke SDP STP Administrative State

The administrative state of STP within a spoke-SDP controls how BPDUs are transmitted and handled when received. The allowable states are:

- Spoke SDP Admin Up

  The default administrative state is up for STP on a spoke-SDP. BPDUs are handled in the normal STP manner on a spoke-SDP that is administratively up.
- Spoke SDP Admin Down

  An administratively down state allows a service provider to prevent a spoke-SDP from becoming operationally blocked. BPDUs will not originate out the spoke-SDP toward the customer.

  If STP is enabled on VPLS level, but disabled on the spoke-SDP, received BPDUs are discarded. Discarding the incoming BPDUs allows STP to continue to operate normally within the VPLS service while ignoring the down spoke-SDP. The specified spoke-SDP will always be in an operationally forwarding state.

Note: The administratively down state allows a loop to form within the VPLS.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`[no] shutdown`

Range: shutdown or no shutdown

Default: no shutdown (spoke-SDP admin up)

Spoke SDP Virtual Port Number

The virtual port number uniquely identifies a spoke-SDP within configuration BPDUs. The internal representation of a spoke-SDP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a spoke-SDP and identifies it with it's own virtual port number that is unique to every other spoke-SDP defined on the VPLS. The virtual port number is assigned at the time that the spoke-SDP is added

to the VPLS.
Since the order in which spoke-SDPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

CLI Syntax: `config>service>vpls>spoke-sdp# stp`
`port-num` *`number`*

Range: 1 to 2047

Default: automatically generated

Restore Default: no port-num

Spoke SDP Priority

Spoke SDP priority allows a configurable tie breaking parameter to be associated with a spoke-SDP. When configuration BPDUs are being received, the configured spoke-SDP priority will be used in some circumstances to determine whether a spoke-SDP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the spoke-SDP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a spoke-SDP within the STP instance. See Spoke SDP Virtual Port Number for details on the virtual port number.

STP computes the actual spoke-SDP priority by taking the configured priority value and masking out the lower four bits. The result is the value that is stored in the spoke-SDP priority parameter. For instance, if a value of 0 was entered, masking out the lower 4 bits would result in a parameter value of 0. If a value of 255 was entered, the result would be 240.

The default value for spoke-SDP priority is 128. This parameter can be modified within a range of 0 to 255, 0 being the highest priority. Masking causes the values actually stored and displayed to be 0 to 240, in increments of 16.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`priority stp-priority`

Range: 0 to 255 (240 largest value, in increments of 16)

Default: 128

Restore Default: no priority

Spoke SDP Path Cost

The spoke-SDP path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that spoke-SDP. When BPDUs are sent out other egress spoke-SDPs, the newly calculated root path cost is used.

STP suggests that the path cost is defined as a function of the link bandwidth. Since spoke-SDPs are controlled by complex queuing dynamics, the STP path cost is a purely static configuration.

The default value for spoke-SDP path cost is 10. This parameter can be modified within a range of 1 to 200000000 (1 is the lowest cost).

CLI Syntax:
```
config>service>vpls>spoke-sdp>stp#
path-cost stp-path-cost
```

Range: 1 to 200000000

Default: 10

Restore Default: no path-cost

Spoke SDP Edge Port

The spoke-SDP edge-port command is used to reduce the time it takes a spoke-SDP to reach the forwarding state when the spoke-SDP is on the edge of the network, and therefore has no further STP bridge to handshake with.

The edge-port command is used to initialize the internal OPER_EDGE variable. At any time, when OPER_EDGE is false on a spoke-SDP, the normal mechanisms are used to transition to the forwarding state (see Forward Delay). When OPER_EDGE is true, STP assumes that the remote end agrees to transition to the forwarding state without actually receiving a BPDU with an agreement flag set.

The OPER_EDGE variable will dynamically be set to false if the spoke-SDP receives BPDUs (the configured edge-port value does not change). The OPER_EDGE variable will dynamically be set to true if auto-edge is enabled and STP concludes there is no bridge behind the spoke-SDP.

When STP on the spoke-SDP is administratively disabled and re-enabled, the OPER_EDGE is re-initialized to the spoke-SDP configured for edge-port.

Valid values for spoke-SDP edge-port are enabled and disabled with disabled being the default.

CLI Syntax:
```
config>service>vpls>spoke-sdp>stp#
[no] edge-port
```

Default: no edge-port

Spoke SDP Auto Edge

The spoke-SDP edge-port command is used to instruct STP to dynamically decide whether the spoke-SDP is connected to another bridge.

If auto-edge is enabled, and STP concludes there is no bridge behind the spoke-SDP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see Spoke SDP Edge Port).

Valid values for spoke-SDP auto-edge are enabled and disabled with enabled being the default.

CLI Syntax:
```
config>service>vpls>spoke-sdp>stp#
[no] auto-edge
```

Default: auto-edge

Spoke SDP Link Type

The spoke-SDP link-type command instructs STP on the maximum number of bridges behind this spoke-SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected by a shared media, their spoke-SDPs should all be configured as shared, and timer-based transitions are used.

Valid values for spoke-SDP link-type are shared and pt-pt with pt-pt being the default.

CLI Syntax:
```
config>service>vpls>spoke-sdp>stp#
link-type {pt-pt|shared}
```

Default: link-type pt-pt

Restore Default: no link-type

3.5.3.8.2 Spoke SDP STP Operational States

The operational state of STP within a spoke-SDP controls how BPDUs are transmitted and handled when received. Defined states are:

- Operationally Disabled
- Operationally Discarding
- Operationally Learning
- Operationally Forwarding

Operationally Disabled

Operationally disabled is the normal operational state for STP on a spoke-SDP in a VPLS that has any of the following conditions:

- VPLS state administratively down
- Spoke SDP state administratively down
- Spoke SDP state operationally down

If the spoke-SDP enters the operationally up state with the STP administratively up and the spoke-SDP STP state is up, the spoke-SDP will transition to the STP spoke-SDP discarding state.

When, during normal operation, the router detects a downstream loop behind a spoke-SDP, BPDUs can be received at a very high rate. To recover from this situation, STP will transition the spoke-SDP to a disabled state for the configured forward-delay duration.

Operationally Discarding

A spoke-SDP in the discarding state only receives and sends BPDUs, building the local correct STP state for each spoke-SDP while not forwarding actual user traffic. The duration of the discarding state is explained in section Forward Delay.

Note: In previous versions of the STP standard, the discarding state was called a blocked state.

Operationally Learning

The learning state allows population of the MAC forwarding table before entering the forwarding state. In this state no user traffic is forwarded.

Operationally Forwarding

Configuration BPDUs are sent out a spoke-SDP in the forwarding state. Layer 2 frames received on the spoke-SDP are source learned and destination forwarded according to the FDB. Layer 2 frames received on other forwarding interfaces and destined for the spoke-SDP are also forwarded.

Spoke SDP BPDU Encapsulation States

IEEE 802.1D (referred as dot1d) and Cisco's per VLAN Spanning Tree (PVST) BPDU encapsulations are supported on a per spoke-SDP basis. STP is associated with a VPLS service like PVST is per VLAN. The main difference resides in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDU.

Table 39 shows differences between dot1D and PVST Ethernet BPDU encapsulations based on the interface encap-type field:

Table 39 **Spoke SDP BPDU Encapsulation States**

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
Destination MAC	01:80:c2:00:00:00	01:80:c2:00:00:00	N/A	01:00:0c:cc:cc:cd
Source MAC	Sending Port MAC	Sending Port MAC	N/A	Sending Port MAC
EtherType	N/A	0x81 00	N/A	0x81 00
Dot1p and CFI	N/A	0xe	N/A	0xe
Dot1q	N/A	VPLS spoke-SDP ID	N/A	VPLS spoke-SDP encap value
Length	LLC Length	LLC Length	N/A	LLC Length
LLC DSAP SSAP	0x4242	0x4242	N/A	0xaaaa (SNAP)
LLC CNTL	0x03	0x03	N/A	0x03
SNAP OUI	N/A	N/A	N/A	00 00 0c (Cisco OUI)

Table 39 **Spoke SDP BPDU Encapsulation States (Continued)**

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
SNAP PID	N/A	N/A	N/A	01 0b
CONFIG or TCN BPDU	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d
TLV: Type and Len	N/A	N/A	N/A	58 00 00 00 02
TLV: VLAN	N/A	N/A	N/A	VPLS spoke-SDP encap value
Padding	As Required	As Required	N/A	As Required

Each spoke-SDP has a Read Only operational state that shows which BPDU encapsulation is currently active on the spoke-SDP. The following states apply:

- **Dot1d** specifies that the switch is currently sending IEEE 802.1D standard BPDUs. The BPDUs will be tagged or non-tagged based on the encapsulation type of the egress interface and the encapsulation value defined in the spoke-SDP. A spoke-SDP defined on an interface with encapsulation type dot1q will continue in the dot1d BPDU encapsulation state until a PVST encapsulated BPDU is received, after which the spoke-SDP will convert to the PVST encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged if the interface encapsulation type is defined to dot1q.
- **PVST** specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The spoke-SDP continues in the PVST BPDU encapsulation state until a dot1d encapsulated BPDU is received, in which case the spoke-SDP reverts to the dot1d encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged with the encapsulation value defined for the spoke-SDP.

Dot1d is the initial and only spoke-SDP BPDU encapsulation state for spoke-SDPs defined on Ethernet interface with encapsulation type set to null.

Each transition between encapsulation types optionally generates an alarm that can be logged and optionally transmitted as an SNMP trap.

3.5.3.8.3 Configuring VPLS Spoke SDPs with Split Horizon

To configure spoke-SDPs with a split horizon group, add the split-horizon-group parameter when creating the spoke-SDP. Traffic arriving on a SAP or spoke-SDP within a split horizon group will not be copied to other SAPs or spoke-SDPs in the same split horizon group.

The following example shows a VPLS configuration with split horizon enabled:

```
-----------------------------------------------
*A:ALA-1>config>service# *A:ALA-1>config>service# info
-----------------------------------------------
...
vpls 800 customer 6001 vpn 700 create
     description "VPLS with split horizon for DSL"
     stp
         shutdown
     exit
     spoke-sdp 51:15 split-horizon-group DSL-group1 create
     exit
     split-horizon-group DSL-group1
         description "Split horizon group for DSL"
     exit
     no shutdown
exit
...
-----------------------------------------------
*A:ALA-1>config>service#
```

3.5.4 Configuring VPLS Redundancy

This section discusses VPLS redundancy service management tasks.

3.5.4.1 Creating a Management VPLS for SAP Protection

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for SAP protection and provides the CLI commands, see Figure 95. The tasks below should be performed on both nodes providing the protected VPLS service.

Before configuring a management VPLS, first read VPLS Redundancy for an introduction to the concept of management VPLS and SAP redundancy.

1. Create an SDP to the peer node.
2. Create a management VPLS.

3. Define a SAP in the m-vpls on the port toward the MTU. The port must be dot1q or qinq tagged. The SAP corresponds to the (stacked) VLAN on the MTU in which STP is active.
4. Optionally modify STP parameters for load balancing.
5. Create a mesh SDP in the m-vpls using the SDP defined in Step 1. Ensure that this mesh SDP runs over a protected LSP (see note below).
6. Enable the management VPLS service and verify that it is operationally up.
7. Create a list of VLANs on the port that are to be managed by this management VPLS.
8. Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 6.

Note: The mesh SDP should be protected by a backup LSP or Fast Reroute. If the mesh SDP were to go down, STP on both nodes would go to “forwarding” state and a loop would occur.

Figure 95 **Example Configuration for Protected VPLS SAP**

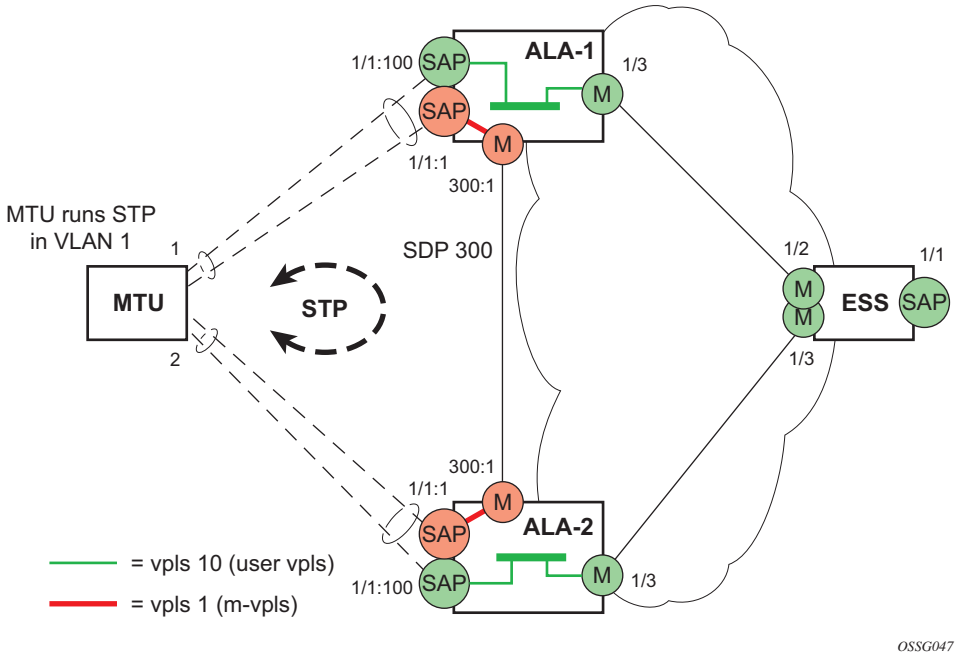

Use the following CLI syntax to create a management VPLS on the 7450 ESS or 7750 SR:

CLI Syntax: `config>service# sdp` *`sdp-id`* `mpls create`

```
far-end ip-address
lsp lsp-name
no shutdown
```

CLI Syntax:

```
vpls service-id customer customer-id [m-vpls] create
description description-string
sap sap-id create
     managed-vlan-list
          range vlan-range
mesh-sdp sdp-id:vc-id create
stp
no shutdown
```

The following example shows a VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
       sdp 300 mpls create
            far-end 10.0.0.20
            lsp "toALA-A2"
            no shutdown
       exit
       vpls 1 customer 1 m-vpls create
            sap 1/1/1:1 create
                managed-vlan-list
                    range 100-1000
                exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

3.5.4.2 Creating a Management VPLS for Spoke SDP Protection

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for spoke-SDP protection and provides the CLI commands, see Figure 96. The tasks below should be performed on all four nodes providing the protected VPLS service.
Before configuring a management VPLS, first read Configuring a VPLS SAP for an introduction to the concept of management VPLS and spoke-SDP redundancy.

1. Create an SDP to the local peer node (node ALA-A2 in the example below).

2. Create an SDP to the remote peer node (node ALA-B1 in the example below).
3. Create a management VPLS.
4. Create a spoke-SDP in the m-vpls using the SDP defined in Step 1. Ensure that this mesh SDP runs over a protected LSP (see note below).
5. Enable the management VPLS service and verify that it is operationally up.
6. Create a spoke-SDP in the m-vpls using the SDP defined in Step 2. Optionally, modify STP parameters for load balancing (see Configuring Load Balancing with Management VPLS).
7. Create one or more user VPLS services with spoke-SDPs on the tunnel SDP defined by Step 2.

As long as the user spoke-SDPs created in step 7are in this same tunnel SDP with the management spoke-SDP created in step 6, the management VPLS will protect them.

The SDP should be protected by, for example, a backup LSP or Fast Reroute. If the SDP were to go down, STP on both nodes would go to “forwarding” state and a loop would occur.

Figure 96 **Example Configuration for Protected VPLS Spoke SDP**

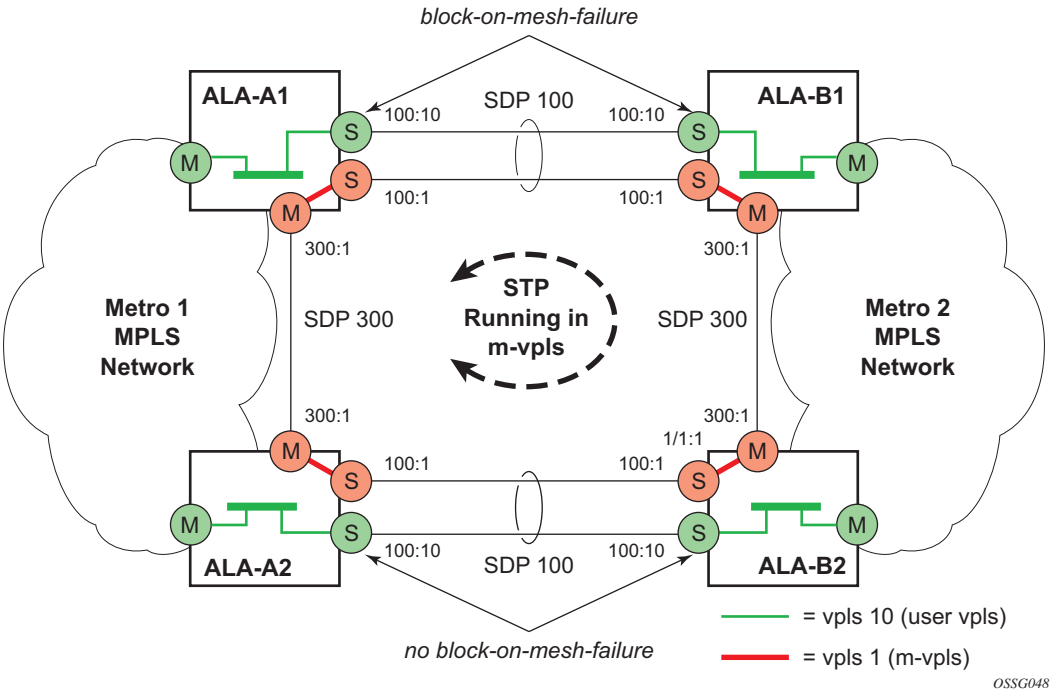

Use the following CLI syntax to create a management VPLS for spoke-SDP protection:

CLI Syntax:
```
config>service# sdp sdp-id mpls create
far-end ip-address
lsp lsp-name
no shutdown
```

CLI Syntax:
```
vpls service-id customer customer-id [m-vpls] create
description description-string
mesh-sdp sdp-id:vc-id create
spoke-sdp sdp-id:vc-id create
stp
no shutdown
```

The following example shows a VPLS configuration:

```
*A:ALA-A1>config>service# info
----------------------------------------------
...
       sdp 100 mpls create
            far-end 10.0.0.30
            lsp "toALA-B1"
            no shutdown
       exit
       sdp 300 mpls create
            far-end 10.0.0.20
            lsp "toALA-A2"
            no shutdown
       exit
       vpls 101 customer 1 m-vpls create
            spoke-sdp 100:1 create
            exit
            meshspoke-sdp 300:1 create
            exit
            stp
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-A1>config>service#
```

3.5.4.3 Configuring Load Balancing with Management VPLS

With the concept of management VPLS, it is possible to load balance the user VPLS services across the two protecting nodes. This is done by creating two management VPLS instances, where both instances have different active QinQ spokes (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two QinQ spokes. Load balancing can be achieved in both the SAP protection and spoke-SDP protection scenarios.

See Figure 97 shows an example configuration for load balancing across two protected VPLS spoke-SDPs.

Figure 97 **Example Configuration for Load Balancing Across Two Protected VPLS Spoke SDPs**

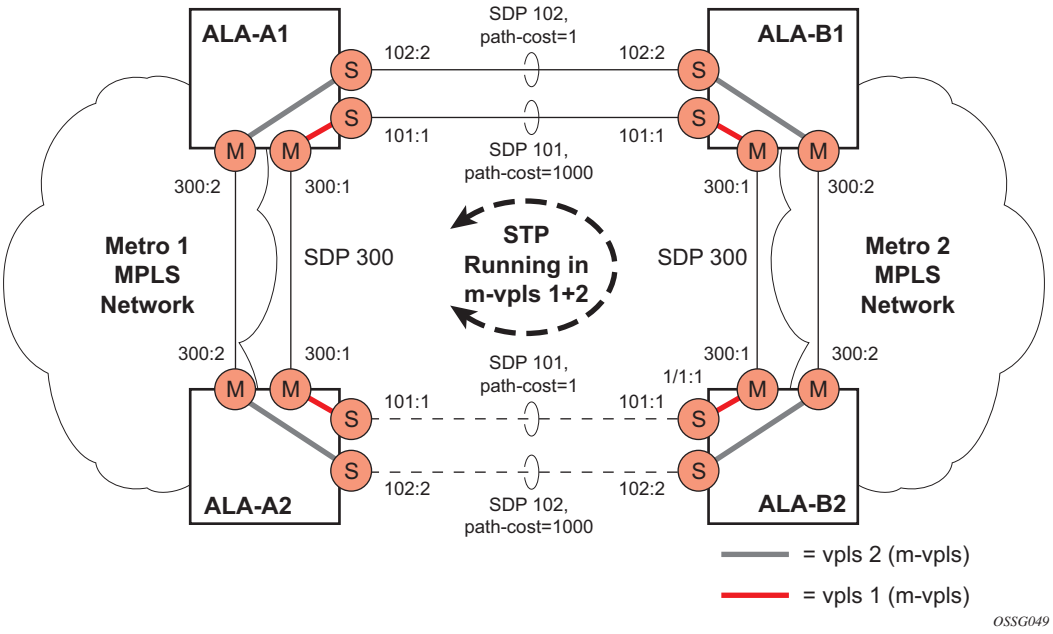

Use the following CLI syntax to create a load balancing across two management VPLS instances:

CLI Syntax:
```
config>service# sdp sdp-id mpls create
far-end ip-address
lsp lsp-name
no shutdown
```

CLI Syntax:
```
vpls service-id customer customer-id [m-vpls] create
description description-string
mesh-sdp sdp-id:vc-id create
```

```
            spoke-sdp sdp-id:vc-id create
                 stp
                      path-cost
            stp
            no shutdown
```

Note: The STP path costs in each peer node should be reversed.

The following example shows the VPLS configuration on ALA-A1 (top left, IP address 10.0.0.10):

```
*A:ALA-A1>config>service# info
----------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.30
            lsp "1toALA-B1"
            no shutdown
      exit
      sdp 102 mpls create
            far-end 10.0.0.30
            lsp "2toALA-B1"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
           no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
----------------------------------------------
*A:ALA-A1>config>service#
```

The following example shows the VPLS configuration on ALA-A2 (bottom left, IP address 10.0.0.20):

```
*A:ALA-A2>config>service# info
----------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.40
            lsp "1toALA-B2"
            no shutdown
      exit
      sdp 102 mpls create
            far-end 10.0.0.40
            lsp "2toALA-B2"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
           no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
----------------------------------------------
*A:ALA-A2>config>service#
```

The following example shows the VPLS configuration on ALA-A3 (top right, IP address 10.0.0.30):

```
*A:ALA-A1>config>service# info
----------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.10
            lsp "1toALA-A1"
            no shutdown
      exit
```

```
      sdp 102 mpls create
            far-end 10.0.0.10
            lsp "2toALA-A1"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
            no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
-------------------------------------------
*A:ALA-A1>config>service#
```

The following example shows the VPLS configuration on ALA-A4 (bottom right, IP address 10.0.0.40):

```
*A:ALA-A2>config>service# info
-------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.20
            lsp "1toALA-B2"
            no shutdown
      exit
      sdp 102 mpls create
            far-end 10.0.0.20
            lsp "2toALA-B2"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:1 create
```

```
            exit
            stp
            exit
            no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
-------------------------------------------
*A:ALA-A2>config>service#
```

3.5.4.4 Configuring Selective MAC Flush

Use the following CLI syntax to enable selective MAC Flush in a VPLS.

CLI Syntax: `config>service# vpls` *`service-id`*
`send-flush-on-failure`

Use the following CLI syntax to disable selective MAC Flush in a VPLS.

CLI Syntax: `config>service# vpls` *`service-id`*
`no send-flush-on-failure`

3.5.4.5 Configuring Multi-Chassis Endpoints

The following output shows configuration examples of multi-chassis redundancy and the VPLS configuration. The configurations in the graphics depicted in Inter-Domain VPLS Resiliency Using Multi-Chassis Endpoints are expressed in this output.

Node mapping to the following examples in this section:

- PE3 = Dut-B
- PE3' = Dut-C
- PE1 = Dut-D
- PE2 = Dut-E

PE3

```
*A:Dut-B>config>redundancy>multi-chassis# info
----------------------------------------------
            peer 3.1.1.3 create
                peer-name "Dut-C"
                description "mcep-basic-tests"
                source-address 2.1.1.2
                mc-endpoint
                    no shutdown
                    bfd-enable
                    system-priority 50
                exit
                no shutdown
            exit
----------------------------------------------
*A:Dut-B>config>redundancy>multi-chassis#


*A:Dut-B>config>service>vpls# info
----------------------------------------------
            fdb-table-size 20000
            send-flush-on-failure
            stp
                shutdown
            exit
            endpoint "mcep-t1" create
                no suppress-standby-signaling
                block-on-mesh-failure
                mc-endpoint 1
                    mc-ep-peer Dut-C
                exit
            exit
            mesh-sdp 201:1 vc-type vlan create
            exit
            mesh-sdp 211:1 vc-type vlan create
            exit
            spoke-sdp 221:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
                precedence 1
            exit
            spoke-sdp 231:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
                precedence 2
            exit
            no shutdown
----------------------------------------------
*A:Dut-B>config>service>vpls#
```

PE3' Dut-C

```
:Dut-C>config>redundancy>multi-chassis# info
```

```
-----------------------------------------------
            peer 2.1.1.2 create
                peer-name "Dut-B"
                description "mcep-basic-tests"
                source-address 3.1.1.3
                mc-endpoint
                    no shutdown
                    bfd-enable
                    system-priority 21
                exit
                no shutdown
            exit
-----------------------------------------------
*A:Dut-C>config>redundancy>multi-chassis#

*A:Dut-C>config>service>vpls# info
-----------------------------------------------
            fdb-table-size 20000
            send-flush-on-failure
            stp
                shutdown
            exit
            endpoint "mcep-t1" create
                no suppress-standby-signaling
                block-on-mesh-failure
                mc-endpoint 1
                    mc-ep-peer Dut-B
                exit
            exit
            mesh-sdp 301:1 vc-type vlan create
            exit
            mesh-sdp 311:1 vc-type vlan create
            exit
            spoke-sdp 321:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
                precedence 3
            exit
            spoke-sdp 331:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
            exit
            no shutdown
-----------------------------------------------
*A:Dut-C>config>service>vpls#
```

PE1 Dut-D

```
*A:Dut-D>config>redundancy>multi-chassis# info
-----------------------------------------------
            peer 5.1.1.5 create
                peer-name "Dut-E"
                description "mcep-basic-tests"
                source-address 4.1.1.4
```

```
                mc-endpoint
                    no shutdown
                    bfd-enable
                    system-priority 50
                    passive-mode
                exit
                no shutdown
            exit
-----------------------------------------------
*A:Dut-D>config>redundancy>multi-chassis#

*A:Dut-D>config>service>vpls# info
-----------------------------------------------
            fdb-table-size 20000
            propagate-mac-flush
            stp
                shutdown
            exit
            endpoint "mcep-t1" create
                block-on-mesh-failure
                mc-endpoint 1
                    mc-ep-peer Dut-E
                exit
            exit
            mesh-sdp 401:1 vc-type vlan create
            exit
            spoke-sdp 411:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
                precedence 2
            exit
            spoke-sdp 421:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
                precedence 1
            exit
            mesh-sdp 431:1 vc-type vlan create
            exit
            no shutdown
-----------------------------------------------
*A:Dut-D>config>service>vpls#
```

PE2 Dut-E

```
*A:Dut-E>config>redundancy>multi-chassis# info
-----------------------------------------------
            peer 4.1.1.4 create
                peer-name "Dut-D"
                description "mcep-basic-tests"
                source-address 5.1.1.5
                mc-endpoint
                    no shutdown
                    bfd-enable
```

```
                    system-priority 22
                    passive-mode
                exit
                no shutdown
            exit
-----------------------------------------------
*A:Dut-E>config>redundancy>multi-chassis#

*A:Dut-E>config>service>vpls# info
-----------------------------------------------
            fdb-table-size 20000
            propagate-mac-flush
            stp
                shutdown
            exit
            endpoint "mcep-t1" create
                block-on-mesh-failure
                mc-endpoint 1
                    mc-ep-peer Dut-D
                exit
            exit
            spoke-sdp 501:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
                precedence 3
            exit
            spoke-sdp 511:1 vc-type vlan endpoint "mcep-t1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
            exit
            mesh-sdp 521:1 vc-type vlan create
            exit
            mesh-sdp 531:1 vc-type vlan create
            exit
            no shutdown
-----------------------------------------------
*A:Dut-E>config>service>vpls#
```

3.5.5 ATM/Frame Relay PVC Access and Termination on a VPLS Service

The application as shown in Figure 98 provides access to a VPLS service to Frame Relay and ATM users connected either directly or through an ATM access network to a 7750 SR PE node. The 7750 SR supports a Frame Relay or an ATM VC-delimited Service Access Point (SAP) terminating on a VPLS service.

Figure 98 **ATM/Frame Relay PVC Access and Termination on a VPLS Example**

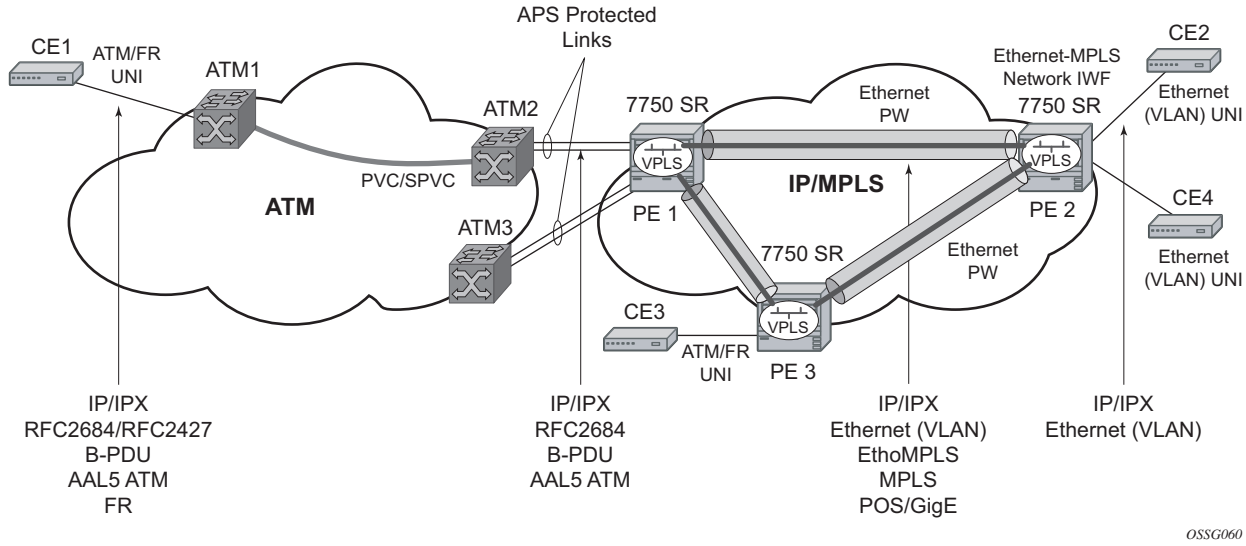

RFC 2427-encapsulated or RFC 2684-encapsulated untagged Ethernet/802.3 frames (with or without Frame Check Sequence (FCS)) or BPDUs from a customer's bridge device are received on a specified SAP over an ATM or Frame Relay interface on the 7750 SR. The Frame Relay or ATM-related encapsulation is stripped and the frames (without FCS) are forwarded toward destination SAPs either locally, or using SDPs associated with the VPLS service (as dictated by destination MAC address VPLS processing). In the egress direction, the received untagged frames are encapsulated into RFC 2427 or RFC 2684 (no Q-tags are added, no FCS in the forwarded frame) and sent over ATM or a FR VC toward the customer CPE.

When AAL5 RFC2427/2684 encapsulated tagged frames are received from the customer's bridge on an FR/ATM SAP, the tags are transparent and the frames are processed as described above with the exception that the frames forwarded toward the destination(s) will have the received tags preserved. Similarly in the egress direction, the received tagged Ethernet frames are encapsulated as is (i.e. Q-tags are again transparent and preserved) into RFC 2427/2684 and sent over the FR/ ATM PVC toward the customer CPE. Since the tagging is transparent, the 7750 SR performs unqualified MAC learning (for example, MAC addresses are learned without reference to VLANs they are associated with). Because of that, MAC addresses used must be unique across all the VLANs used by the customer for a specified VPLS service instance. If a customer wants a per-VLAN separation, then the VLAN traffic that needs to be separated must come on different VCs (different SAPs) associated with different VPLS service instances.

All VPLS functionality available on the 7750 SR is applicable to FR and ATM-delimited VPLS SAPs. For example, bridged PDUs received over ATM SAP can be tunneled through or dropped; all Forwarding Information Base functionality applies; packet level QoS and MAC filtering applies; etc. Also, split horizon groups are applicable to ATM SAPs terminating on VPLS. In other words, frame forwarding between ATM SAPs, also referred to as VCI-to-VCI forwarding, within the same group is disabled.

The Ethernet pseudowire is established using Targeted LDP (TLDP) signaling and uses the ether, vlan, or vpls VC type on the SDP. The SDP can be an MPLS or a GRE type.

3.5.6 Configuring BGP Auto-Discovery

This section provides important information to explain the different configuration options used to populate the required BGP AD and generate the LDP generalized pseudowire-ID FEC fields. There are a large number of configuration options that are available with the this feature. Not all these configurations option are required to start using BGP AD. At the end of this section, it will be apparent that a very simple configuration will automatically generate the required values used by BGP and LDP. In most cases, deployments will provide full mesh connectivity between all nodes across a VPLS instance. However, capabilities are available to influence the topology and build hierarchies or hub and spoke models.

3.5.6.1 Configuration Steps

Using Figure 99, assume PE6 was previously configured with VPLS 100 as indicated by the configurations lines in the upper right. The BGP AD process will commence after PE134 is configured with the VPLS 100 instance as shown in the upper left. This shows a very basic and simple BGP AD configuration. The minimum requirement for enabling BGP AD on a VPLS instance is configuring the VPLS-ID and point to a pseudowire template.

Figure 99 **BGP AD Configuration Example**

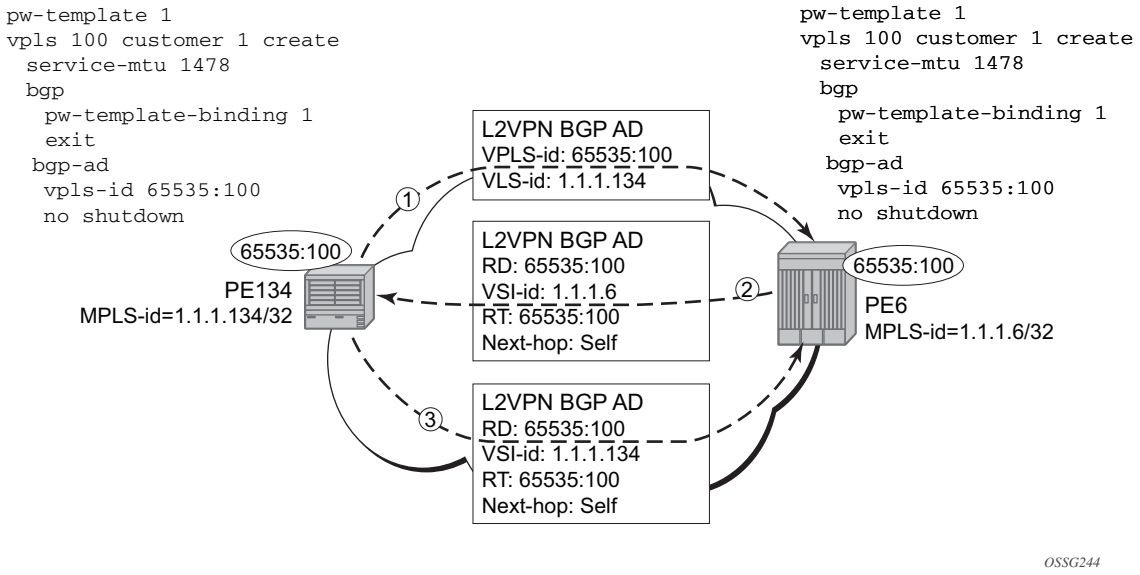

In many cases, VPLS connectivity is based on a pseudowire mesh. To reduce the configuration requirement, the BGP values can be automatically generated using the VPLS-ID and the MPLS router-ID. By default, the lower six bytes of the VPLS-ID are used to generate the RD and the RT values. The VSI-ID value is generated from the MPLS router-ID. All of these parameters are configurable and can be coded to suit requirements and build different topologies.

```
PE134>config>service>vpls>bgp-ad#
[no] shutdown - Administratively enable/disable BGP auto-discovery
vpls-id - Configure VPLS-ID
vsi-id + Configure VSI-id
```

A helpful command shows the service information, the BGP parameters and the SDP bindings in use. When the discovery process is completed successfully each endpoint will have an entry for the service.

```
PE134>#  show service l2-route-table
===============================================================================
Services: L2 Route Information - Summary Service
===============================================================================
Svc Id     L2-Routes (RD-Prefix)                 Next Hop        Origin
               Sdp Bind Id
-------------------------------------------------------------------------------
100        65535:100-1.1.1.6                     1.1.1.6         BGP-L2
               17406:4294967295
-------------------------------------------------------------------------------
No. of L2 Route Entries: 1
===============================================================================
PERs6>#

PERs6>#  show service l2-route-table
```

```
===============================================================================
Services: L2 Route Information - Summary Service
===============================================================================
Svc Id     L2-Routes (RD-Prefix)                Next Hop        Origin
               Sdp Bind Id
-------------------------------------------------------------------------------
100        65535:100-1.1.1.134                  1.1.1.134       BGP-L2
               17406:4294967295
-------------------------------------------------------------------------------
No. of L2 Route Entries: 1
===============================================================================
PERs6>#
```

When only one of the endpoints has an entry for the service in the l2-routing-table, it is most likely a problem with the RT values used for import and export. This would most likely happen when different import and export RT values are configured using a router policy or the route-target command.

Service specific commands continue to be available to show service specific information, including status.

```
PERs6# show service sdp-using
===============================================================================
SDP Using
===============================================================================
SvcId       SdpId               Type    Far End        Opr S* I.Label  E.Label
-------------------------------------------------------------------------------
100         17406:4294967295    BgpAd   1.1.1.134      Up     131063   131067
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

BGP AD will advertise the VPLS-ID in the extended community attribute, VSI-ID in the NLRI and the local PE ID in the BGP next hop. At the receiving PE, the VPLS-ID is compared against locally provisioned information to determine whether the two PEs share a common VPLS. If it is found that they do, the BGP information is used in the signaling phase (see Configuring BGP VPLS).

3.5.6.2 LDP Signaling

T-LDP is triggered when the VPN endpoints have been discovered using BGP. The T-LDP session between the PEs is established when one does not exist. The far-end IP address required for the T-LDP identification is learned from the BGP AD next hop information. The pw-template and pw-template-binding configuration statements are used to establish the automatic SDP or to map to the appropriate SDP. The FEC129 content is built using the following values:

- AGI from the locally configured VPLS-ID.

- The SAII from the locally configured VSI-ID.
- The TAII from the VSI-ID contained in the last 4 bytes of the received BGP NLRI.

Figure 100 shows the different detailed phases of the LDP signaling path, post BGP AD completion. It also indicates how some fields can be auto generated when they are not specified in the configuration.

Figure 100 **BGP AD Triggering LDP Functions**

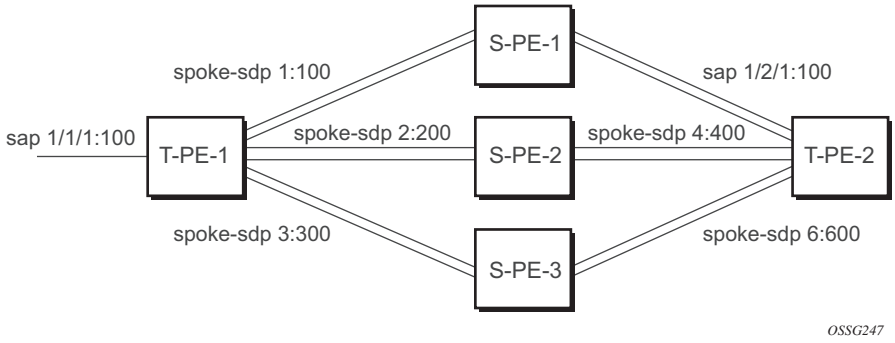

The first command shows the LDP peering relationships that have been established (Figure 101). The type of adjacency is displayed in the “Adj Type” column. In this case the type is “Both” meaning link and targeted sessions have been successfully established.

Figure 101 **Show Router LDP Session Output**

```
PERs6# show router ldp session
LDP Sessions
Peer LDP Id          Adj Type    State          Msg Sent  Msg Recv  Up Time
1.1.1.134:0          Both        Established    21482     21482     0d 15:38:44
No. of Sessions: 1
```

0988

The second command shows the specific LDP service label information broken up per FEC element type, 128 or 129, basis (Figure 102). The information for FEC element 129 includes the AGI, SAII and the TAII.

Figure 102 **Show Router LDP Bindings FEC-Type Services**

```
PERs6# show router ldp bindings fec-type services
LDP LSR ID: 1.1.1.6
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up,  D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, C - Cpipe Service
        TLV - (Type, Length: Value)
LDP Service FEC 128 Bindings
Type   VCId       SvcId      SDPId  Peer            IngLbl  EgrLbl  LMTU  RMTU
No Matching Entries Found

LDP Service FEC 129 Bindings
AGI                                    SAII                    TAII
Type               SvcId      SDPId  Peer            IngLbl  EgrLbl  LMTU  RMTU
65535:100                             1.1.1.6                 1.1.1.134
V-Eth              100        17406  1.1.1.134       131063U 131067S 1464  1464
No. of FEC 129s: 1
```

0989

3.5.6.3 Pseudowire Template

The pseudowire template is defined under the top-level service command (**config>service> pw-template**) and specifies whether to use an automatically generated SDP or manually configured SDP. It also provides the set of parameters required for establishing the pseudowire (SDP binding) as displayed below:

```
PERs6>config>service# pw-template 1 create
 -[no] pw-template <policy-id> [use-provisioned-sdp | prefer-provisioned-sdp]
<policy-id> : [1..2147483647]
<use-provisioned-s*> : keyword
<prefer-provisioned*> : keyword


[no] accounting-pol*    - Configure accounting-policy to be used
[no] auto-learn-mac*    - Enable/Disable automatic update of MAC protect list
[no] block-on-peer-*    - Enable/Disable block traffic on peer fault
[no] collect-stats      - Enable/disable statistics collection
[no] control word       - Enable/Disable the use of Control Word
[no] disable-aging      - Enable/disable aging of MAC addresses
[no] disable-learni*    - Enable/disable learning of new MAC addresses
[no] discard-unknow*    - Enable/disable discarding of frames with unknown source
                          MAC address
     egress             + Spoke SDP binding egress configuration
[no] force-qinq-vc-*    - Forces qinq-vc-type forwarding in the data-path
[no] force-vlan-vc-*    - Forces vlan-vc-type forwarding in the data-path
[no] hash-label         - Enable/disable use of hash-label
     igmp-snooping      + Configure IGMP snooping parameters
     in gr ess          + Spoke SDP binding ingress configuration
[no] l2pt-terminati*    - Configure L2PT termination on this spoke SDP
```

```
[no] limit-mac-move     - Configure mac move
[no] mac-pinning        - Enable/disable MAC address pinning on this spoke SDP
[no] max-nbr-mac-ad*    - Configure the maximum number of MAC entries in the FDB
                          from this SDP
[no] restrict-prote*    - Enable/disable protected src MAC restriction
[no] sdp-exclude        - Configure excluded SDP group
[no] sdp-include        - Configure included SDP group
[no] split-horizon-*    + Configure a split horizon group
     stp                + Configure STP parameters
     vc-type            - Configure VC type
[no] vlan-vc-tag        - Configure VLAN VC tag
```

A **pw-template-binding** command configured within the VPLS service under the **bgp-ad** sub-command is a pointer to the pw-template that should be used. If a VPLS service does not specify an import-rt list, then that binding applies to all route targets accepted by that VPLS. The **pw-template-bind** command can select a different template on a per import-rt basis. It is also possible to specify specific pw-templates for some route targets with a VPLS service and use the single **pw-template-binding** command to address all unspecified but accepted imported targets.

Figure 103 **PW-Template-Binding CLI Syntax**

```
PERs6>config>service>vpls>bgp-ad# pw-template-binding
  - pw-template-binding <policy-id> [split-hozion-group <group-name>] [import-
rt
    {ext-community, ...(upto 5 max)}]
  - no pw-template-binding <policy-id>

 <policy-id>          : [1..2147483647]
 <group-name>         : [32 chars max]
 <ext-community>      : target:{<ip-addr:comm-val>|<as-number:ext-comm-val>}
                        ip-addr      - a.b.c.d
                        comm-val     - [0..65535]
                        as-number    - [1..65535]
                        ext-comm-val - [0..4294967295]
```

0990

It is important understand the significance of the split horizon group used by the pw-template. Traditionally, when a VPLS instance was manually created using mesh-sdp bindings, these were automatically placed in a common split horizon group to prevent forwarding between the pseudowire in the VPLS instances. This prevents loops that would have otherwise occurred in the Layer 2 service. When automatically discovering VPLS service using BGP AD the service provider has the option of associating the auto-discovered pseudowire with a split horizon group to control the forwarding between pseudowires.

3.5.7 Configuring BGP VPLS

This section gives a configuration example required to bring up BGP VPLS in the VPLS PEs depicted in Figure 104:

Figure 104 **BGP VPLS Example**

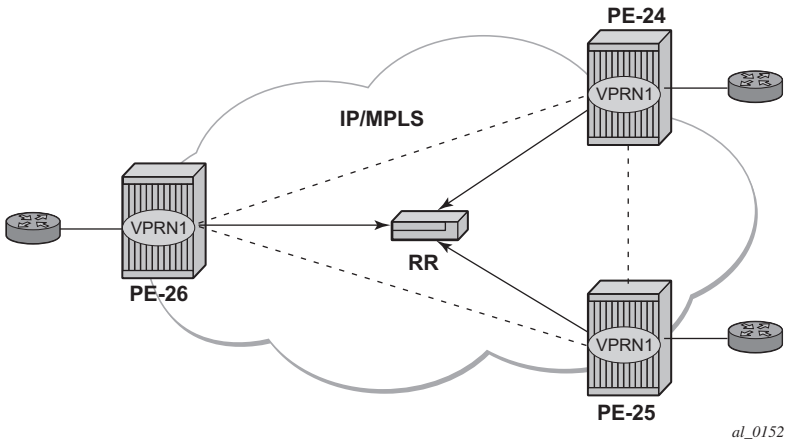

The red BGP VPLS is configured in the PE24, PE25 and PE26 using the commands shown in the following CLI examples.

```
*A:PE24>config>service>vpls# info
----------------------------------------------
            bgp
                route-distinguisher 65024:600
                route-target export target:65019:600 import target:65019:600
                pw-template-binding 1
            exit
            bgp-vpls
                max-ve-id 100
                ve-name 24
                    ve-id 24
                exit
                no shutdown
            exit
            sap 1/1/20:600.* create
            exit
            no shutdown
----------------------------------------------
*A:PE24>config>service>vpls#


*A:PE25>config>service>vpls# info
----------------------------------------------
            bgp
                route-distinguisher 65025:600
                route-target export target:65019:600 import target:65019:600
                pw-template-binding 1
```

```
            exit
            bgp-vpls
                max-ve-id 100
                ve-name 25
                    ve-id 25
                exit
                no shutdown
            exit
            sap 1/1/19:600.* create
            exit
            no shutdown
-----------------------------------------------
*A:PE25>config>service>vpls#


*A:PE26>config>service>vpls# info
-----------------------------------------------
            bgp
                route-distinguisher 65026:600
                route-target export target:65019:600 import target:65019:600
                pw-template-binding 1
            exit
            bgp-vpls
                max-ve-id 100
                ve-name 26
                    ve-id 26
                exit
                no shutdown
            exit
            sap 5/2/20:600.* create
            exit
            no shutdown
-----------------------------------------------
*A:PE26>config>service>vpls#
```

3.5.7.1 Configuring a VPLS Management Interface

Use the following CLI syntax to create a VPLS management interface.

CLI Syntax:

```
config>service>vpls# interface ip-int-name
address ip-address[/mask] [netmask]
arp-timeout seconds
description description-string
mac ieee-address
no shutdown
static-arp ip-address ieee-address
```

The following example shows the configuration.

```
A:ALA-49>config>service>vpls>interface# info detail
----------------------------------------------
            no description
            mac 14:31:ff:00:00:00
```

```
                address 123.231.10.10/24
                no arp-timeout
                no shutdown
-------------------------------------------
A:ALA-49>config>service>vpls>interface#
```

3.5.8 Configuring Policy-Based Forwarding for Deep Packet Inspection (DPI) in VPLS

The purpose of policy-based forwarding is to capture traffic from a customer and perform a deep packet inspection (DPI) and forward traffic, if allowed, by the DPI on the 7450 ESS or 7750 SR.

In the following example, the split horizon groups are used to prevent flooding of traffic. Traffic from customers enter at SAP 1/1/5:5. Due to the mac-filter 100 that is applied on ingress, all traffic with dot1p 07 marking will be forwarded to SAP 1/1/22:1, which is the DPI.

DPI performs packet inspection/modification and either drops the traffic or forwards the traffic back into the box through SAP 1/1/21:1. Traffic will then be sent to spoke-sdp 3:5.

SAP 1/1/23:5 is configured to see if the VPLS service is flooding all the traffic. If flooding is performed by the router then traffic would also be sent to SAP 1/1/23:5 (which it should not).

Figure 105 shows an example to configure policy-based forwarding for deep packet inspection on a VPLS service. For information about configuring filter policies, refer to the *7450 ESS, 7750 SR, and 7950 XRS Router Configuration Guide*.

Figure 105 **Policy-Based Forwarding For Deep Packet Inspection**

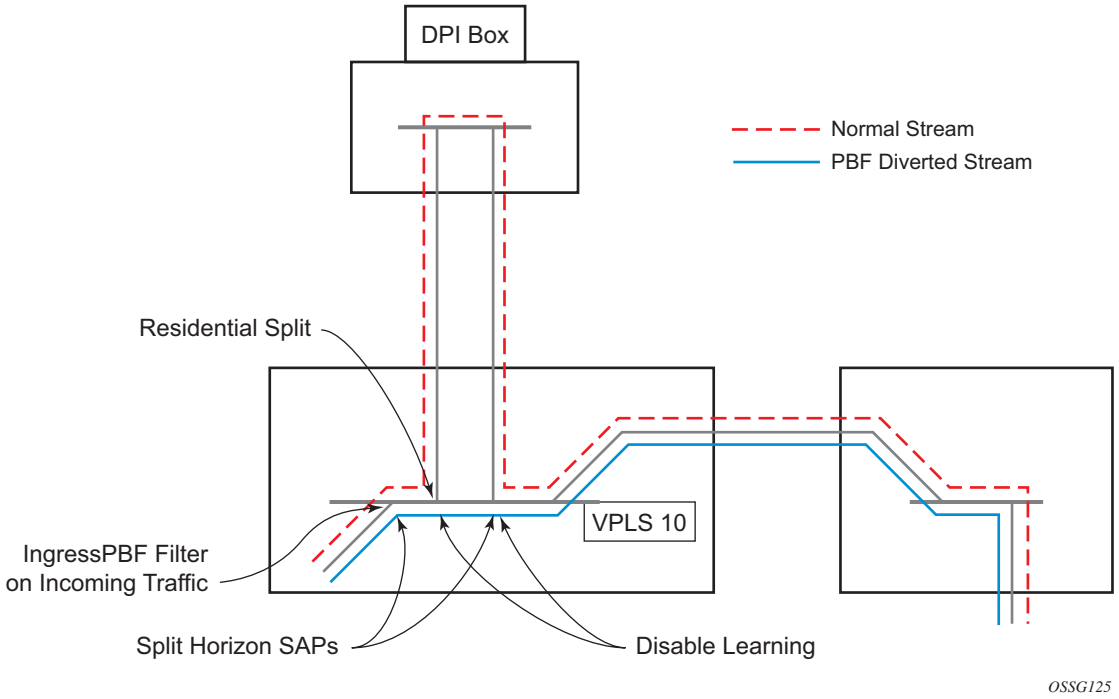

The following example shows the service configuration:

```
*A:ALA-48>config>service# info
----------------------------------------------
...
        vpls 10 customer 1 create
            service-mtu 1400
            split-horizon-group "dpi" residential-group create
            exit
            split-horizon-group "split" create
            exit
            stp
                shutdown
            exit
            igmp-host-tracking
                expiry-time 65535
                no shutdown
            exit
            sap 1/1/21:1 split-horizon-group "split" create
                disable-learning
                static-mac 00:00:00:31:11:01 create
            exit
            sap 1/1/22:1 split-horizon-group "dpi" create
                disable-learning
                static-mac 00:00:00:31:12:01 create
            exit
            sap 1/1/23:5 create
                static-mac 00:00:00:31:13:05 create
            exit
```

```
            no shutdown
        exit
...
-----------------------------------------------
*A:ALA-48>config>service#
```

The following example shows the MAC filter configuration:

```
*A:ALA-48>config>filter# info
-----------------------------------------------
...
        mac-filter 100 create
            default-action forward
            entry 10 create
                match
                    dot1p 7 7
                exit
                log 101
                action forward sap 1/1/22:1
            exit
        exit
...
-----------------------------------------------
*A:ALA-48>config>filter#
```

The following example shows the service configuration with a MAC filter:

```
*A:ALA-48>config>service# info
-----------------------------------------------
...
        vpls 10 customer 1 create
            service-mtu 1400
            split-horizon-group "dpi" residential-group create
            exit
            split-horizon-group "split" create
            exit
            stp
                shutdown
            exit
            igmp-host-tracking
                expiry-time 65535
                no shutdown
            exit
            sap 1/1/5:5 split-horizon-group "split" create
                ingress
                    filter mac 100
                exit
                static-mac 00:00:00:31:15:05 create
            exit
            sap 1/1/21:1 split-horizon-group "split" create
                disable-learning
                static-mac 00:00:00:31:11:01 create
            exit
            sap 1/1/22:1 split-horizon-group "dpi" create
                disable-learning
                static-mac 00:00:00:31:12:01 create
            exit
```

```
            sap 1/1/23:5 create
                static-mac 00:00:00:31:13:05 create
            exit
            spoke-sdp 3:5 create
            exit
            no shutdown
        exit
....
-------------------------------------------
*A:ALA-48>config>service#
```

3.5.9 Configuring VPLS E-Tree Services

When configuring a VPLS E-Tree service the **etree** keyword must be specified when the VPLS service is created. This is the first operation required before any SAPs or SDPs are added to the service, since the E-Tree service type affects the operations of the SAPs and SDP bindings.

When configuring AC SAPs the configuration model is very similar to normal SAPs. Since the VPLS service must be designated as an E-Tree, the default AC SAP is a root AC SAP. An E-Tree service with all root AC behaves just as a regular VPLS service. A leaf AC SAP must be configured for leaf behavior.

For root-leaf-tag SAPs, the SAP is created with both root and leaf VIDs. The 1/1/1:x.* or 1/1/1:x would be the typical format where x designates the root tag. A leaf-tag is configured at SAP creation and replaces the x with a leaf-tag VID. Combined statistics for root and leaf SAPs are reported under the SAP. There are no individual statistics shown for root and leaf.

The following example illustrates the configuration of a VPLS E-Tree service with root AC (default configuration for SAPs and SDP binds) and leaf AC interfaces, as well as a root leaf tag SAP and SDP bind.

In the example, the SAP 1/1/7:2006.200 is configured using the root-leaf-tag parameter, where the outer VID 2006 is used for root traffic and the outer VID 2007 is used for leaf traffic.

```
*A:ALA-48>config>service# info
-----------------------------------------------
...
        service vpls 2005 etree customer 1 create
            sap 1/1/1:2005 leaf-ac create
            exit
            sap 1/1/7:2006.200 root-leaf-tag leaf-tag 2007 create
            exit
            sap 1/1/7:0.* create
            exit
            spoke-sdp 12:2005 vc-type vlan root-leaf-tag create
                no shutdown
```

```
            exit
            spoke-sdp 12:2006 leaf-ac create
                no shutdown
            exit
            no shutdown
        exit
....
*A:ALA-48>config>service# info
----------------------------------------------
```

3.6 Service Management Tasks

This section describes VPLS service management tasks.

3.6.1 Modifying VPLS Service Parameters

You can change existing service parameters. The changes are applied immediately. To display a list of services, use the **show service service-using vpls** command. Enter the parameter such as description, SAP, SDP, and/or service-MTU command syntax, and then enter the new information.

The following shows a modified VPLS configuration.

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
            description "This is a different description."
            disable-learning
            disable-aging
            discard-unknown
            local-age 500
            remote-age 1000
            stp
                shutdown
            exit
            sap 1/1/5:22 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

3.6.2 Modifying Management VPLS Parameters

To modify the range of VLANs on an access port that are to be managed by an existing management VPLS, first the new range should be entered and afterwards the old range removed. If the old range is removed before a new range is defined, all customer VPLS services in the old range will become unprotected and may be disabled.

CLI Syntax:

```
config>service# vpls service-id
     sap sap-id
          managed-vlan-list
               [no] range vlan-range
```

3.6.3 Deleting a Management VPLS

As with normal VPLS service, a management VPLS cannot be deleted until SAPs and SDPs are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following CLI syntax to delete a management VPLS service:

CLI Syntax:
```
config>service
[no] vpls service-id
     shutdown
     [no] spoke-sdp sdp-id
     [no] mesh-sdp sdp-id
          shutdown
     [no] sap sap-id
          shutdown
```

3.6.4 Disabling a Management VPLS

You can shut down a management VPLS without deleting the service parameters.

When a management VPLS is disabled, all associated user VPLS services are also disabled (to prevent loops). If this is not needed, first un-manage the user's VPLS service by removing them from the managed-vlan-list or moving the spoke-SDPs on to another tunnel SDP.

CLI Syntax:
```
config>service
vpls service-id
     shutdown
```

Example:
```
config>service# vpls 1
config>service>vpls# shutdown
config>service>vpls# exit
```

3.6.5 Deleting a VPLS Service

A VPLS service cannot be deleted until SAPs and SDPs are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following CLI syntax to delete a VPLS service:

CLI Syntax:
```
config>service
```

```
[no] vpls service-id
     shutdown
     [no] mesh-sdp sdp-id
          shutdown
     sap sap-id [split-horizon-group group-name]
     no sap sap-id
          shutdown
```

3.6.6 Disabling a VPLS Service

You can shut down a VPLS service without deleting the service parameters.

CLI Syntax:
```
config>service> vpls service-id
[no] shutdown
```

Example:
```
config>service# vpls 1
config>service>vpls# shutdown
config>service>vpls# exit
```

3.6.7 Re-enabling a VPLS Service

Use the following CLI syntax to re-enable a VPLS service that was shut down.

CLI Syntax:
```
config>service> vpls service-id
[no] shutdown
```

Example:
```
config>service# vpls 1
config>service>vpls# no shutdown
config>service>vpls# exit
```

3.7 VPLS Service Configuration Command Reference

This section describes the VPLS service configuration command reference.

3.7.1 Command Hierarchies

3.7.1.1 Global Commands

```
config
    — service
        — system
            — fdb-table-size table-size
            — no fdb-table-size
        — vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls] [etree]
          [name name] [create]
        — no vpls service-id
            — [no] allow-ip-int-bind
                — [no] forward-ipv4-multicast-to-ip-int
                — [no] forward-ipv6-multicast-to-ip-int
                — igmp-snooping
                    — [no] mrouter-port
                — mld-snooping
                    — [no] mrouter-port
                — [no] vxlan-ipv4-tep-ecmp
            — backbone-vpls service-id[:isid]
            — no backbone-vpls
                — [no] stp
            — bgp
                — pw-template-binding policy-id [split-horizon-group group-name]
                  [import-rt ext-community [ext-community ... (up to 5 max)]]
                — no pw-template-binding policy-id
                    — [no] bfd-enable
                    — bfd-template [name]
                    — no bfd-template
                    — monitor-oper-group group-name
                    — no monitor-oper-group
                    — oper-group group-name
                    — no oper-group
                — route-target {ext-community | export ext-community [import ext-
                  community] | import ext-community}
                — no route-target
                — route-distinguisher rd
                — no route-distinguisher
                — route-distinguisher auto-rd
```

```
            — vsi-export policy-name [policy-name...(up to 5 max)]
            — no vsi-export
            — vsi-import policy-name [policy-name...(up to 5 max)]
            — no vsi-import
        — [no] bgp-ad
            — vpls-id vpls-id
            — vsi-id
                — prefix low-order-vsi-id
                — no prefix
        — bgp-evpn
            — [no] mac-advertisement
            — mac-duplication
                — detect num-moves num-moves window minutes
                — [no] retry minutes
            — [no] unknown-mac-route
            — vxlan
                — [no] shutdown
        — bgp-vpls
            — max-ve-id value
            — no max-ve-id
            — ve-name name
            — no ve-name
                — ve-id ve-id-value
                — no ve-id
            — [no] shutdown
        — [no] def-mesh-vc-id vc-id
        — mcr-default-gtw
            — ip ip-address
            — no ip
            — mac ieee-address
            — no mac
        — description description-string
        — no description
        — [no] disable-aging
        — [no] disable-learning
        — [no] discard-unknown
        — endpoint endpoint-name [create]
        — no endpoint
            — [no] auto-learn-mac-protect
            — [no] block-on-mesh-failure
            — description description-string
            — no description
            — [no] ignore-standby-signaling
            — [no] mac-pinning
            — max-nbr-mac-addr table-size
            — no max-nbr-mac-addr
            — [no] mc-endpoint
                — mc-ep-peer name
                — mc-ep-peer ip-address
                — no mc-ep-peer
            — restrict-protected-src alarm-only
            — restrict-protected-src [discard-frame]
            — no restrict-protected-src
            — revert-time revert-time | infinite
```

```
            — no revert-time
            — static-mac ieee-address [create]
            — no static-mac
            — [no] suppress-standby-signaling
        — eth-cfm
            — [no] mep mep-id domain md-index association ma-index
                — [no] ccm-enable
                — ccm-ltm-priority priority
                — no ccm-ltm-priority
                — description description-string
                — no description
                — [no] eth-test-enable
                    — [no] test-pattern {all-zeros | all-ones} [crc-enable]
                — grace
                    — eth-ed
                        — max-rx-defect-window seconds
                        — no max-rx-defect-window
                        — priority priority
                        — no priority
                        — [no] rx-eth-ed
                        — [no] tx-eth-ed
                    — eth-vsm-grace
                        — [no] rx-eth-vsm-grace
                        — [no] tx-eth-vsm-grace
                — low-priority-defect {allDef | macRemErrXcon | remErrXcon |
                      errXcon | xcon | noXcon}
                — mac-address mac-address
                — no mac-address
                — one-way-delay-threshold seconds
                — [no] shutdown
            — tunnel-fault [accept | ignore]
        — [no] fdb-table-high-wmark high-water-mark
        — [no] fdb-table-low-wmark low-water-mark
        — fdb-table-size table-size
        — no fdb-table-size [table-size]
        — gsmp
            — [no] group name [create]
                — ancp
                    — [no] dynamic-topology-discover
                    — [no] line-configuration
                    — [no] oam
                — description description-string
                — no description
                — hold-multiplier multiplier
                — no hold-multiplier
                — keepalive seconds
                — no keepalive
                — [no] neighbor ip-address
                    — description description-string
                    — no description
                    — local-address ip-address
                    — no local-address
                    — priority-marking dscp dscp-name
                    — priority-marking prec ip-prec-value
```

```
                        — no priority-marking
                        — [no] shutdown
                    — [no] shutdown
                — [no] idle-filter
                — persistency-database
                — no persistency-database
                — [no] shutdown
            — host-connectivity-verify source-ip ip-address [source-mac ieee-address]
              [interval interval] [action {remove | alarm}]
            — igmp-host-tracking
                — expiry-time expiry-time
                — no expiry-time
                — [no] shutdown
            — igmp-snooping
                — mvr
                    — description description-string
                    — no description
                    — group-policy policy-name
                    — no group-policy
                    — [no] shutdown
                — query-interval seconds
                — no query-interval
                — query-src-ip ip-address
                — no query-src-ip
                — report-src-ip ip-address
                — no report-src-ip
                — robust-count robust-count
                — no robust-count
                — [no] shutdown
            — [no] interface ip-int-name
                — address ip-address[/mask] [netmask]
                — no address
                — arp-timeout seconds
                — no arp-timeout
                — description long-description-string
                — no description
                — hold-time
                    — up ip seconds
                    — no up ip
                    — up ipv6 seconds
                    — no up ipv6
                    — down ip seconds [init-only]
                    — no down ip
                    — down ipv6 seconds [init-only]
                    — no down ipv6
                — mac ieee-address
                — no mac
                — [no] shutdown
                — static-arp ieee-mac-addr unnumbered
                — no static-arp unnumbered]
                — unnumbered [ip-int-name | ip-address]
                — no unnumbered
            — isid-policy
                — entry range-entry-id [create]
```

```
            — no entry range-entry-id
                — [no] advertise-local
                — range isid [to isid]
                — no range
                — [no] use-def-mcast
        — load-balancing
            — [no] per-service-hashing
            — [no] spi-load-balancing
            — [no] teid-load-balancing
        — local-age aging-timer
        — no local-age [aging-time]
        — [no] mac-move
            — move-frequency frequency
            — no move-frequency
            — number-retries number-retries
            — no number-retries
            — primary-ports
                — cumulative-factor cumulative-factor
                — no cumulative-factor
                — [no] sap sap-id
                — [no] spoke-sdp spoke-id
                — [no] cumulative-factor factor
            — retry-timeout timeout
            — no retry-timeout
            — secondary-ports
                — cumulative-factor cumulative-factor
                — no cumulative-factor
                — [no] sap sap-id
                — [no] spoke-sdp spoke-id
                — [no] cumulative-factor factor
            — [no] shutdown
        — mac-notification
            — count value
            — no count
            — interval deci-seconds
            — no interval
            — renotify seconds
            — no renotify
            — [no] shutdown
        — mac-protect
            — [no] mac ieee-address
        — mac-subnet-length subnet-length
        — no mac-subnet-length
        — mcast-ipv6-snooping-scope {mac-based | sg-based}
        — no mcast-ipv6-snooping-scope
        — mfib-table-high-wmark high-water-mark
        — no mfib-table-high-wmark
        — mfib-table-low-wmark low-water-mark
        — no mfib-table-low-wmark
        — mfib-table-size table-size
        — no mfib-table-size
        — mld-snooping
            — mvr
                — description description-string
```

- **no description**
- **group-policy** *policy-name*
- **no group-policy**
- [**no**] **shutdown**

- **query-interval** *seconds*
- **no query-interval**
- **query-src-ip** *ipv6-address*
- **no query-src-ip**
- **report-src-ip** *ipv6-address*
- **no report-src-ip**
- **robust-count** *robust-count*
- **no robust-count**
- [**no**] **shutdown**

- **mrp**
  - [**no**] **attribute-table-size**
  - [**no**] **attribute-table-high-wmark**
  - [**no**] **attribute-table-low-wmark**
  - **flood-time** *flood-time*
  - **no flood-time**
  - [**no**] **shutdown**
  - **mvrp**
    - [**no**] **attribute-table-size**
    - [**no**] **attribute-table-high-wmark**
    - [**no**] **attribute-table-low-wmark**
    - **flood-time** *flood-time*
    - **no flood-time**
    - **flood-time**
    - [**no**] **hold-time** *value*
    - [**no**] **shutdown**
- **multicast-info-policy** *policy-name*
- **no multicast-info-policy**
- [**no**] **pim-snooping**
  - **group-policy** *grp-policy-name* [.. *grp-policy-name*]
  - **no group-policy**
  - **hold-time** *seconds*
  - **no hold-time**
  - [**no**] **ipv4-multicast-disable**
  - [**no**] **ipv6-multicast-disable**
  - **mode** *mode*
- **remote-age** *aging-timer*
- **no remote-age**
- [**no**] **selective-learned-fdb**
- **send-bvpls-flush** {[**all-but-mine**] | [**all-from-me**]}
- **no send-bvpls-flush**
- [**no**] **send-flush-on-failure**
- [**no**] **send-flush-on-failure**
- **service-mtu** *octets*
- **no service-mtu**
- **service-name** *service-name*
- **no service-name**
- [**no**] **shutdown**
- **site** *name* [**create**]
- **no site** *name*
  - **boot-timer** *seconds*

- **no boot-timer**
- **failed-threshold** [1 to 1000]
- **failed-threshold all**
- [**no**] **mesh-sdp-binding**
- **monitor-oper-group** *name*
- **no monitor-oper-group**
- **sap** *sap-id*
- **no sap**
- [**no**] **shutdown**
- **site-activation-timer** *seconds*
- **no site-activation-timer**
- **site-min-down-timer** *min-down-time*
- **no site-min-down-timer**
- **site-id** *value*
- **no site-id**
- **split-horizon-group** *group-name*
- **no split-horizon-group**
- **spoke-sdp** *sdp-id:vc-id*
- **no spoke-sdp**

— **spb** [*isis-instance*] [**fid** *fid*] [create]
— **no spb**

- **level** [1..1]
  - **bridge-priority** *bridge-priority*
  - **no bridge-priority**
  - **ect-algorithm fid-range** *fid-range* {**low-path-id**|**high-path-id**}
  - **no ect-algorithm fid-range fid-range**
  - **forwarding-tree-topology unicast** {**spf** | **st**}
  - **hello-interval** *seconds*
  - **no hello-interval**
  - **hello-multiplier** *multiplier*
  - **no hello-multiplier**
  - **metric** ipv4-metric
  - **no metric**
- **lsp-lifetime** *seconds*
- **no lsp-lifetime**
- **lsp-refresh-interval** [*seconds*] [**half-lifetime** {**enable** | **disable**}]
- **no lsp-refresh-interval**

— [**no**] **split-horizon-group** *group-name* [*residential-group*]

- [**no**] **auto-learn-mac-protect**
- **description** *description-string*
- **no description**
- **restrict-protected-src alarm-only**
- **restrict-protected-src** [**discard-frame**]
- **no restrict-protected-src**

— **static-mac**

- **mac** *ieee-address* [**create**] **black-hole**
- **mac** *ieee-address* [**create**] **sap** *sap-id* **monitor** {*fwd-status*}
- **mac** *ieee-address* [**create**] **spoke-sdp** *sdp-id:vc-id* **monitor** {*fwd-status*}
- **no mac** *ieee-address*

— **stp**

- **forward-delay** *forward-delay*
- **no forward-delay**
- **hello-time** *hello-time*

```
            — no hello-time
            — hold-count BDPU tx hold count
            — no hold-count
            — max-age max-info-age
            — no max-age
            — mode {rstp | comp-dot1w | dot1w | mstp | pmstp}
            — no mode
            — [no] mst-instance mst-inst-number
                — mst-priority bridge-priority
                — no mst-priority
                — [no] vlan-range vlan-range
            — mst-max-hops hops-count
            — no mst-max-hops
            — mst-name region-name
            — no mst-name
            — mst-revision revision-number
            — no mst-revision
            — priority bridge-priority
            — no priority
            — [no] shutdown
        — timers
            — lsp-wait lsp-wait [lsp-initial-wait initial-wait] [lsp-second-wait second-
              wait]
            — no lsp-wait
            — spf-wait spf-wait [spf-initial-wait initial-wait] [spf-second-wait second-
              wait]
            — no spf-wait
        — vpls-group id
            — service-range startid-endid [vlan-id startvid]
            — vpls-template-binding name/id
            — no vpls-template-binding
            — sap-template-binding name/id
            — no sap-template-binding
            — [no] mvrp-control
        — vxlan vni vni-id create
        — no vxlan vni vni-id
            — network
                — ingress
                    — qos network-policy-id fp-redirect-group queue-group-
                      name instance instance-id
                    — no qos
            — restrict-protected-src discard-frame
            — no restrict-protected-src
```

3.7.1.2 Oper Group Commands

```
config
    — service
        — vpls service-id
            — [no] interface ip-int-name
                — monitor-oper-group name
```

— **no monitor-oper-group**

3.7.1.3 SAP Commands

config
— **service**
— **vpls** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**m-vpls**] [**b-vpls** | **i-vpls**] [**etree**] [**create**]
— **no vpls** *service-id*
— **sap** *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**] [root-leaf-tag **leaf-tag-vid** | leaf-ac]
— [**no**] **sap** *eth-tunnel-tunnel-id* [*:eth-tunnel-sap-id*]
— **no sap** *sap-id*
— **accounting-policy** *acct-policy-id*
— **no accounting-policy**
— [**no**] **auto-learn-mac-protect**
— **anti-spoof** {**ip** | **mac** | **ip-mac**}
— **no anti-spoof**
— **app-profile** *app-profile-name*
— **no app-profile**
— **arp-host**
— **host-limit** *max-num-hosts*
— **no host-limit**
— **min-auth-interval** *min-auth-interval*
— **no min-auth-interval**
— [**no**] **shutdown**
— **arp-reply-agent** [**sub-ident**]
— **no arp-reply-agent**
— **atm**
— **egress**
— **traffic-desc** *traffic-desc-profile-id*
— **no traffic-desc**
— **encapsulation** *atm-encap-type*
— **ingress**
— **traffic-desc** *traffic-desc-profile-id*
— **no traffic-desc**
— **oam**
— [**no**] **alarm-cells**
— **authentication-policy** *name*
— **no authentication-policy**
— **bpdu-translation** {**auto** | **auto-rw** | **pvst** | **pvst-rw** | **stp**}
— **no bpdu-translation**
— **calling-station-id** {**mac** | **remote-id** | **sap-id** | **sap-string**}
— **no calling-station-id**
— [**no**] **cflowd**
— [**no**] **collect-stats**
— **cpu-protection** *policy-id* {[**mac-monitoring**] | [**eth-cfm-monitoring** [**aggregate**] [**car**]]}
— **no cpu-protection**
— **default-msap-policy** *policy-name*
— **no default-msap-policy**

```
— description description-string
— no description
— dhcp
    — description description-string
    — no description
    — lease-populate [nbr-of-entries]
    — no lease-populate
    — [no] option
        — action [dhcp-action]
        — no action
        — circuit-id
        — circuit-id [ascii-tuple | vlan-ascii-tuple]
        — circuit-id hex [0x0..0xFFFFFFFF...(64 hex nibbles)]
        — no circuit-id
        — remote-id
        — remote-id [mac | string string]
        — remote-id hex [0x0..0xFFFFFFFF...(64 hex nibbles)]
        — no remote-id
        — [no] vendor-specific-option
            — [no] client-mac-address
            — [no] sap-id
            — [no] service-id
            — string text
            — no string
            — [no] system-id
    — proxy-server
        — emulated-server ip-address
        — no emulated-server
        — lease-time [days days] [hrs hours] [min minutes] [sec
            seconds] [radius-override]
        — no lease-time
        — [no] shutdown
    — [no] shutdown
    — [no] snoop
— dhcp-python-policy policy-name
— no dhcp-python-policy
— dhcp6-user-db local-user-db-name
— no dhcp6-user-db
— dhcp6
    — description description-string
    — no description
    — [no] option
        — interface-id
        — interface-id ascii-tuple
        — interface-id vlan-ascii-tuple
        — no interface-id
        — remote-id
        — remote-id mac
        — remote-id string [32 chars max]
        — no remote-id
    — [no] shutdown
    — [no] snoop
— [no] disable-aging
— [no] disable-learning
```

```
— [no] discard-unknown-source
— dist-cpu-protection policy-name
— no dist-cpu-protection
— egress
    — [no] agg-rate
        — [no] limit-unused-bandwidth
        — [no] queue-frame-based-accounting
        — rate {max | rate}
        — no rate
    — dest-mac-rewrite
    — no dest-mac-rewrite
    — encap-defined-qos
            — encap-group group-name [type group-type]
              [qos-per-member] [create]
            — no encap-group group-name
                — [no] agg-rate
                    — [no] limit-unused-bandwidth
                    — [no] queue-frame-based-
                      accounting
                    — rate {max | rate}
                    — no rate
                — [no] member encap-id [to encap-id]
                — qos policy-id
                — no qos
                — scheduler-policy scheduler-policy-name
                — no scheduler-policy
    — filter ip ip-filter-id
    — filter ipv6 ipv6-filter-id
    — filter mac mac-filter-id
    — no filter
    — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
    — [no] hsmda-queue-override
        — secondary-shaper secondary-shaper-name
        — no secondary-shaper
        — wrr-policy hsmda-wrr-policy-name
        — no wrr-policy
        — packet-byte-offset {add add-bytes | subtract sub-bytes}
        — no packet-byte-offset
        — queue queue-id
        — no queue queue-id
            — wrr-weight weight
            — no wrr-weight
            — mbs size {[bytes | kilobytes] | default}
            — no mbs
            — rate pir-rate
            — no rate
            — slope-policy hsmda-slope-policy-name
              allowable
            — no slope-policy
    — policer-control-override [create]
    — no policer-control-override
        — max-rate {rate | max}
        — priority-mbs-thresholds
            — min-thresh-separation size [bytes | kilobytes]
```

```
                    — [no] priority level
                        — mbs-contribution size [bytes | kilobytes]
            — policer-control-policy policy-name
            — no policer-control-policy
            — [no] policer-override
                — policer policer-id [create]
                — no policer policer-id
                    — cbs size [bytes | kilobytes]
                    — no cbs
                    — mbs {size [bytes | kilobytes] | default}
                    — no mbs
                    — packet-byte-offset {add add-bytes | subtract
                        sub-bytes}
                    — rate {rate | max} [cir {max | rate}]
                    — stat-mode stat-mode
                    — no stat-mode
            — [no] qinq-mark-top-only
            — qos policy-id [port-redirect-group queue-group-name instance
                instance-id]]
            — no qos
            — [no] queue-override
                — [no] queue queue-id
                    — adaptation-rule [pir adaptation-rule] [cir
                        adaptation-rule]
                    — no adaptation-rule
                    — avg-frame-overhead percentage
                    — no avg-frame-overhead
                    — burst-limit {default | size [bytes | kilobytes]}
                    — no burst-limit
                    — cbs size-in-kbytes
                    — no cbs
                    — drop-tail
                        — low
                            — percent-reduction-from-mbs
                                percent
                            — no percent-reduction-from-mbs
                    — mbs {size [bytes | kilobytes] | default}
                    — no mbs
                    — parent {[weight weight] [cir-weight cir-weight]}
                    — no parent
                    — percent-rate pir-percent [cir cir-percent]
                    — no percent-rate
                    — rate pir-rate [cir cir-rate]
                    — no rate
            — [no] scheduler-override
                — [no] scheduler scheduler-name
                    — parent [weight weight] [cir-weight cir-weight]
                    — no parent
                    — rate pir-rate [cir cir-rate]
                    — no rate
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
    — eth-cfm
        — [no] collect-lmm-stats
```

```
— collect-lmm-fc-stats
    — fc fc-name [fc-name ... (up to 8 max)]
    — no fc
    — fc-in-profile fc-name [fc-name ... (up to 8 max)]
    — no fc-in-profile
— mep mep-id domain md-index association ma-index [direction
      {up |down}] primary-vlan-enable [vlan vlan-id]
— no mep mep-id domain md-index association ma-index
    — [no] ais-enable
            — [no] interface-support-enable
        — [no] interface-support-enable
        — ccm-padding-size ccm-padding
        — no ccm-padding-size ccm-padding
        — [no] csf-enable
            — multiplier multiplier-value
            — no multiplier
        — client-meg-level [level [level...]]
        — no client-meg-level
        — [no] description
        — interval {1 | 60}
        — no interval
        — priority priority-value
        — no priority
    — [no] ccm-enable
    — ccm-ltm-priority priority
    — no ccm-ltm-priority
    — [no] eth-test-enable
        — test-pattern {all-zeros | all-ones} [crc-enable]
        — no test-pattern
    — fault-propagation-enable {use-if-tlv | suspend-ccm}
    — no fault-propagation-enable
    — grace
        — eth-ed
            — max-rx-defect-window seconds
            — no max-rx-defect-window
            — priority priority
            — no priority
            — [no] rx-eth-ed
            — [no] tx-eth-ed
        — eth-vsm-grace
            — [no] rx-eth-vsm-grace
            — [no] tx-eth-vsm-grace
    — [no] lbm-svc-act-responder
    — low-priority-defect {allDef | macRemErrXcon |
        remErrXcon | errXcon | xcon | noXcon}
    — mac-address mac-address
    — no mac-address
    — one-way-delay-threshold seconds
    — [no] shutdown
— [no] squelch-ingress-levels [md-level [md-level…]]
— [no] mip [{mac mac-address | default-mac}]
— tunnel-fault [accept | ignore]
— [no] vmep-extensions
— [no] vmep-filter
```

- **eth-tunnel**
  - **path** *path-index* **tag** *qtag*[.*qtag*]
  - **no path** *path-index*
  - [**no**] **mip**
- **fault-propagation-bmac** [*mac-name* | *ieee-address*] [**create**]
- **no fault-propagation-bmac** [*mac-name* | *ieee-address*
- [**no**] **force-l2pt-boundary**
- **frame-relay**
  - [**no**] **frf-12**
    - **ete-fragment-threshold** *threshold*
    - **no ete-fragment-threshold**
    - [**no**] **interleave**
  - **scheduling-class** *class-id*
  - **no scheduling-class**
- **host-connectivity-verify source-ip** *ip-address* [**source-mac** *ieee-address*] [**interval** *interval*] [**action** {**remove** | **alarm**}]
- **igmp-host-tracking**
  - [**no**] **disable-router-alert-check**
  - **expiry-time** *expiry-time*
  - **no expiry-time**
  - **import** *policy-name*
  - **no import**
  - **max-num-groups** *max-num-groups*
  - **no max-num-groups**
  - **max-num-sources** *max-num-sources*
  - **no max-num-sources**
  - **max-num-grp-sources** [1..32000]
  - **no max-num-grp-sources**
- **igmp-snooping**
  - [**no**] **disable-router-alert-check**
  - [**no**] **fast-leave**
  - **import** *policy-name*
  - **no import**
  - **last-member-query-interval** *interval*
  - **no last-member-query-interval**
  - **max-num-groups** *max-num-groups*
  - **no max-num-groups**
  - **max-num-sources** *max-num-sources*
  - **no max-num-sources**
  - **max-num-grp-sources** [1 to 32000]
  - **no max-num-grp-sources**
  - **mcac**
    - **if-policy** *mcac-if-policy-name*
    - **no if-policy**
    - **mc-constraints**
      - **level** *level-id* **bw** *bandwidth*
      - **no level** *level-id*
      - **number-down** *number-lag-port-down* **level** *level-id*
      - **no number-down**
      - [**no**] **shutdown**
      - **no use-lag-port-weight**
    - **policy** *policy-name*
    - **no policy**

```
                    — unconstrained-bw bandwidth mandatory-bw
                        mandatory-bw
                    — no unconstrained-bw
                — [no] mrouter-port
                — mvr
                    — from-vpls vpls-id
                    — no from-vpls
                    — to-sap sap-id
                    — no to-sap
                — query-interval interval
                — no query-interval
                — query-response-interval interval
                — no query-response-interval
                — robust-count count
                — no robust-count
                — [no] send-queries
                — static
                    — [no] group group-address
                        — [no] source ip-address
                        — [no] starg
                — version version
                — no version
            — ingress
                — filter ip ip-filter-id
                — filter ipv6 ipv6-filter-id
                — filter mac mac-filter-id
                — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
                — match-qinq-dot1p {top | bottom}
                — no match-qinq-dot1p de
                — policer-control-override [create]
                — no policer-control-override
                    — max-rate {rate | max}
                    — priority-mbs-thresholds
                        — min-thresh-separation size [bytes | kilobytes]
                        — [no] priority level
                            — mbs-contribution size [bytes | kilobytes]
                — policer-control-policy policy-name
                — no policer-control-policy
                — [no] policer-override
                    — policer policer-id [create]
                    — no policer policer-id
                        — cbs size [bytes | kilobytes]
                        — no cbs
                        — mbs {size [bytes | kilobytes] | default}
                        — no mbs
                        — packet-byte-offset {add add-bytes | subtract
                            sub-bytes}
                        — rate {rate | max} [cir {max | rate}]
                        — stat-mode stat-mode
                        — no stat-mode
                — qos policy-id [shared-queuing | multipoint-shared] [fp-redirect-
                    group queue-group-name instance instance-id]
                — no qos
                — [no] queue-override
```

```
— [no] queue queue-id
    — adaptation-rule [pir {max | min | closest}] [cir
        {max | min | closest}]
    — no adaptation-rule
    — cbs size-in-kbytes
    — no cbs
    — drop-tail
        — low
            — percent-reduction-from-mbs
                percent
            — no percent-reduction-from-mbs
    — mbs {size [bytes | kilobytes] | default}
    — no mbs
    — parent {[weight weight] [cir-weight cir-weight]}
    — no parent
    — percent-rate pir-percent [cir cir-percent]
    — no percent-rate
    — rate pir-rate [cir cir-rate]
    — no rate
— [no] scheduler-override
    — [no] scheduler scheduler-name
        — parent [weight weight] [cir-weight cir-weight]
        — no parent
        — rate pir-rate [cir cir-rate]
        — no rate
    — scheduler-policy scheduler-policy-name
    — no scheduler-policy
— l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp]
— no l2pt-termination
— lag-link-map-profile link-map-profile-id
— no lag-link-map-profile
— lag-per-link-hash class {1 | 2 | 3} weight [1 to 1024]
— no lag-per-link-hash
— leaf-ac
— limit-mac-move [blockable | non-blockable]
— no limit-mac-move
— [no] mac-pinning
— managed-vlan-list
    — [no] default-sap
    — [no] range vlan-range
— max-nbr-mac-addr table-size
— no max-nbr-mac-addr
— mld-snooping
    — [no] disable-router-alert-check
    — [no] fast-leave
    — import policy-name
    — no import
    — last-member-query-interval interval
    — no last-member-query-interval
    — max-num-groups max-num-groups
    — no max-num-groups
    — mcac
        — if-policy mcac-if-policy-name
        — no if-policy
```

```
                        — mc-constraints
                            — level level-id bw bandwidth
                            — no level level-id
                            — number-down number-lag-port-down level level-id
                            — no number-down
                            — [no] shutdown
                            — no use-lag-port-weight
                        — policy policy-name
                        — no policy
                        — unconstrained-bw bandwidth mandatory-bw mandatory-bw
                        — no unconstrained-bw
                    — mvr
                        — fast-leave
                        — no fast-leave
                        — to-sap sap-id
                        — no to-sap
                    — query-interval seconds
                    — no query-interval
                    — query-response-interval seconds
                    — no query-response-interval
                    — robust-count robust-count
                    — no robust-count
                    — [no] send-queries
                    — static
                        — [no] group group-address
                            — [no] source ip-address
                            — [no] starg
                    — version version
                    — no version
                — mrp
                    — [no] join-time value
                    — [no] leave-all-time value
                    — [no] leave-time value
                    — [no] mrp-policy policy-name
                    — [no] periodic-time value
                    — [no] periodic-time
                    — mvrp
                        — endstation-vid-group id vlan-id startvid-endvid
                        — no endstation-vid-group id
                        — [no] shutdown
                — msap-defaults
                    — [no] service service-id
                    — policy msap-policy-name
                    — no policy
                — monitor-oper-group group-name
                — no monitor-oper-group
                — oper-group group-name
                — no oper-group
                — multi-service-sitecustomer-site-name
                — no multi-service-site
                — pim-snooping
                    — max-num-groups num-groups
```

- — **no max-num-groups**
- — [**no**] **process-cpm-traffic-on-sap-down**
- — **pppoe-policy** *pppoe-policy-name*
- — **no pppoe-policy**
- — **restrict-protected-src alarm-only**
- — **restrict-protected-src** [**discard-frame**]
- — **no restrict-protected-src**
- — **restrict-unprotected-dst**
- — **no restrict-unprotected-dst**
- — **shcv-policy-ipv4** *policy-name*
- — **no shcv-policy-ipv4**
- — [**no**] **shutdown**
- — **spb** [*isis-instance*] [**fid** *fid*] [create]
- — **no spb**
  - — **level** [1 to 1]
    - — **hello-interval** *seconds*
    - — **no hello-interval**
    - — **hello-multiplier** *multiplier*
    - — **no hello-multiplier**
    - — **metric** *ipv4-metric*
    - — **no metric**
  - — **lsp-pacing-interval** *milli-seconds*
  - — **no lsp-pacing-interval**
  - — **retransmit-interval** seconds
  - — **no retransmit-interval**
  - — [**no**] **shutdown**
- — **static-host ip** *ip-address* [**mac** *ieee-address*] [**create**]
- — **static-host mac** *ieee-address* [**create**]
- — **no static-host** [**ip** *ip-address*] **mac** *ieee-address*
- — **no static-host all** [**force**]
- — **no static-host ip** *ip-address*
  - — **ancp-string** *ancp-string*
  - — **no ancp-string**
  - — **app-profile** *app-profile-name*
  - — **no app-profile**
  - — **inter-dest-id** *intermediate-destination-id*
  - — **no inter-dest-id**
  - — [**no**] **shutdown**
  - — **sla-profile** *sla-profile-name*
  - — **no sla-profile**
  - — **sub-profile** *sub-profile-name*
  - — **no sub-profile**
  - — **subscriber** *sub-ident*
  - — **no subscriber**
  - — [**no**] **subscriber-sap-id**
- — [**no**] **static-isid range** *entry-id isid* [**to** *isid*] [**create**]
- — [**no**] **static-mac** *ieee-address*
- — **stp**
  - — [**no**] **auto-edge**
  - — [**no**] **edge-port**
  - — **link-type** {**pt-pt** | **shared**}
  - — **no link-type**
  - — **mst-instance** *mst-inst-number*
    - — **mst-path-cost** *inst-path-cost*

```
                        — no mst-path-cost
                        — mst-priority bridge-priority
                        — no mst-priority
                    — path-cost sap-path-cost
                    — no path-cost
                    — [no] port-num virtual-port-number
                    — priority stp-priority
                    — no priority
                    — [no] vpls-group
                    — [no] shutdown
                — [no] sub-sla-mgmt
                    — def-sla-profile default-sla-profile-name
                    — no def-sla-profile
                    — def-sub-profile default-subscriber-profile-name
                    — no def-sub-profile
                    — [no] mac-da-hashing
                    — multi-sub-sap [subscriber-limit]
                    — [no] shutdown
                    — single-sub-parameters
                        — non-sub-traffic sub-profile sub-profile-name sla-
                            profile sla-profile-name [subscriber sub-ident-string]
                        — no non-sub-traffic
                        — [no] profiled-traffic-only
                    — sub-ident-policy sub-ident-policy-name
                    — no sub-ident-policy
                — trigger-packet [dhcp] [pppoe] [arp] [dhcp6] [ppp] [data]
                — no trigger-packet
```

3.7.1.4 Template Commands

```
config
    — service
        — template
            — vpls-template name/id create
                — [no] temp-flooding flood-time
                — [no] disable-aging
                — [no] disable-learning
                — [no] discard-unknown
                — [no] fdb-table-high-wmark high-water-mark
                — [no] fdb-table-low-wmark low-water-mark
                — fdb-table-size table-size
                — no fdb-table-size [table-size]
                — local-age aging-timer
                — load-balancing
                    — [no] per-service-hashing
                    — [no] spi-load-balancing
                    — [no] teid-load-balancing
                — no local-age
                — [no] mac-move
                    — move-frequency frequency
                    — no move-frequency
```

```
                    — number-retries number-retries
                    — no number-retries
                    — primary-ports
                        — cumulative-factor cumulative-factor
                        — no cumulative-factor
                    — retry-timeout timeout
                    — no retry-timeout
                    — secondary-ports
                        — cumulative-factor cumulative-factor
                        — no cumulative-factor
                    — [no] shutdown
                — [no] per-service-hashing
                — remote-age aging-timer
                — no remote-age
                — service-mtu octets
                — no service-mtu
                — stp
                    — forward-delay forward-delay
                    — no forward-delay
                    — hello-time hello-time
                    — no hello-time
                    — hold-count BDPU tx hold count
                    — no hold-count
                    — max-age max-info-age
                    — no max-age
                    — mode {rstp | comp-dot1w | dot1w | mstp | pmstp}
                    — no mode
                    — priority bridge-priority
                    — no priority
                    — [no] shutdown
            — vpls-sap-template name/id create
                — l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp]
                — no l2pt-termination
                — bpdu-translation {auto | auto-rw | pvst | pvst-rw | stp}
                — no bpdu-translation
                — [no] collect-stats
                — cpu-protection policy-id [mac-monitoring]
                — no cpu-protection
                — eth-cfm
                    — [no] mip primary-vlan-enable [vlan vlan-id]
                    — [no] squelch-ingress-levels [md-level [md-level…]]
                — [no] disable-aging
                — [no] disable-learning
                — [no] discard-unknown-source
                — egress
                    — [no] agg-rate
                        — rate {max | rate}
                        — no rate
                        — [no] limit-unused-bandwidth
                        — [no] queue-frame-based-accounting
                    — filter ip ip-filter-id
                    — filter ipv6 ipv6-filter-id
                    — filter mac mac-filter-id
                    — no filter
```

- no filter [ip *ip-filter-id*] [mac *mac-filter-id*] [ipv6 *ipv6-filter-id*]
- policer-control-policy *policy-name*
- no policer-control-policy
- [no] qinq-mark-top-only
- qos *policy-id* [shared-queuing | multipoint-shared]
- no qos
- scheduler-policy *scheduler-policy-name*
- no scheduler-policy

ingress
- filter ip *ip-filter-id*
- filter ipv6 *ipv6-filter-id*
- filter mac *mac-filter-id*
- no filter [ip *ip-filter-id*] [mac *mac-filter-id*] [ipv6 *ipv6-filter-id*]
- match-qinq-dot1p {top | bottom}
- no match-qinq-dot1p de
- policer-control-policy *policy-name*
- no policer-control-policy
- qos *policy-id* [*shared-queuing* | *multipoint-shared*]
- no qos
- scheduler-policy *scheduler-policy-name*
- no scheduler-policy

limit-mac-move [blockable | non-blockable]
no limit-mac-move
max-nbr-mac-addr *table-size*
no max-nbr-mac-addr
stp
- [no] auto-edge
- [no] edge-port
- link-type {pt-pt | shared}
- no link-type [pt-pt | shared]
- path-cost *sap-path-cost*
- no path-cost
- priority *stp-priority*
- no priority
- [no] vpls-group
- [no] shutdown

[no] mac-move-level

3.7.1.5 Mesh SDP Commands

config
- service
  - [no] vpls *service-id* [customer *customer-id*] [vpn *vpn-id*] [mvpls] [b-vpls | i-vpls] [etree] [create]
    - mesh-sdp *sdp-id*[:*vc-id*] [vc-type {ether | vlan}] [root-leaf-tag | leaf-ac]
    - no mesh-sdp *sdp-id*[:*vc-id*]
      - accounting-policy *acct-policy-id*
      - no accounting-policy
      - [no] auto-learn-mac-protect
      - [no] bfd-enable

```
— bfd-template name
— no bfd-template
— [no] collect-stats
— [no] control-word
— description description-string
— no description
— dhcp
    — description description-string
    — no description
    — [no] snoop
— egress
    — filter ip ip-filter-id
    — filter ipv6 ipv6-filter-id
    — filter mac mac-filter-id
    — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
    — qos network-policy-id port-redirect-group queue-group-name
          [instance instance-id]
    — no qos
    — mfib-allowed-mda-destinations
        — [no] mda mda-id
    — vc-label egress-vc-label
    — no vc-label [egress-vc-label]
— [no] entropy-label
— eth-cfm
    — [no] collect-lmm-stats
    — collect-lmm-fc-stats
        — fc fc-name [fc-name ... (up to 8 max)]
        — no fc
        — fc-in-profile fc-name [fc-name ... (up to 8 max)]
        — no fc-in-profile
    — mep mep-id domain md-index association ma-index [direction
          {up | down}] [primary-vlan-enable]
    — no mep mep-id domain md-index association ma-index
          primary-vlan-enable [vlan vlan-id]
        — [no] ais-enable
            — client-meg-level [level [level...]]
            — no client-meg-level
            — [no] interface-support-enable
            — interval {1 | 60}
            — no interval
            — low-priority-defect {allDef | macRemErrXcon}
            — priority priority-value
            — no priority
        — [no] ccm-enable
        — ccm-padding-size ccm-padding
        — no ccm-padding-size
        — ccm-ltm-priority priority
        — no ccm-ltm-priority
        — [no] csf-enable
                — multiplier multiplier-value
                — no multiplier
        — description description-string
        — no description
        — [no] eth-test-enable
```

```
                        — test-pattern {all-zeros | all-ones} [crc-enable]
                        — no test-pattern
                    — fault-propagation-enable {use-if-tlv | suspend-ccm}
                    — no fault-propagation-enable
                    — grace
                        — eth-ed
                            — max-rx-defect-window seconds
                            — no max-rx-defect-window
                            — priority priority
                            — no priority
                            — [no] rx-eth-ed
                            — [no] tx-eth-ed
                        — eth-vsm-grace
                            — [no] rx-eth-vsm-grace
                            — [no] tx-eth-vsm-grace
                    — [no] lbm-svc-act-responder
                    — low-priority-defect {allDef | macRemErrXcon |
                        remErrXcon | errXcon | xcon | noXcon}
                    — mac-address mac-address
                    — no mac-address
                    — [no] shutdown
                — [no] mip [mac mac-address] [primary-vlan-enable vlan-id]
                — [no] squelch-ingress-levels [md-level [md-level...]]
                — [no] vmep-filter
            — fault-propagation-bmac [mac-name | ieee-address] [create]
            — no fault-propagation-bmac [mac-name | ieee-address]
            — [no] force-qinq-vc-forwarding
            — [no] force-vlan-vc-forwarding
            — [no] hash-label
            — igmp-snooping
                — [no] disable-router-alert-check
                — [no] fast-leave
                — import policy-name
                — no import
                — last-member-query-interval interval
                — no last-member-query-interval
                — max-num-groups max-num-groups
                — no max-num-groups
                — max-num-grp-sources [1 to 32000]
                — no max-num-grp-sources
                — mcac
                    — if-policy mcac-if-policy-name
                    — no if-policy
                    — policy policy-name
                    — no policy
                    — unconstrained-bw bandwidth mandatory-bw
                        mandatory-bw
                    — no unconstrained-bw
                — [no] mrouter-port
                — query-response-interval interval
                — no query-response-interval
                — robust-count count
                — no robust-count
                — [no] send-queries
```

```
            — static
                — [no] group grp-ip-address
                    — [no] source ip-address
                    — [no] starg
            — version version
            — no version
        — ingress
            — filter ip ip-filter-id
            — filter ipv6 ipv6-filter-id
            — filter mac mac-filter-id
            — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
            — qos network-policy-id fp-redirect-group queue-group-name
                  instance instance-id
            — no qos
            — mfib-allowed-mda-destinations
                — [no] mda mda-id
            — vc-label ingress-vc-label
            — no vc-label [ingress-vc-label]
        — [no] mac-pinning
        — mld-snooping
            — [no] disable-router-alert-check
            — [no] fast-leave
            — import policy-name
            — no import
            — last-member-query-interval interval
            — no last-member-query-interval
            — max-num-groups max-num-groups
            — no max-num-groups
            — mcac
                — if-policy mcac-if-policy-name
                — no if-policy
                — policy policy-name
                — no policy
                — unconstrained-bw bandwidth mandatory-bw
                    mandatory-bw
                — no unconstrained-bw
            — mvr
                — [no] fast-leave
                — to-sap sap-id
                — no to-sap
            — query-interval seconds
            — no query-interval
            — query-response-interval seconds
            — no query-response-interval
            — robust-count robust-count
            — no robust-count
            — [no] send-queries
            — static
                — [no] group
                    — [no] source src-ipv6-address
                    — [no] starg
                — [no] group grp-ipv6-address
            — version version
            — no version
```

```
— mrp
    — [no] join-time value
    — [no] leave-all-time value
    — [no] leave-time value
    — [no] mrp-policy policy-name
    — [no] periodic-time value
    — [no] periodic-time
— restrict-protected-src alarm-only
— restrict-protected-src [discard-frame]
— no restrict-protected-src
— [no] shutdown
— [no]static-mac ieee-address
— vlan-vc-tag 0 to 4094
— no vlan-vc-tag [0 to 4094]
```

3.7.1.6 Spoke SDP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [mvpls] [b-vpls | i-vpls]
          [etree] [create]
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-
              name] [root-leaf-tag | leaf-ac]
            — no spoke-sdp sdp-id[:vc-id]
                — accounting-policy acct-policy-id
                — no accounting-policy
                — app-profile app-profile-name
                — no app-profile
                — [no] auto-learn-mac-protect
                — [no] bfd-enable
                — bfd-template name
                — no bfd-template
                — [no] block-on-mesh-failure
                — bpdu-translation {auto | auto-rw | pvst | pvst-rw | stp}
                — no bpdu-translation
                — [no] collect-stats
                — [no] control-channel-status
                    — [no] acknowledgment
                    — refresh-timer value
                    — no refresh-timer
                    — request-timer timer1 retry-timer timer2 [timeout-multiplier
                      multiplier]
                    — no request-timer
                — [no] control-word
                — description description-string
                — no description
                — dhcp
                    — description description-string
                    — no description
                    — [no] snoop
                — [no] disable-aging
```

- [no] disable-learning
- [no] discard-unknown-source
- [no] entropy-label
- eth-cfm
  - [no] collect-lmm-stats
  - collect-lmm-fc-stats
    - fc *fc-name* [*fc-name* ... (up to 8 max)]
    - no fc
    - fc-in-profile *fc-name* [*fc-name* ... (up to 8 max)]
    - no fc-in-profile
  - mep *mep-id* domain *md-index* association *ma-index* [direction {up | down}] [primary-vlan-enable]
  - no mep *mep-id* domain *md-index* association *ma-index*
    - [no] ais-enable
      - [no] interface-support-enable
    - client-meg-level [*level* [*level*...]]
    - no client-meg-level
    - interval {1 | 60}
    - no interval
    - priority *priority-value*
    - no priority
    - [no] ccm-enable
    - ccm-ltm-priority *priority*
    - no ccm-ltm-priority
    - ccm-padding-size *ccm-padding*
    - no ccm-padding-size *ccm-padding*
    - [no] csf-enable
      - multiplier *multiplier-value*
      - no multiplier
    - [no] description
    - [no] eth-test-enable
      - test-pattern {all-zeros | all-ones} [crc-enable]
      - no test-pattern
    - fault-propagation-enable {use-if-tlv | suspend-ccm}
    - no fault-propagation-enable
    - grace
      - eth-ed
        - max-rx-defect-window *seconds*
        - no max-rx-defect-window
        - priority *priority*
        - no priority
        - [no] rx-eth-ed
        - [no] tx-eth-ed
      - eth-vsm-grace
        - [no] rx-eth-vsm-grace
        - [no] tx-eth-vsm-grace
    - [no] lbm-svc-act-responder
    - low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
    - mac-address *mac-address*
    - no mac-address
    - [no] description
    - [no] shutdown
  - [no] mip [mac *mac-address*] [primary-vlan-enable *vlan-id*]

```
            — [no] squelch-ingress-levels [md-level [md-level...]]
            — vmep-filter
        — egress
            — filter ip ip-filter-id
            — filter ipv6 ipv6-filter-id
            — filter mac mac-filter-id
            — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
            — qos network-policy-id port-redirect-group queue-group-name
                  [instance instance-id]
            — no qos
            — mfib-allowed-mda-destinations
                — [no] mda mda-id
            — vc-label egress-vc-label
            — no vc-label [egress-vc-label]
        — fault-propagation-bmac [mac-name | ieee-address] [create]
        — no fault-propagation-bmac [mac-name | ieee-address
        — [no] force-vlan-vc-forwarding
        — hash-label
        — no hash-label
        — igmp-snooping
            — [no] disable-router-alert-check
            — [no] fast-leave
            — import policy-name
            — no import
            — last-member-query-interval interval
            — no last-member-query-interval
            — max-num-groups max-num-groups
            — no max-num-groups
            — max-num-grp-sources [1 to 32000]
            — no max-num-grp-sources
            — mcac
                — if-policy mcac-if-policy-name
                — no if-policy
                — policy policy-name
                — no policy
                — unconstrained-bw bandwidth mandatory-bw
                      mandatory-bw
                — no unconstrained-bw
            — [no] mrouter-port
            — query-interval interval
            — no query-interval
            — query-response-interval interval
            — no query-response-interval
            — robust-count count
            — no robust-count
            — [no] send-queries
            — static
                — [no] group group-address
                    — [no] source ip-address
                    — [no] starg
            — version version
            — no version
        — [no] ignore-standby-signaling
        — ingress
```

```
            — filter ip ip-filter-id
            — filter ipv6 ipv6-filter-id
            — filter mac mac-filter-id
            — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
            — qos network-policy-id fp-redirect-group queue-group-name
                  instance instance-id
            — no qos
            — mfib-allowed-mda-destinations
                  — [no] mda mda-id
            — vc-label egress-vc-label
            — no vc-label [egress-vc-label]
      — [no] l2pt-termination
      — limit-mac-move [blockable | non-blockable]
      — no limit-mac-move
      — [no] mac-pinning
      — max-nbr-mac-addr table-size
      — no max-nbr-mac-addr
      — mld-snooping
            — [no] disable-router-alert-check
            — [no] fast-leave
            — import policy-name
            — no import
            — last-member-query-interval interval
            — no last-member-query-interval
            — max-num-groups max-num-groups
            — no max-num-groups
            — mcac
                  — if-policy mcac-if-policy-name
                  — no if-policy
                  — policy policy-name
                  — no policy
                  — unconstrained-bw bandwidth mandatory-bw
                        mandatory-bw
                  — no unconstrained-bw
            — query-interval seconds
            — no query-interval
            — query-response-interval seconds
            — no query-response-interval
            — robust-count robust-count
            — no robust-count
            — [no] send-queries
            — static
                  — [no] group group-address
                        — [no] source ip-address
                        — [no] starg
            — version version
            — no version
      — monitor-oper-group group-name
      — no monitor-oper-group
      — oper-group group-name
      — no oper-group
      — mrp
            — [no] join-time value
            — [no] leave-all-time value
```

- [no] leave-time *value*
- [no] mrp-policy *policy-name*
- [no] periodic-time *value*

- oper-group *group-name*
- no oper-group
- pim-snooping
  - max-num-groups *num-groups*
  - no max-num-groups
- precedence *precedence-value* | primary
- no precedence
- [no] pw-path-id
  - agi agi
  - no agi
  - saii-type2 *global-id*:*node-id*:*ac-id*
  - no saii-type2
  - taii-type2 *global-id*:*node-id*:*ac-id*
  - no taii-type2
- [no] pw-status-signaling
- propagate-mac-flush [*precedence-value* | primary]
- no propagate-mac-flush
- [no] shutdown
- spb [*isis-instance*] [fid *fid*] [create]
- no spb
  - level [1 to 1]
    - hello-interval *seconds*
    - no hello-interval
    - hello-multiplier *multiplier*
    - no hello-multiplier
    - metric *ipv4-metric*
    - no metric
  - lsp-pacing-interval *milli-seconds*
  - no lsp-pacing-interval
  - retransmit-interval seconds
  - no retransmit-interval
  - [no] shutdown
- [no] static-isid range *entry-id isid* [to *isid*] [create]
- [no] static-mac *ieee-address*
- stp
  - [no] auto-edge
  - [no] edge-port
  - link-type {pt-pt | shared}
  - no link-type
  - path-cost *sap-path-cost*
  - no path-cost
  - [no] port-num *virtual-port-number*
  - priority *stp-priority*
  - no priority
  - [no] vpls-group
  - [no] shutdown
- transit-policy prefix *prefix-aasub-policy-id*
- no transit-policy
- vlan-vc-tag *0 to 4094*
- no vlan-vc-tag [*0 to 4094*]

3.7.1.7 Provider Tunnel Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [mvpls] [b-vpls | i-vpls]
          [create]
            — provider-tunnel
                — inclusive
                    — data-delay-interval seconds
                    — no data-delay-interval
                    — mldp
                    — [no] mldp
                    — [no] root-and-leaf
                    — [no] rsvp
                        — lsp-template p2mp-lsp-template-name
                        — no lsp-template
                    — [no] shutdown
```

3.7.1.8 Routed VPLS Commands

```
config
    — service
        — vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls] [create]
            — service-name service-name
            — no service-name
            — [no] allow-ip-int-bind
                — [no] forward-ipv4-multicast-to-ip-int
            — [no] igmp-snooping
                — [no] mrouter-port
            — [no] mld-snooping
                — [no] mrouter-port
```

3.7.1.9 Multi-Chassis Redundancy Commands

```
— config
    — redundancy
        — multi-chassis
            — [no] peer ip-address
                — [no] sync
                    — [no] igmp-snooping
                    — [no] mld-snooping
                    — pim-snooping [saps] [spoke-sdps]
                    — no pim-snooping
                    — port [port-id | lag-id] [sync-tag sync-tag] [create]
                    — no port [port-id | lag-id]
                        — range encap-range [sync-tag sync-tag]
                        — no range encap-range
```

```
— sdp sdp-id [sync-tag sync-tag] [create]
— no sdp sdp-id
    — range vc-id-range [sync-tag sync-tag]
    — no range vc-id-range
— [no] shutdown
```

3.7.2 Command Descriptions

3.7.2.1 Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>redundancy>multi-chassis>peer>sync
config>service>vpls
config>service>vpls>bgp-ad
config>service>vpls>eth-cfm>mep
config>service>vpls>gsmp
config>service>vpls>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vpls>igmp-snooping
config>service>vpls>igmp-snooping>mvr
config>service>vpls>interface
config>service>vpls>mac-move
config>service>vpls>mesh-sdp
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>mld-snooping
config>service>vpls>mld-snooping>mvr
config>service>vpls>mrp
config>service>vpls>sap
config>service>vpls>sap>arp-host
config>service>vpls>sap>dhcp>proxy
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>sap>igmp-snooping>mcac>mc-constraints
config>service>vpls>sap>mld-snooping>mcac>mc-constraints
config>service>vpls>sap>spb
config>service>vpls>sap>stp
config>service>vpls>sap>sub-sla-mgmt
config>service>vpls>spb>level
config>service>vpls>split-horizon-group
config>service>vpls>spoke-sdp
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>spoke-sdp>spb

config>service>vpls>spoke-sdp>stp
config>service>vpls>stp
config>redundancy>multi-chassis>peer>sync

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases

Service Admin State — Bindings to an SDP within the service will be put into the out-of-service state when the service is shut down. While the service is shut down, all customer packets are dropped and counted as discards for billing and debugging purposes.

Service Operational State — A service is regarded as operational providing that two SAPs or one SDP is operational.

SDP (global) — When an SDP is shut down at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

SDP Keepalives — Enables SDP connectivity monitoring keepalive messages for the SDP ID. Default state is disabled (shutdown), in which case the operational state of the SDP ID is not affected by the keepalive message state.

VPLS SAPs and SDPs — SAPs are created in a VPLS and SDPs are bound to a VPLS in the administratively up default state. The created SAP will attempt to enter the operationally up state. An SDP will attempt to go into the in-service state when bound to the VPLS.

Routed VPLS with forward-ipv4-multicast-to-ip-int and IGMP Snooping — To enable IGMP snooping (configured using **igmp-snooping no shutdown**) in a routed VPLS supporting the forwarding of multicast traffic from the VPLS to the IP interface (configured using **forward-ipv4-multicast-to-ip-int**), it is necessary to enable IGMP on the routed VPLS IP interface.

Routed VPLS with forward-ipv6-multicast-to-ip-int and MLD Snooping — To enable MLD snooping (configured using **mld-snooping no shutdown**) in a routed VPLS supporting the forwarding of multicast traffic from the VPLS to the IP interface (configured using forward-ipv6-multicast-to-ip-int) it is necessary to enable MLD on the routed VPLS IP interface.

description

Syntax **description** *description-string*
no description

Context config>service>vpls
config>service>vpls>sap>dhcp6
config>service>vpls>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vpls>igmp-snooping>mvr
config>service>vpls>split-horizon-group
config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>vpls>sap>dhcp
config>service>vpls>mld-snooping>mvr

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

description

Syntax **description** *long-description-string*
no description

Context config>service>vpls>interface

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *long-description-string* — The description character string. Allowed values are any string up to 160 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

3.7.2.2 VPLS Service Commands

fdb-table-size

Syntax **fdb-table-size** *table-size*
no fdb-table-size

Context config>service>system

Description This command configures the maximum system FDB table size, which is dependent on the chassis type. CPMs with at least 16 GB of memory are required when exceeding 500k MAC addresses in a system. The table size cannot be reduced below its default value, which is also chassis-dependent.

The maximum system FDB table size also limits the maximum FDB table size of any card within the system.

The **no** version of this command sets the table size to its default.

The command default depends on the chassis type and available memory.

Parameters *table-size* — Specifies the maximum system FDB table size.

Values 255999 to 2047999

vpls

Syntax **vpls** *service-id* **customer** *customer-id* **vpn** *vpn-id* [**m-vpls**] [**bvpls** | **i-vpls**] [**etree**] [**name** *name*] [**create**]
no vpls *service-id*

Context config>service

Description This command creates or edits a Virtual Private LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

To create a management VPLS on the 7450 ESS, the m-vpls keyword must be specified. See section **Hierarchical VPLS Redundancy** for an introduction to the concept of management VPLS.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *service-id* — Specifies unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 2147483648
svc-name: 64 characters maximum

customer *customer-id* — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number

Values 1 to 2147483647

Default null (0)

m-vpls — Specifies a management VPLS

e-tree — Specifies a VPLS service as an E-Tree VPLS. E-Tree VPLS services have root and leaf-ac SAPs/SDP bindings and root-leaf-tag SAPs/SDP bindings for E-Tree interconnection. The access (AC) root SAP behaves as a normal VPLS SAP. The AC leaf SAP is restricted to communication only with root-connected services. AC leaf and root SAPs are externally normal SAPs. The AC root SDP bind behaves as a normal VPLS SDP bind. The AC leaf SDP bind is restricted to communication only with root-connected services. AC leaf and root SDP bindings are externally normal SDPs bindings.

In the E-Tree VPLS, the root-ac SAP/SDP bindings can communicate with other root-ac and leaf-ac SAP/SDP bind services locally and remotely. Root originated traffic is marked internally with a root indication and root tagged externally on tag SAP/SDP binds. The leaf-ac SAP/SDP bindings can communicate with other root SAP/SDP bindings locally and remotely. Leaf originated traffic is marked internally with a leaf indication and tagged externally on leaf tag SAP/SDP bindings.

There may be any number of AC SAPs of root or leaf up to typical SAP limits. Network Side tag SAPs for root-leaf use additional resources. These tag SAPs used two tags one for Root and one for Leaf. Network side tag SDPs use a hard coded tag of 1 for root and 2 for leaf. AC SDP bindings are designated as root or leaf SDP bindings but carry no tags marking traffic on the egress frames.

An E-Tree SAP types are specified at creation time. To change the type of a E-Tree SAP the SAP must be removed and re-created.

b-vpls | i-vpls — Creates a backbone-vpls or ISID-vpls

name *name* — A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration data but unused by SR OS. The name is tied to the **service-name** in the service context (setting either **service-name** or **name** will cause the other to change as well).

backbone-vpls

Syntax **backbone-vpls** *vpls-id*[:*isid*]
no backbone-vpls

Context config>service>vpls>pbb

Description This command associated the I-VPLS with the B-VPLS service. The ISID value is used to mux/demux packets for the VPLS flowing through the B-VPLS.

Parameters *vpls-id* — This value represents the VPLS ID value associated with the B-VPLS

isid — Defines ISID associated with the I-VPLS

Default The default is the service-id

Values 0 to 16777215

stp

Syntax [**no**] stp

Context config>service>vpls>pbb>backbone-vpls

Description This command enables STP on the backbone VPLS service.

The **no** form of the command disables STP on the backbone VPLS service.

block-on-mesh-failure

Syntax [**no**] **block-on-mesh-failure**

Context config>service>vpls>spoke-sdp
config>service>vpls>endpoint

Description This command enables blocking (brings the entity to an operationally down state) after all configured SDPs or endpoints are in operationally down state. This event is signaled to corresponding T-LDP peer by withdrawing service label (status-bit-signaling non-capable peer) or by setting “PW not forwarding” status bit in T-LDP message (status-bit-signaling capable peer).

Default disabled

bpdu-translation

Syntax **bpdu-translation** {**auto** | **auto-rw** | **pvst** | **pvst-rw** | **stp**}
no bpdu-translation

Context config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command enables the translation of BPDUs to a specified format, meaning that all BPDUs transmitted on a specified SAP or spoke-SDP will have a specified format.

The **no** form of this command reverts to the default.

Default no bpdu-translation

Parameters **auto** — Specifies that appropriate format will be detected automatically, based on type of BPDUs received on such port.

auto-rw — Specifies that appropriate format will be detected automatically and the VLAN-ID will be rewritten as follows:

- BPDU sent on egress of dot1q SAP will contain the VLAN-ID of the SAP in BDPU-PVID TLV

- BPDU sent on egress of default QinQ SAP will contain the outer VLAN-ID of the SAP in BDPU-PVID TLV
- BPDU sent on egress of QinQ SAP will contain the inner VLAN-ID of the SAP in BDPU-PVID TLV

pvst — Specifies the BPDU-format as PVST. Note: the correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).

pvst-rw — Specifies the BPDU-format as PVST. The VLAN-ID will be rewritten as follows:

- BPDU sent on egress of dot1q SAP will contain the VLAN-ID of the SAP in BDPU-PVID TLV
- BPDU sent on egress of default QinQ SAP will contain the outer VLAN-ID of the SAP in BDPU-PVID TLV
- BPDU sent on egress of QinQ SAP will contain the inner VLAN-ID of the SAP in BDPU-PVID TLV

stp — Specifies the BPDU-format as STP.

calling-station-id

Syntax **calling-station-id** {**mac** | **remote-id** | **sap-id** | **sap-string**}
no calling-station-id

Context config>service>vpls>sap

Description This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages.

The **no** form of the command reverts to the default.

Default no calling-station-id

Parameters **mac** — Specifies that the mac-address will be sent.

remote-id — Specifies that the remote-id will be sent.

sap-id — Specifies that the sap-id will be sent.

sap-string — Specifies that the value is the inserted value set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute will not be sent.

lag-link-map-profile

Syntax **lag-link-map-profile** *link-map-profile-id*
no lag-link-map-profile

Context config>service>vpls>sap

Description This command assigns a pre-configured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/unassigned, the SAP/network interface egress traffic will be re-hashed over LAG as required by the new configuration.

The **no** form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.

Default no lag-link-map-profile

Parameters *link-map-profile-id* — An integer from 1 to 64 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

l2pt-termination

Syntax **l2pt-termination** [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]
no l2pt-termination

Context config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command enables Layer 2 Protocol Tunneling (L2PT) termination on a specified SAP or spoke-SDP.

This feature can be enabled only if STP is disabled in the context of the specified VPLS service.

The **no** form of the command reverts to the default.

Default no l2pt-termination

Parameters **cdp** — Specifies the Cisco discovery protocol

dtp — Specifies the dynamic trunking protocol

pagp — Specifies the port aggregation protocol

stp — Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default)

udld — Specifies unidirectional link detection

vtp — Specifies the virtual trunk protocol

def-mesh-vc-id

Syntax [**no**] **def-mesh-vc-id** *vc-id*

Context config>service>vpls

Description This command configures the value used by each end of a tunnel to identify the VC. If this command is not configured, then the service ID value is used as the VC-ID.

This VC-ID is used instead of a label to identify a virtual circuit.The VC-ID is significant between peer nodes on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

The **no** form of this command disables the VC-ID.

Default none

Parameters *vc-id* — Specifies the default mesh vc-id

Values 1 to 4294967295

mcr-default-gtw

Syntax **mcr-default-gtw**

Context config>service>vpls

Description This command enters the context to configure the default gateway information when using Dual Homing in L2-TPSDA. The IP and MAC address of the default gateway used for subscribers on an L2 MC-Ring are configured in this context. After a ring heals or fails, the system will send out a gratuitous ARP on an active ring SAP in order to attract traffic from subscribers on the ring with connectivity to that SAP.

ip

Syntax **ip** *address*
no ip

Context config>service>vpls>mcr-default-gtw

Description This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the IP address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

The **no** form of the command reverts to the default.

Default no ip

Parameters *address* — Specifies the IP address in a.b.c.d. format.

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>vpls>mcr-default-gtw

Description This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the MAC address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

The **no** form of the command reverts to the default.

Default no mac

Parameters *ieee-address* — Specifies the address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format (cannot be all zeros)

dhcp-python-policy

Syntax **dhcp-python-policy** *policy-name*
no dhcp-python-policy

Context config>service>vpls>sap

Description This command specifies the name of the Python policy. The Python policy is created in the **config>python>python-policy** *policy-name* context.

The **no** form of the command reverts to the default.

Default none

Parameters *policy-name* — Specifies a Python policy name up to 32 characters in length

dhcp6-user-db

Syntax **dhcp6-user-db** *local-user-db-name*
no dhcp6-user-db

Context config>service>vpls>sap

Description This command assigns a local user database for DHCP6 clients (capture SAP only).

The **no** form of the command removes the name from the configuration.

Default none

Parameters *local-user-db-name* — Specifies a local user database name up to 32 characters in length

dhcp6

Syntax **dhcp6**

Context config>service>vpls>sap

Description This command enters the context to configure DHCP6 parameters for this SAP.

interface-id

Syntax **interface-id** [**ascii-tuple**]
interface-id ifindex
interface-id sap-id
interface-id string *string*
no interface-id

Context config>service>ies>if>ipv6>dhcp6>option

Description This command configure the interface-id sub-option of the DHCP6 Relay packet.

The **no** form of the command disables the sending of interface ID options in the DHCPv6 relay packet

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by "|"

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits as well as what is included in ascii-tuple already. The format is supported on dot1q-encapsulated ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

mac — This keyword specifies the MAC address of the remote end is encoded in the sub-option

disable-aging

Syntax [**no**] **disable-aging**

Context config>service>vpls
config>service>vpls>spoke-sdp
config>service>vpls>sap
config>service>template>vpls-template

Description This command disables MAC address aging across a VPLS service or on a VPLS service SAP or spoke-SDP.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB). The **disable-aging** command turns off aging for local and remote learned MAC addresses.

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specific SAPs and/or spoke-SDPs by entering the **disable-aging command at the appropriate level.**

When the **disable-aging** command is entered at the VPLS level, the **disable-aging** state of individual SAPs or SDPs will be ignored.

The **no** form of this command enables aging on the VPLS service.

Default no disable-aging

disable-learning

Syntax [**no**] **disable-learning**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>template>vpls-template

Description This command disables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance, SAP instance or spoke-SDP instance.

When **disable-learning** is enabled, new source MAC addresses will not be entered in the VPLS service forwarding database. This is true for both local and remote MAC addresses.

When **disable-learning is** disabled, new source MAC addresses will be learned and entered into the VPLS forwarding database.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses.

Default no disable-learning (Normal MAC learning is enabled)

discard-unknown

Syntax [**no**] **discard-unknown**

Context config>service>vpls
config>service>template>vpls-template

Description By default, packets with unknown destination MAC addresses are flooded. If **discard-unknown** is enabled at the VPLS level, packets with unknown destination MAC address will be dropped instead (even when configured FDB size limits for VPLS or SAP are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

Default no discard-unknown — Packets with unknown destination MAC addresses are flooded.

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>service>vpls>sap

Description This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid existing DCP policy can be assigned to a SAP or a network interface (this rule does not apply to templates, such as an msap-policy template).

Default If no dist-cpu-protection policy is assigned to a SAP, then the default access DCP policy (_default-access-policy) is used.

If no DCP functionality is required on the SAP, then an empty DCP policy can be created and explicitly assigned to the SAP.

Parameters *policy-name* — Specifies the name of the DCP policy up to 32 characters in length

endpoint

Syntax **endpoint** *endpoint-name* [**create**]
no endpoint

Context config>service>vpls

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name up to 32 characters in length

create — This keyword is mandatory while creating a service endpoint

description

Syntax **description** *description-string*
no description

Context config>service>vpls>endpoint

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

auto-learn-mac-protect

Syntax [**no**] **auto-learn-mac-protect**

Context config>service>vpls>endpoint
config>service>vpls>mesh-sdp
config>service>vpls>sap
config>service>vpls>split-horizon-group
config>service>vpls>spoke-sdp

Description This command specifies whether to enable automatic population of the MAC protect list with source MAC addresses learned on the associated with this SHG. For more information, see Auto-Learn MAC Protect.

The **no** form of the command disables the automatic population of the MAC protect list.

Default auto-learn-mac-protect

ignore-standby-signaling

Syntax [**no**] **ignore-standby-signaling**

Context config>service>vpls>endpoint
config>service>vpls>spoke-sdp

Description When this command is enabled, the node ignores the standby-bit received from the TLDP peers for the specific spoke-SDP and performs internal tasks without taking it into account.

This command is present at the endpoint level and the spoke-SDP level. If the spoke-SDP is part of the explicit-endpoint, this setting cannot be changed at the spoke-SDP level. The existing spoke-SDP will become part of the explicit-endpoint only if the setting is not conflicting. The newly created spoke-SDP, which is a part of the specified explicit-endpoint, will inherit this setting from the endpoint configuration.

Default disabled

revert-time

Syntax **revert-time** *revert-time* | **infinite**
no revert-time

Context config>service>vpls>endpoint

Description This command configures the time to wait before reverting to primary spoke-SDP.

In a regular endpoint the revert-time setting affects just the pseudowire defined as primary (precedence 0). For a failure of the primary pseudowire followed by restoration the revert-timer is started. After it expires the primary pseudowire takes the active role in the endpoint. This behavior does not apply for the case when both pseudowires are defined as secondary. For example, if the active secondary pseudowire fails and is restored it will stay in standby until a configuration change or a force command occurs.

Parameters *revert-time* — Specifies the time to wait, in seconds, before reverting back to the primary spoke-SDP defined on this service endpoint, after having failed over to a backup spoke-SDP

Values 0 to 600

infinite — Specifying this keyword makes endpoint non-revertive

static-mac

Syntax **static-mac** *ieee-address* [**create**]
no static-mac

Context config>service>vpls>endpoint

Description This command assigns a static MAC address to the endpoint. In the FDB, the static MAC is then associated with the active spoke-SDP.

Default none

Parameters *ieee-address* — Specifies the static MAC address to the endpoint

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) Cannot be all zeros

create — This keyword is mandatory while creating a static MAC

suppress-standby-signaling

Syntax [**no**] **suppress-standby-signaling**

Context config>service>vpls>endpoint

Description When this command is enabled, the pseudowire standby bit (value 0x00000020) will not be sent to T-LDP peer when the specified spoke is selected as a standby. This allows faster switchover as the traffic will be sent over this SDP and discarded at the blocking side of the connection. This is particularly applicable to multicast traffic.

Default enabled

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description This command enabled propagation of mac-flush messages received from the specified T-LDP on all spoke and mesh-SDPs within the context of the VPLS service. The propagation will follow split-horizon principles and any data-path blocking in order to avoid looping of these messages.

Default disabled

fdb-table-high-wmark

Syntax [**no**] **fdb-table-high-wmark** *high-water-mark*

Context config>service>vpls
config>service>template>vpls-template

Description This command specifies the value to send logs and traps when the threshold is reached.

Parameters *high-water-mark* — Specifies the value to send logs and traps when the threshold is reached

Values 0 to 100

Default 95%

fdb-table-low-wmark

Syntax [**no**] **fdb-table-low-wmark** *low-water-mark*

Context config>service>vpls
config>service>template>vpls-template

Description This command specifies the value to send logs and traps when the threshold is reached.

Parameters *low-water-mark* — Specifies the value to send logs and traps when the threshold is reached

Values 0 to 100

Default 90%

fdb-table-size

Syntax **fdb-table-size** *table-size*
no fdb-table-size [*table-size*]

Context config>service>vpls
config>service>template>vpls-template

Description This command specifies the maximum number of MAC entries in the forwarding database (FDB) for the VPLS instance on this node.

The **fdb-table-size** specifies the maximum number of forwarding database entries for both learned and static MAC addresses for the VPLS instance.

The **no** form of this command returns the maximum FDB table size to default.

Default fdb-table-size 250

Parameters *table-size* — Specifies the number of entries permitted in the forwarding database for this VPLS instance

Values 7450 ESS, 7950 XRS, or 7750 SR-7, SR-12, SR-12e: 1 to 511999
7750 SR-e, SR-a: 1 to 250000

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>service>vpls

Description This command creates an IP interface.

address

Syntax **address** *ip-address*[/*mask*] [*netmask*]
no address

Context config>service>vpls>interface

Description This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

The local subnet that the **address** command defines must be part of the services address space within the routing context using the **config router service-prefix** command. The default is to disallow the complete address space to services. Once a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the **config router interface** CLI context for network core connectivity with the **exclude** option in the **config router service-prefix** command.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin State	Oper State
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.

Parameters *ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP netmask

The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 to 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses.

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vpls>interface

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

For the 7450 ESS or 7750 SR, when the **arp-populate** and **lease-populate** commands are enabled on an interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured **arp-timeout** value has no effect.

The default value for **arp-timeout** is 14400 seconds (4 hours).

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

hold-time

Syntax **hold-time**

Context config>service>vpls>interface

Description This command creates the CLI context to configure interface level hold-up and hold-down timers for the associated IP interface.

The **up** timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the deactivation of the associated interface for the specified amount of time.

The **down** timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the activation of the associated interface for the specified amount of time

up

Syntax **up ip** *seconds*
no up ip

up ipv6 *seconds*
no up ipv6

Context config>service>vpls>interface>hold-time

Description This command will cause a delay in the deactivation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface down.

The **no** form of the command removes the command from the active configuration and removes the delay in deactivating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it expires.

Parameters *seconds* — The time delay, in seconds, to make the interface operational.

Values 1 to 1200

down

Syntax **down ip** *seconds* [**init-only**]
no up ip
up ipv6 *seconds* [**init-only**]
no up ipv6

Context config>service>vpls>interface>hold-time

Description This command will cause a delay in the activation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up, unless the **init-only** option is configured. If the **init-only** option is configured, the delay is only applied when the IP interface is first configured or after a system reboot.

The **no** form of the command removes the command from the active configuration and removes the delay in activating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it completes.

Parameters *seconds* — The time delay, in seconds, to make the interface operational

Values 1 to 1200

init-only — Specifies that the **down** delay is only applied when the interface is configured or after a reboot

Values 1 to 1200

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>vpls>interface

Description This command assigns a specific MAC address to a VPLS IP interface.

For Routed Central Office (CO), a group interface has no IP address explicitly configured but inherits an address from the parent subscriber interface when needed. For example, a MAC will respond to an ARP request when an ARP is requested for one of the IPs associated with the subscriber interface through the group interface.

The **no** form of the command returns the MAC address of the IP interface to the default value.

Default The system chassis MAC address.

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

static-arp

Syntax **static-arp** *ieee-mac-addr unnumbered*
no static-arp *unnumbered*

Context config>service>vpls>interface

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of the command removes a static ARP entry.

Default None

Parameters *ip-address* — Specifies the IP address for the static ARP in dotted decimal notation

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered — Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

static-mac

Syntax **static-mac**

Context config>service>vpls

Description A set of conditional static MAC addresses can be created within a VPLS supporting bgp-evpn. Conditional static macs are also supported in B-VPLS with SPBM. Conditional Static MACs are dependent on the SAP/SDP state.

This command allows assignment of a set of conditional static MAC addresses to a SAP/spoke-SDP. In the FDB, the static MAC is then associated with the active SAP or spoke-SDP.

Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

Static MACs configured in a bgp-evpn service are advertised as protected (EVPN will signal the mac as protected).

mac

Syntax **mac** *ieee-address* [**create**] **black-hole**
mac *ieee-address* [**create**] **sap** *sap-id* **monitor** {**fwd-status**}
mac *ieee-address* [**create**] **spoke-sdp** *sdp-id:vc-id* **monitor** {**fwd-status**}
no mac ieee-address

Context config>service>vpls>static-mac

Description This command assigns a conditional static MAC address entry to an SPBM B-VPLS SAP/spoke-SDP allowing external MACs for single and multi-homed operation.

For the 7450 ESS or 7750 SR, this command also assigns a conditional static MAC address entry to an EVPN VPLS SAP/spoke-SDP.

Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

Default none

Parameters *ieee-address* — Specifies the static MAC address to an SPBM/sdp-binding interface

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) Cannot be all zeros.

sap-id — Specifies the SAP identifier

sdp-id — Specifies the SDP identifier

Values 1 to 17407

vc-id — Specifies the virtual circuit identifier

Values 1 to 4294967295

create — This keyword is mandatory while creating a static MAC

fwd-status — Specifies that this static mac is based on the forwarding status of the SAP or spoke-SDP for multi-homed operation.

black-hole — Specifies for TLS FDB entries defined on a local SAP the value 'sap', remote entries defined on an SDP have the value 'sdp'.

unnumbered

Syntax **unnumbered** [*ip-int-name* | *ip-address*]
no unnumbered

Context config>service>vpls>if

Description This command configures the interface as an unnumbered interface.

Parameters *ip-int-name* — Specifies the name of the IP interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes

ip-address — Specifies an IP address which must be a valid address of another interface

isid-policy

Syntax **isid-policy**
no isid-policy

Context config>service>vpls

Description This command configures ISID policies for individual ISIDs or ISID ranges in a B-VPLS using SPBM. The ISIDs may belong to I-VPLS services or may be static-isids defined on this node. Multiple entry statements are allowed under a **isid-policy**. ISIDs that are declared as static do not require and **isid-policy** unless the ISIDs are not to be advertised.

isid-policy allows finer control of ISID multicast but is not typically required for SPBM operation. Use of ISID policies can cause additional flooding of multicast traffic.

Default no default

entry

Syntax **entry** *range-entry-id* [**create**]
no entry *range-entry-id*

Context config>service>vpls>isid-policy

Description This command creates or edits an ISID policy entry. Multiple entries can be created using unique entry-id numbers within the ISID policy.

entry-id — Specifies an entry-id uniquely identifies a ISID range and the corresponding actions. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

The following rules govern the usage of multiple entry statements:

- overlapping values are allowed:
  - isid from 301 to 310
  - isid from 305 to 315
  - isid 316
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “isid from 301 to 316” statement.
- there is no consistency check with the content of ISID statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation t o execute the associated action for that entry.

no isid - removes all the previous statements under one entry.

no isid value | from value to higher-value - removes a specific ISID value or range. Must match a previously used positive statement: for example, if the command “isid 16 to 100” was used using “no isid 16 to 50”, it will not work but “no isid 16 to 100 will be successful.

Values 1 to 65535

Default No entry

Parameters *range-entry-id* — Specifies the ID of the ISID policy to be created or edited

Values 1 to 8191

create — Required when first creating the configuration context. Once the context is created, one can navigate into the context without the create keyword.

advertise-local

Syntax [**no**] **advertise-local**

Context config>service>vpls>isid-policy>entry

Description The **no advertise-local** option prevents the advertisement of any locally defined I-VPLS ISIDs or static-isids in the range in a B-VPLS. For I-VPLS services or static-isids that are primarily unicast traffic, the use-def-mcast and no advertise-local options allows the forwarding of ISID based multicast frames locally using the default multicast. The **no advertise-local** option also suppresses this range of ISIDs from being advertised in ISIS. When using the **use-def-mcast** and **no advertise-local** policies, the ISIDs configured under this **static-isid** declarations SPBM treats the ISIDs as belonging to the default tree.

Default advertise-local

range

Syntax **range** *isid* [**to** *isid*]

Context config>service>vpls>isid-policy>entry

Description This command specifies an ISID or a Range of ISIDs in a B-VPLS. One range is allowed per entry.

Default no range

Parameters *isid* — Specifies the ISID value in 24 bits. When singular, ISID identifies a particular ISID to be used for matching

Values 0 to 16777215

to *isid* — Identifies upper value in a range of ISIDs to be used as matching criteria

use-def-mcast

Syntax [**no**] **use-def-mcast**

Context config>service>vpls>isid-policy>entry

Description The **use-def-mcast** option prevents local installation of the ISIDs in the range in the MFIB and uses the default multicast tree instead for a B-VPLS. In a node that does not have I-VPLS or static-isids, this command prevents the building of an MFIB entry for this ISID when received in a SPBM TLV and allows the broadcast of ISID based traffic on the default multicast tree. If an **isid-policy** exists, the core nodes can have this policy to prevent connectivity problems when some nodes are advertising an ISID and others are not. In a I-VPLS service if the customer MAC (C-MAC) is unknown, a frame will have the Multicast DA for an ISID (PBB-OUI + ISID) flooded on the default multicast tree and not pruned.

Default no use-def-mcast

load-balancing

Syntax **load-balancing**

Context config>service>vpls
config>service>template>vpls-template

Description This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Default not applicable

per-service-hashing

Syntax [**no**] **per-service-hashing**

Context config>service>vpls>load-balancing
config>service>template>vpls-template>load-balancing

Description This command enables on a per service basis, consistent per-service hashing for Ethernet services over LAG, over Ethernet tunnel (eth-tunnel) using loadsharing protection-type or over CCAG. Specifically, it enables the new hashing procedures for Epipe, VPLS, regular or PBB services.

The following algorithm describes the hash-key used for hashing when the new option is enabled:

- If the packet is PBB encapsulated (contains an I-TAG ethertype) at the ingress side, use the ISID value from the I-TAG
- If the packet is not PBB encapsulated at the ingress side
  - For regular (non-PBB) VPLS and Epipe services, use the related service ID
  - If the packet is originated from an ingress IVPLS or PBB Epipe SAP
- If there is an ISID configured use the related ISID value
- If there is no ISID yet configured use the related service ID
  - For BVPLS transit traffic use the related flood list id
- Transit traffic is the traffic going between BVPLS endpoints
- An example of non-PBB transit traffic in BVPLS is the OAM traffic
- The above rules apply regardless of traffic type
  - Unicast, BUM flooded without MMRP or with MMRP, IGMP snooped

The **no** form of this command implies the use of existing hashing options.

Default no per-service-hashing

spi-load-balancing

Syntax [**no**] **spi-load-balancing**

Context config>service>vpls>load-balancing
config>service>template>vpls-template>load-balancing

Description This command enables use of the SPI in hashing for ESP/AH encrypted IPv4/v6 traffic. This is a per interface setting.

The **no** form disables the SPI function.

Default disabled

teid-load-balancing

Syntax [**no**] **teid-load-balancing**

Context config>service>vpls>load-balancing
config>service>template>vpls-template>load-balancing

Description This command enables inclusion of TEID in hashing for GTP-U/C encapsulates traffic for GTPv1/GTPv2. The **no** form of this command ignores TEID in hashing.

Default disabled

local-age

Syntax **local-age** *aging-timer*
no local-age

Context config>service>vpls
config>service>template>vpls-template

Description Specifies the aging time for locally learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance. In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **local-age** timer specifies the aging time for local learned MAC addresses.

The **no** form of this command returns the local aging timer to the default value.

Default local age 300 — Local MACs aged after 300 seconds.

Parameters *aging-timer* — Specifies the aging time for local MACs expressed in seconds

Values 60 to 86400

mac-move

Syntax [**no**] **mac-move**

Context config>service>vpls
config>service>template>vpls-template

Description This command enters the context to configure MAC move attributes. A sustained high re-learn rate can be a sign of a loop somewhere in the VPLS topology. Typically, STP detects loops in the topology, but for those networks that do not run STP, the mac-move feature is an alternative way to protect your network against loops.

When enabled in a VPLS, **mac-move** monitors the re-learn rate of each MAC. If the rate exceeds the configured maximum allowed limit, it disables the SAP where the source MAC was last seen. The SAP can be disabled permanently (until a **shutdown**/**no shutdown** command is executed) or for a length of time that grows linearly with the number of times the specified SAP was disabled. You have the option of marking a SAP as non-blockable in the **config>service>vpls>sap>limit-mac-move** or **config>service>vpls>spoke-sdp>limit-mac-move** contexts. This means that when the re-learn rate has exceeded the limit, another (blockable) SAP will be disabled instead.

The **mac-move** command enables the feature at the service level for SAPs and spoke-SDPs, as only those objects can be blocked by this feature. Mesh SDPs are never blocked, but their re-learn rates (sap-to-mesh/spoke-to-mesh or vice versa) are still measured.

The operation of this feature is the same on the SAP and spoke-SDP. For example, if a MAC address moves from SAP to SAP, from SAP to spoke-SDP, or between spoke-SDPs, one will be blocked to prevent thrashing. If the MAC address moves between a SAP and mesh SDP or spoke-SDP and mesh SDP combinations, the respective SAP or spoke-SDP will be blocked.

mac-move will disable a VPLS port when the number of relearns detected has reached the number of relearns needed to reach the move-frequency in the 5-second interval. For example, when the move-frequency is configured to 1 (relearn per second) mac-move will disable one of the VPLS ports when 5 relearns were detected during the 5-second interval because then the average move-frequency of 1 relearn per second has been reached. This can already occur in the first second if the real relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

mac-protect

Syntax **mac-protect**

Context config>service>vpls

Description This command indicates whether or not this MAC is protected on the MAC protect list. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP, spoke-SDP or mesh-SDP that has restricted learning enabled. The MAC protect list is used in conjunction with **restrict-protected-src**, **restrict-unprotected-dst** and **auto-learn-mac-protect**.

Default disabled

mac

Syntax [**no**] **mac** *ieee-address*

Context config>service>vpls>mac-protect

Description This command adds a protected MAC address entry.

Parameters *ieee-address* — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers

mac-subnet-length

Syntax **mac-subnet-length** *subnet-length*
no mac-subnet-length

Context config>service>vpls

Description This command specifies the number of bits to be considered when performing MAC learning (MAC source) and MAC switching (MAC destination). Specifically, this value identifies how many bits, starting from the beginning of the MAC address are used. For example, if the mask-value of 28 is used, MAC learning will only do a lookup for the first 28 bits of the source MAC address when comparing with existing FDB entries. Then, it will install the first 28 bits in the FDB while zeroing out the last 20 bits of the MAC address. When performing switching in the reverse direction, only the first 28 bits of the destination MAC address will be used to perform a FDB lookup to determine the next hop.

The **no** form of this command switches back to full MAC lookup.

Parameters *subnet-length* — Specifies the number of bits to be considered when performing MAC learning or MAC switching

Values 24 to 48

mac-notification

Syntax **mac-notification**

Context config>service>vpls

Description This command controls the settings for the MAC notification message.

The mac-notification message must be generated under the following events:

1. When enabled in the BVPLS using **no shutdown**, a MAC notification will be sent for every active MC-LAG link. The following 3 cases assume no shutdown in the BVPLS.
2. Whenever a related MC-LAG link becomes active (the related MC-LAG link has at least 1 SAP associated with the BVPLS) if the MC-LAG peering is initialized and the PE peers are synchronized.
3. First SAP on an active MC-LAG is associated (via IVPLS/Epipe) with the BVPLS
4. The link between IVPLS/Epipe and BVPLS is configured and there are I-SAPs configured on an active MC-LAG link.

The MAC notification is not sent for the following events:

1. Change of source-bmac or source-bmac-lsb
2. On changes of use-sap-bmac parameter
3. If MC-LAG peering is not (initialized and in sync).

count

Syntax [**no**] **count** *value*

Context config>service>vpls>mac-notification

Description This command configures how often MAC notification messages are sent.

Parameters *value* — Specifies, in seconds, how often MAC notification messages are sent

Values 1 to 10

Default Inherits the chassis level configuration from config>service>mac-notification

interval

Syntax [**no**] **interval** *deci-seconds*

Context config>service>vpls>mac-notification

Description This command controls the frequency of subsequent MAC notification messages.

Default Inherits the chassis level configuration from config>service>mac-notification

Parameters *deci-seconds* — Specifies the frequency of subsequent MAC notification messages, in deciseconds

Values 1 to 100

renotify

Syntax **renotify** *value*
no renotify

Context config>service>vpls>mac-notification

Description This command controls the periodic interval at which sets of MAC notification messages are sent. At each expiration of the renotify timer, a new burst of notification messages is sent, specifically <count> frames at <interval> deci-seconds.

Default no renotify

Parameters *value* — Specifies the time interval between re-notification, in seconds

Values 240 to 840

move-frequency

Syntax **move-frequency** *frequency*
no move-frequency

Context config>service>vpls>mac-move
config>service>template>vpls-template>mac-move

Description This command indicates the maximum rate at which MACs can be re-learned in the VPLS service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MACs.

The **no** form of the command reverts to the default value.

Default 2 (when mac-move is enabled). For example, 10 relearns in a 5 second period.

Parameters *frequency* — Specifies the rate, in 5-second intervals for the maximum number of relearns

Values 1 to 100

number-retries

Syntax **number-retries** *number-retries*
no number-retries

Context config>service>vpls>mac-move
config>service>template>vpls-template>mac-move

Description This command configures the number of times retries are performed for reenabling the SAP/SDP.

Parameters *number-retries* — Specifies number of retries for reenabling the SAP/SDP. A zero (0) value indicates unlimited number of retries.

Values 0 to 255

primary-ports

Syntax **primary-ports**

Context config>service>vpls>mac-move
config>service>template>vpls-template>mac-move

Description This command enters the context to define primary VPLS ports. VPLS ports that were declared as secondary prior to the execution of this command will be moved from secondary port-level to primary port-level. Changing a port to the tertiary level can only be done by first removing it from the secondary port-level.

cumulative-factor

Syntax **cumulative-factor** *cumulative-factor*
no cumulative-factor

Context configure->service->vpls->mac-move->primary-ports
configure->service->vpls->mac-move->secondary-ports
config>service>template>vpls-template>mac-move>primary-ports
config>service>template>vpls-template>mac-move>secondary-ports

Description This command configures a factor for the primary or secondary ports defining how many MAC relearn periods should be used to measure the MAC relearn rate. This rate must be exceeded during consecutive periods before the corresponding ports (SAP and/or spoke-SDP) are blocked by the MAC-move feature.

Parameters *cumulative-factor* — Specifies a MAC relearn period to be used for MAC relearn rate

Values 3 to 10

sap

Syntax **sap** [**split-horizon-group** *group-name*] [**create**] [**capture-sap**]
no sap *sap-id*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>secondary-ports

Description This command declares a specified SAP as a primary (or secondary) VPLS port.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

spoke-sdp

Syntax [**no**] **spoke-sdp** *spoke-id*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>secondary-ports

Description This command declares a specified spoke-SDP as a primary (or secondary) VPLS port.

Parameters *spoke-id* — Specifies the SDP ID to configure as the primary VPLS port

Values 1 to 17407

vc-id — Specifies the virtual circuit identifier

Values 1 to 4294967295

cumulative-factor

Syntax [**no**] **cumulative-factor** *factor*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>seocndary-ports

Description This command defines a factor defining how many mac-relearn measurement periods can be used to measure mac-relearn rate. The rate must be exceeded during the defined number of consecutive periods before the corresponding port is blocked by the mac-move feature. The cumulative-factor of primary ports must be higher than cumulative-factor of secondary ports.

Default 2 — secondary ports

3 — primary ports

Parameters *factor* — Specifies the factor defining the number of mac-relearn measurement periods can be used to measure mac-relearn rate

Values 2 to 10

secondary-ports

Syntax **secondary-ports**

Context config>service>vpls>mac-move
config>service>template>vpls-template>mac-move

Description This command opens configuration context for defining secondary vpls-ports. VPLS ports that were declared as primary prior to the execution of this command will be moved from primary port-level to secondary port-level. Changing a port to the tertiary level can only be done by first removing it from the primary port-level.

retry-timeout

Syntax **retry-timeout** *timeout*
no retry-timeout

Context config>service>vpls>mac-move
config>service>template>vpls-template>mac-move

Description This indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

It is recommended that the retry-timeout value is larger or equal to 5s * cumulative factor of the highest priority port so that the sequential order of port blocking will not be disturbed by re-initializing lower priority ports.

A zero value indicates that the SAP will not be automatically re-enabled after being disabled. If, after the SAP is reenabled it is disabled again, the retry timeout is increased with the provisioned retry timeout in order to avoid thrashing. For example, when retry-timeout is set to 15, it increments (15,30,45,60...).

The **no** form of the command reverts to the default value.

Default retry-timeout 10 (when mac-move is enabled)

Parameters *timeout* — Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled

Values 0 to 120

mcast-ipv6-snooping-scope

Syntax **mcast-ipv6-snooping-scope** {**mac-based** | **sg-based**}

Context config>service>vpls

Description This command specifies the forwarding scope used for IPv6 multicast traffic when PIM snooping for IPv6 is enabled.

By default, the scope is **mac-based**; IPv6 snooped multicast traffic is forwarded is based on the low-order 32 bits of the destination IPv6 address.

When the scope is configured as **sg-based**, the IPv6 snooped multicast traffic is forwarded based on both its full source (if specified in the join) and destination IPv6 address. SG-based forwarding is only supported on FP3-based line cards.

PIM snooping for IPv6 must be disabled to change the forwarding mode within a VPLS service between **mac-based** and **sg-based**.

Default mcast-ipv6-snooping-scope mac-based

Parameters **mac-based** — Enables forwarding for PIM-snooped IPv6 multicast traffic based on the low-order 32 bits of its destination IPv6 address.

sg-based — Enables forwarding for PIM-snooped IPv6 multicast traffic based on its full source (if specified in the join) and destination IPv6 address.

mfib-table-high-wmark

Syntax [**no**] **mfib-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and/or a log entry is added.

Parameters *high-water-mark* — Specifies the multicast FIB high watermark as a percentage

Values 1 to 100

Default 95

mfib-table-low-wmark

Syntax [**no**] **mfib-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB low watermark. When the percentage filling level of the Multicast FIB drops below the configured value, the corresponding trap is cleared and/or a log entry is added.

Parameters *low-water-mark* — Specifies the multicast FIB low watermark as a percentage

Values 1 to 100

Default 90

mfib-table-size

Syntax **mfib-table-size** *size*
no mfib-table-size

Context config>service>vpls

Description This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.

The *mfib-table-size* parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance. When a table-size limit is set on the mfib of a service which is lower than the current number of dynamic entries present in the mfib then the number of entries remains above the limit.

The **no** form of this command removes the configured maximum MFIB table size.

Default none

Parameters *size* — Specifies the maximum number of (s,g) entries allowed in the Multicast FIB

Values 1 to 16383

mld-snooping

Syntax **mld-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>bind

Description This command configures MLD snooping parameters.

remote-age

Syntax **remote-age** *seconds*
no remote-age

Context config>service>vpls
config>service>template>vpls-template

Description Specifies the aging time for remotely learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance. In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **remote-age** timer specifies the aging time for remote learned MAC addresses. To reduce the amount of signaling required between switches configure this timer larger than the **local-age** timer.

The **no** form of this command returns the remote aging timer to the default value.

Default remote age 900 — Remote MACs aged after 900 seconds

Parameters *seconds* — Specifies the aging time for remote MACs expressed in seconds

Values 60 to 86400

selective-learned-fdb

Syntax [**no**] **selective-learned-fdb**

Context config>service>vpls

Description This command determines which line cards FDB entries are allocated on for MAC addresses in the VPLS service in which the command is configured.

By default, FDB entries for MAC addresses in VPLS services are allocated on all line cards in the system. Enabling **selective-learned-fdb** causes FDB entries to be allocated only on the line cards on which the service has a configured object, which includes all line cards:

- on which a SAP is configured
- which have ports configured in a LAG SAP
- which have ports configured in an Ethernet tunnel SAP
- which have ports configured on a network interface (which also may be on a LAG) when the service has a mesh or spoke-SDP, VXLAN or EVPN-MPLS configured

Only MAC addresses with a type “L” or “Evpn” in the **show** output displaying the FDB can be allocated selectively, unless a MAC address configured as a conditional static MAC address is learned dynamically on an object other than its monitored object; this can be displayed with type “L” or “Evpn” but is allocated as global because of the conditional static MAC configuration.

The **no** form of this command returns the FDB MAC address entry allocation mode to its default where FDB entries for MAC addresses are allocated on all line cards in the system.

Default no selective-learned-fdb

send-bvpls-flush

Syntax **send-bvpls-flush** {[**all-but-mine**] [**all-from-me**]}

no send-bvpls-flush

Context config>service>vpls

Description This command enables generation of LDP MAC withdrawal “flush-all-from-me” in the B-VPLS domain when the following triggers occur in the related IVPLS:

- MC-LAG failure
- Failure of a local SAP
- Failure of a local pseudowire/SDP binding

A failure means transition of link SAP/pseudowire to either down or standby status.

This command does not require send-flush-on-failure in B-VPLS to be enabled on an IVPLS trigger to send an MAC flush into the BVPLS.

Default no send-bvpls-flush

Parameters **all-but-mine** — Specifies to send an LDP flush all-but-mine and also sent into the B-VPLS. Both parameters can be set together.

all-from-me — Specifies to send an LDP flush-all-from and when STP initiates a flush, it is sent into the B-VPLS using LDP MAC flush all-from-me. Both parameters can be set together.

send-flush-on-failure

Syntax [**no**] **send-flush-on-failure**

Context config>service>vpls

Description This command enables sending out flush-all-from-me messages to all LDP peers included in affected VPLS, in the event of physical port failures or “operationally down” events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, in the case of link failure, is resolved by RSTP running between a PE router and CE devices. If the endpoint is configured within the VPLS and send-flush-on-failure is enabled, flush-all-from-me messages will be sent out only when all spoke-SDPs associated with the endpoint go down.

This feature cannot be enabled on management VPLS.

Default no send-flush-on-failure

service-mtu

Syntax **service-mtu** *octets*
no service-mtu

Context config>service>vpls
config>service>template>vpls-template

Description This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU. The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

For i-VPLS and Epipes bound to a b-VPLS, the service-mtu must be at least 18 bytes smaller than the b-VPLS service MTU to accommodate the PBB header.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Default VPLS: 1514

Parameters *octets* — The following table displays MTU values for specific VC types

VC-Type	Example Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

The size of the MTU in octets, expressed as a decimal integer

Values 1 to 9194

service-name

Syntax **service-name** *service-name*
no service-name

Context config>service>vpls

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a specified service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a specified service when it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0 to 9).

allow-ip-int-bind

Syntax [**no**] **allow-ip-int-bind**

Context config>service>vpls

Description The allow-ip-int-bind command that sets a flag on the VPLS or I-VPLS service that enables the ability to attach an IES or VPRN IP interface to the VPLS service in order to make the VPLS service routable. When the allow-ip-int-bind command is not enabled, the VPLS service cannot be attached to an IP interface.

VPLS Configuration Constraints for Enabling allow-ip-int-bind

When attempting to set the allow-ip-int-bind VPLS flag, the system first checks to see if the correct configuration constraints exist for the VPLS service and the network ports. The following VPLS features must be disabled or not configured for the allow-ip-int-bind flag to set:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- The VPLS service type cannot be B-VPLS or M-VPLS
- MVR from Routed VPLS and to another SAP is not supported
- Enhanced and Basic Subscriber Management (ESM and BSM) features
- Network domain on SDP bindings

Once the VPLS allow-ip-int-bind flag is set on a VPLS service, the above features cannot be enabled on the VPLS service.

Network Port Hardware Constraints

The system also checks to ensure that all ports configured in network mode are associated with FlexPath2 forwarding planes. If a port is currently in network mode and the port is associated with a FlexPath1 forwarding plane, the allow-ip-int-bind command will fail. Once the allow-ip-int-bind flag is set on any VPLS service, attempting to enable network mode on a port associated with a FlexPath1 forwarding plane will fail.

VPLS SAP Hardware Constraints

Besides VPLS configuration and network port hardware association, the system also checks to that all SAPs within the VPLS are created on Ethernet ports and the ports are associated with FlexPath2 forwarding planes. Certain Ethernet ports and virtual Ethernet ports are not supported which include HSMDA ports and CCAG virtual ports (VSM based). If a SAP in the VPLS exists on an unsupported port type or is associated with a FlexPath1 forwarding plane, the allow-ip-int-bind command will fail. Once the allow-ip-int-bind flag is set on the VPLS service, attempting to create a VPLS SAP on the wrong port type or associated with a FlexPath1 forwarding plane will fail.

VPLS Service Name Bound to IP Interface without allow-ip-int-bind flag Set

In the event that a service name is applied to a VPLS service and that service name is also bound to an IP interface but the allow-ip-int-bind flag has not been set on the VPLS service context, the system attempt to resolve the service name between the VPLS service and the IP interface will fail. After the allow-ip-int-bind flag is successfully set on the VPLS service, either the service name on the VPLS service must be removed and reapplied or the IP interface must be re-initialized using the **shutdown** / **no shutdown** commands. This will cause the system to reattempt the name resolution process between the IP interface and the VPLS service.

The **no** form of the command resets the allow-ip-int-bind flag on the VPLS service. If the VPLS service currently has an IP interface from an IES or VPRN service attached, the no allow-ip-int-bind command will fail. Once the allow-ip-int-bind flag is reset on the VPLS service, the configuration and hardware restrictions associated with setting the flag are removed. The port network mode hardware restrictions are also removed.

forward-ipv4-multicast-to-ip-int

Syntax [**no**] **forward-ipv4-multicast-to-ip-int**

Context config>service>vpls>allow-ip-int-bind

Description This command enables support for forwarding IPv4 multicast traffic from sources connected to the VPLS service of a routed VPLS to the IP interface of the routed VPLS service. It can only be enabled after the routed VPLS service has been bound to an IP interface.

Default no forward-ipv4-multicast-to-ip-int

forward-ipv6-multicast-to-ip-int

Syntax [**no**] **forward-ipv6-multicast-to-ip-int**

Context config>service>vpls>allow-ip-int-bind

Description This command enables support for forwarding IPv6 multicast traffic from sources connected to the VPLS service of a routed VPLS to the IP interface of the routed VPLS service. It can only be enabled after the routed VPLS service has been bound to an IP interface.

Default no forward-ipv6-multicast-to-ip-int

vxlan-ipv4-tep-ecmp

Syntax [**no**] **vxlan-ipv4-tep-ecmp**

Context config>service>vpls>allow-ip-int-bind

Description This command enables and disables ECMP on VXLAN IPv4 destinations for R-VPLS services. When this command is enabled, packets entering a VPRN connected to an R-VPLS that is terminating on a VXLAN IPv4 destination are looked up in the routing table. If the next hop is a VXLAN IPv4 TEP, the packets are distributed based on per-flow load-balancing.

This command can only be used in FP3-only routers. R-VPLS per-flow load-balancing for VXLAN IPv6 destinations works by default without this command.

The **no** version of this command reverts the process to the default behavior of per-remote VTEP load-balancing.

Default no vxlan-ipv4-tep-ecmp

site

Syntax **site** *name* [**create**]
no site *name*

Context config>service>vpls

Description This command configures a VPLS site.

The no form of the command removes the name from the configuration.

Parameters *name* — Specifies a site name up to 32 characters in length

create — This keyword is mandatory while creating a VPLS service

boot-timer

Syntax **boot-timer** *seconds*
no boot-timer

Context config>service>vpls>site

Description This command configures for how long the service manger waits after a node reboot before running the DF election algorithm. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged.

The **no** form of the command reverts the default.

Default 10

Parameters *seconds* — Specifies the site boot-timer in seconds

Values 0 to 100

failed-threshold

Syntax **failed-threshold** [1 to 1000]
failed-threshold all

Context config>service>vpls>site

Description This command defines the number of objects should be down for the site to be declared down. Both administrative and operational status must be evaluated and if at least one is down, the related object is declared down.

Default failed-threshold all

Parameters *1 to 1000* — Specifies the threshold for the site to be declared down

mesh-sdp-binding

Syntax [**no**] **mesh-sdp-binding**

Context config>service>vpls>site

Description This command enables applications to all mesh SDPs.

The **no** form of reverts the default.

Default no mesh-sdp-binding

monitor-oper-group

Syntax **monitor-oper-group** *group-name*
no monitor-oper-group

Context config>service>vpls>site
config>service>vpls>spoke-sdp
config>service>vpls>sap
config>service>vpls>bgp>pw-template-binding

Description This command specifies the operational group to be monitored by the object under which it is configured. The **oper-group** *name* must be already configured under the **config>service** context before its name is referenced in this command.

The **no** form of the command removes the association.

Parameters *group-name* — Specifies the name of the operational group. 32 characters maximum

sap

Syntax **sap** *sap-id*
no sap

Context config>service>vpls>site

Description This command configures a SAP for the site.

The **no** form of the command removes the SAP ID from the configuration.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

site-activation-timer

Syntax **site-activation-timer** *seconds*
no site-activation-timer

Context config>service>vpls>site

Description This command configures the time-period the system keeps the local sites in standby status, waiting for BGP updates from remote PEs before running the DF (designated-forwarder) election algorithm to decide whether the site should be unblocked. This timer if terminated if an update is received for which the remote PE has transitioned from DF to non-DF.

The no form of the command removes the value from the configuration.

Default 2

Parameters *seconds* — Specifies the site activation timer in seconds

Values 0 to 100

site-min-down-timer

Syntax **site-min-down-timer** *min-down-time*
no site-min-down-timer

Context config>service>vpls>site

Description This command configures the BGP multi-homing site minimum down time. When set to a non-zero value, if the site goes operationally down it will remain operationally down for at least the length of time configured for the **site-min-down-timer**, regardless of whether other state changes would have caused it to go operationally up. This timer is restarted every time that the site transitions from up to down. Setting this parameter to zero allows the minimum down timer to be disabled for this service.

The above operation is optimized in the following circumstances:

- If the site goes down on the designated forwarder but there are no BGP multi-homing peers with the same site in an UP state, then the **site-min-down-timer** is not started and is not used.
- If the site goes down on the designated forwarder but there are no active BGP multi-homing peers, then the **site-min-down-timer** is not started and is not used.
- If the s**ite-min-down-timer** is active and a BGP multi-homing update is received from the designated forwarder indicating its site has gone down, the **site-min-down-timer** is immediately terminated and this PE becomes the designated forwarder if the BGP multi-homing algorithm determines it should be the designated forwarder.

The **no** form of this command reverts to the default value.

Default Taken from the value of **site-min-down-timer** configured for Multi-Chassis BGP multi-homing under the **config>redundancy>bgp-multi-homing** context.

Parameters *min-down-time* — Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down

Values 0 to 100 seconds

site-id

Syntax **site-id** *value*
no site-id

Context config>service>vpls>site

Description This command configures the identifier for the site in this service.

Parameters *value* — Specifies the site identifier

Values 1 to 65535

split-horizon-group

Syntax **split-horizon-group** *group-name*
no split-horizon-group

Context config>service>vpls>site

Description This command configures the value of split-horizon group associated with this site.

The **no** form of the command reverts the default.

Default no split-horizon-group

Parameters *group-name* — Specifies a split-horizon group name

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id*
no spoke-sdp

Context config>service>vpls>site

Description This command binds a service to an existing Service Distribution Point (SDP).

The **no** form of the command removes the parameter from the configuration.

Parameters *sdp-id* — Specifies the SDP identifier

Values 1 to 17407

vc-id — Specifies the virtual circuit identifier.

Values 1 to 429496729

spb

Syntax **spb** [*isis-instance*] [**fid** *fid*] [**create**]
no spb

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command configures Shortest Path Bridging.

Parameters *isis-instance* — Specifies the ISIS instance

Values 1024 to 2047

fid — Specifies the FID

Values 1 to 4095

create — This keyword is mandatory when creating an SPB instance

level

Syntax **level** [1 to 1]

Context config>service>vpls>spb
config>service>vpls>sap>spb
config>service>vpls>spoke-sdp>spb

Description This command enters the context to configure SPB level information.

bridge-priority

Syntax **bridge-priority** *bridge-priority*
no bridge-priority

Context config>service>vpls>spb>level

Description This command configures the level 1 four bit bridge priority associated with this Shortest Path Bridging context in this VPLS service.

Default 8

Parameters *bridge-priority* — Specifies the bridge priority

Values 0 to 15

ect-algorithm

Syntax **ect-algorithm fid-range** *fid-range* {**low-path-id** | **high-path-id**}
no ect-algorithm fid-range *fid-range*

Context config>service>vpls>spb>level

Description This command configures the ECT algorithm of forwarding range.

Parameters *fid-range* — Specifies the FID range in the form *start-end*; for example, 50-120

Values *start* — 1 to 4095
end — 1 to 4095

low-path-id — Keyword specifies the low path ID

high-path-id — Keyword specifies the high path ID

forwarding-tree-topology

Syntax **forwarding-tree-topology unicast** {**spf** | **st**}

Context config>service>vpls>spb>level

Description This command specifies level 1 unicast forwarding to follow the shortest path tree or to follow a single tree for this Shortest Path Bridging context in this VPLS service.

Default spf

Parameters **spf** — Follows the shortest path tree.

st — Follows a single tree.

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>service>vpls>spb>level
config>service>vpls>sap>spb>level
config>service>vpls>spoke-sdp>spb>level

Description This command configures the interval in seconds between hello messages issued on this interface at this level. This command is valid only for interfaces on control B-VPLS.

The **no** form of the command to reverts to the default value.

Default hello-interval 3 (for the designated intersystem)

hello-interval 9 (for non-designated intersystems)

Parameters *seconds* — Specifies the hello interval, in seconds, expressed as a decimal integer

Values 1 to 20000

hello-multiplier

Syntax **hello-multiplier** *multiplier*
no hello-multiplier

Context config>service>vpls>spb>level
config>service>vpls>sap>spb>level

config>service>vpls>spoke-sdp>spb>level

Description This command configures the number of missing hello PDUs from a neighbor SPB declares the adjacency down. This command is valid only for interfaces on control B-VPLS.

The no form of the command reverts to the default value.

Default hello-multiplier 3

Parameters *multiplier* — Specifies the multiplier for the hello interval expressed as a decimal integer

Values 2 to 100

lsp-lifetime

Syntax **lsp-lifetime** *seconds*
no lsp-lifetime

Context config>service>vpls>spb

Description This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the **lsp-lifetime** expires unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from this formula: lsp-lifetime/2.

The **no** form of the command reverts to the default value.

Default lsp-lifetime 1200

Parameters *seconds* — Specifies the time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain

Values 350 to 65535

lsp-refresh-interval

Syntax **lsp-refresh-interval** [*seconds*] [**half-lifetime** [**enable** | **disable**]
no lsp-refresh-interval

Context config>service>vpls>spb

Description This command configures the LSP refresh timer interval. When configuring the LSP refresh interval, the value that is specified for **lsp-lifetime** must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.

The no form of the command reverts to the default (600 seconds), unless this value is greater than 90% of the LSP lifetime. For example, if the LSP lifetime is 400, then the **no lsp-refresh-interval** command will be rejected.

Default 600, half-lifetime enable

Parameters *seconds* — Specifies the refresh interval.

Values 150 to 65535

half-lifetime — Sets the refresh interval to always be half the **lsp-lifetime** value. When this parameter is set to **enable**, the configured refresh interval is ignored.

Values enable, disable

timers

Syntax [**no**] **timers**

Context config>service>vpls>spb

Description This command configures the IS-IS timer values.

Default disabled

lsp-wait

Syntax **lsp-wait** *lsp-wait* [**lsp-initial-wait** *initial-wait*] [**lsp-second-wait** *second-wait*]]

Context config>service>vpls>spb>timers

Description This command is used to customize LSP generation throttling. Timers that determine when to generate the first, second and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

Note: The IS-IS timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Parameters *lsp-max-wait* — Specifies the maximum interval in milliseconds between two consecutive occurrences of an LSP being generated.

Values 10 to 120000

Default 50000

initial-wait — Specifies the initial LSP generation delay in milliseconds. Values < 100 ms are internally rounded down to 0, so that there is no added initial LSP generation delay.

Values 10 to 100000

Default 10

second-wait — Specifies the hold time in milliseconds between the first and second LSP generation.

Values 10 to 100000

Default 1000

spf-wait

Syntax [**no**] **spf-wait** *spf-wait* [**spf-initial-wait** *initial-wait*] [**spf-second-wait** *second-wait*]]

Context config>service>vpls>spb>timers

Description This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second and subsequent SPF calculations after a topology change occurs can be controlled with this command.

Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, and so on, until it reaches the *spf-wait* value. The SPF interval will stay at *spf-wait* value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

Note: The IS-IS timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default no spf-wait

Parameters *spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

Values 10 to 120000

Default 10000

initial-wait — Specifies the initial SPF calculation delay in milliseconds after a topology change

Values 10 to 100000

Default 1000

second-wait — Specifies the hold time in milliseconds between the first and second SPF calculation

Values 10 to 100000

Default 1000

overload-on-boot

Syntax **overload-on-boot** [**timeout** *seconds*]
no overload-on-boot

Context config>service>vpls>spb

Description When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

1. The timeout timer expires.
2. A manual override of the current overload state is entered with the **config>router>isis>no overload** command.

The **no overload** command does not affect the **overload-on-boot** function.

If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:

- L1 LSDB Overload : Manual on boot (Indefinitely in overload)
- L2 LSDB Overload : Manual on boot (Indefinitely in overload)

This state can be cleared with the **config>router>isis>no overload** command.

When specifying a timeout value, IS-IS will go into overload for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

- L1 LSDB Overload : Manual on boot (Overload Time Left : 17)
- L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **config>router>isis>no overload** command.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Default no overload-on-boot

Use **show router ospf status** or **show router isis status** commands to display the administrative and operational state as well as all timers.

Parameters *seconds* — Specifies the timeout timer for overload-on-boot, in seconds

Values 60 to 1800

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>service>vpls>spb

Description This command administratively sets the router to operate in the overload state for a specific time period, in seconds, or indefinitely.

During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The **overload** command can be useful in circumstances where the router is overloaded or used prior to executing a **shutdown** command to divert traffic around the router.

The **no** form of the command causes the router to exit the overload state.

Default no overload

Parameters *seconds* — Specifies the time, in seconds, that the router must operate in an overload state

Default infinity (overload state maintained indefinitely)

Values 60 to 1800

metric

Syntax **metric** *ipv4-metric*
no metric

Context config>service>vpls>spb>level
config>service>vpls>sap>spb>level
config>service>vpls>spoke-sdp>spb>level

Description This command configures the IS-IS interface metric for IPv4 unicast.

Parameters *ipv4-metric* — Specifies the IS-IS interface metric for IPv4 unicast

Values 1 to 16777215

lsp-pacing-interval

Syntax **lsp-pacing-interval** *milli-seconds*
no lsp-pacing-interval

Context config>service>vpls>sap>spb
config>service>vpls>spoke-sdp>spb

Description This command configures the interval during which LSPs are sent from the interface.

To avoid overwhelming neighbors that have less CPU processing power with LSPs, the pacing interval can be configured to limit how many LSPs are sent during an interval. LSPs may be sent in bursts during the interval up to the configured limit. If a value of 0 is configured, no LSPs are sent from the interface.

The **no** form of the command reverts to the default value.

Note: The IS-IS timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default lsp-pacing-interval 100

Parameters *milli-seconds* — Specifies the interval, in milliseconds, during which IS-IS LSPs are sent from the interface expressed as a decimal integer

Values 0 to 65535

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>service>vpls>sap>spb
config>service>vpls>spoke-sdp>spb

Description This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.

The **no** form of the command reverts to the default value.

Default retransmit-interval 100

Parameters *seconds* — Specifies the interval in seconds that IS-IS LSPs can be sent on the interface

Values 1 to 65535

split-horizon-group

Syntax [**no**] **split-horizon-group** [*group-name*] [*residential-group*]

Context config>service>vpls

Description This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke-SDP within this split horizon group will not be copied to other SAPs or spoke-SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke-SDPs can be assigned to the group.

The split horizon group is defined within the context of a single VPLS. The same group-name can be re-used in different VPLS instances.

Up to 30 split horizon groups can be defined per VPLS instance. Half are supported in i-VPLS.

The **no** form of the command removes the group name from the configuration.

Default A split horizon group is by default not created as a residential-group.

Parameters *group-name* — Specifies the name of the split horizon group to which the SDP belongs

residential-group — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

a) SAPs which are members of this Residential Split Horizon Group will have:

- Double-pass queuing at ingress as default setting (can be disabled)
- STP disabled (cannot be enabled)
- ARP reply agent enabled per default (can be disabled)
- MAC pinning enabled per default (can be disabled)
- Downstream broadcast packets are discarded thus also blocking the unknown, flooded traffic
- Downstream multicast packets are allowed when IGMP snooping is enabled

b) Spoke SDPs which are members of this Residential Split Horizon Group will have:

- Downstream multicast traffic supported
- Double-pass queuing is not applicable
- STP is disabled (can be enabled)
- ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke-SDPs)
- MAC pinning enabled per default (can be disabled)

process-cpm-traffic-on-sap-down

Syntax **process-cpm-traffic-on-sap-down**
[**no**] **process-cpm-traffic-on-sap-down**

Context config>service>vpls>sap

Description This command is applicable to simple SAPs configured on LAGs that are not part of any "endpoint" configurations or complicated resiliency schemes like MC-LAG with inter-chassis-backup (ICB) configurations. When configured, a simple LAG SAP will not be removed from the forwarding plane and flooded traffic (unknown unicast, broadcast and multicast) will be dropped on egress. This allows applicable control traffic that is extracted at the egress interface to be processed by the CPM. This command will not prevent a VPLS service from entering an Operational Down state if it is the last active connection to enter a non-operational state. By default, without this command, when a SAP on a LAG enters a non-operational state it is removed from the forwarding plane and no forwarding occurs to the egress.

The **no** form of the command means a SAP over a LAG that is not operational will be removed from the forwarding process.

Default no process-cpm-traffic-on-sap-down

pppoe-policy

Syntax **pppoe-policy** *pppoe-policy-name*
no pppoe-policy

Context config>service>vpls>sap

Description For the 7450 ESS or 7750 SR, this command specifies an existing PPPoE policy. These policies are referenced from interfaces configured for PPPoE. Multiple PPPoE policies may be configured.

Default none

Parameters *pppoe-policy-name* — Specifies an existing PPPoE policy name up to 32 characters in length

auto-learn-mac-protect

Syntax [**no**] **auto-learn-mac-protect**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls >mesh-sdp
config>service>vpls>split-horizon-group
config>service>vpls>endpoint
config>service>pw-template
config>service>pw-template>split-horizon-group

Description This command enables the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with **restrict-protected-src**, **restrict-unprotected-dst** and **mac-protect**. When this command is applied or removed, the MAC addresses are cleared from the related object.

When the **auto-learn-mac-protect** is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke-SDPs in the SHG). To enable this function for spoke-SDPs within a SHG, the **auto-learn-mac-protect** must be enabled explicitly under the spoke-SDP. If required, **auto-learn-mac-protect** can also be enabled explicitly under specific SAPs within the SHG.

Default no auto-learn-mac-protect

restrict-protected-src

Syntax **restrict-protected-src** [{**alarm-only** | **discard-frame**}]
no restrict-protected-src

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>split-horizon-group
config>service>vpls>endpoint
config>service>pw-template>
config>service>pw-template>split-horizon-group

Description This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP, spoke-SDP, mesh-SDP, or any SAP that is part of the configured split horizon group (SHG) will be verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the restrict-protected-src command, namely:

- No parameter — The packet will be discarded, an alarm will be generated and the SAP, spoke-SDP or mesh SDP will be set operationally down. The SAP, spoke-SDP or mesh-SDP must be shutdown and enabled (no shutdown) for this state to be cleared.
- alarm-only — The packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP, spoke-SDP or mesh SDP.
- discard-frame — The packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes in a specified VPLS service. This parameter is only applicable to automatically protected MAC addresses.

When the **restrict-protected-src** is enabled on an SHG, the action only applies to the associated SAPs (no action is taken by default for spoke-SDPs in the SHG) and is displayed in the SAP show output as the oper state unless it is overridden by the configuration of **restrict-protected-src** on the SAP itself. To enable this function for spoke-SDPs within a SHG, the **restrict-protected-src** must be enabled explicitly under the spoke-SDP. If required, **restrict-protected-src** can also be enabled explicitly under specific SAPs within the SHG.

When this command is applied or removed, with either the **alarm-only** or **discard-frame** parameters, the MAC addresses are cleared from the related object.

The use of **restrict-protected-src discard-frame** is mutually exclusive with the configuration of manually protected MAC addresses within a specified VPLS.

The **alarm-only** parameter is not supported on the 7750 SR-a, 7750 SR-1e/2e/3e, or 7950 XRS.

Default no restrict-protected-src

Parameters **alarm-only** — Specifies that the packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP, spoke-SDP or mesh SDP

discard-frame — Specifies that the packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes within a specified VPLS service

restrict-unprotected-dst

Syntax **restrict-unprotected-dst**
no restrict-unprotected-dst

Context config>service>pw-template>split-horizon-group
config>service>vpls>split-horizon-group
config>service>vpls>sap

Description This command indicates how the system will forward packets destined for an unprotected MAC address, either manually added using the **mac-protect** command or automatically added using the **auto-learn-mac-protect** command. While enabled all packets entering the configured SAP or SAPs within a split horizon group (but not spoke or mesh-SDPs) will be verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it will be discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state. No alarms are generated.

If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with restrict-unprotected-dst enabled, it will be flooded.

Default no restrict-unprotected-dst

shcv-policy-ipv4

Syntax **shcv-policy-ipv4** *policy-name*
no shcv-policy-ipv4

Context config>service>vpls>sap

Description This command specifies the Subscriber Host Connectivity Verification (SHCV) policy for IPv4 only.

The **no** form of the command removes the policy name from the SAP configuration.

vpls-group

Syntax [**no**] **vpls-group** *id*

Context config>service>vpls

Description This command defines a vpls-group index. Multiple vpls-group commands can be specified to allow the use of different VPLS and SAP templates for different ranges of service ids. A vpls-group can be deleted only in shutdown state. Multiple commands under different vpls-group ids can be issued and can be in progress at the same time.

Default no vpls-group

Parameters *id* — Specifies the ID associated with the VPLS group

Values 1 to 4094

service-range

Syntax **service-range** *startid-endid* [**start-vlan-id** *startvid*]
no service-range *startid-endid*

Context config>service>vpls>vpls-group

Description This command configures the service ID and implicitly the VLAN-ID ranges to be used as input variables for related VPLS and SAP templates to pre-provision “data” VPLS instances and related SAPs using the service ID specified in the command. If the start-vlan-id is not specified then the service-range values are used for vlan-ids. The data SAPs will be instantiated on all the ports used to specify SAP instances under the related control VPLS.

Modifications of the service id and vlan ranges are allowed with the following restrictions.

- service-range increase can be achieved in two ways:
  - Allowed when vpls-group is in shutdown state
  - By creating a new vpls-group
- service-range decrease can be achieved in two ways:

- Allowed when vpls-group is in shutdown state; when **shutdown** command is executed the associated service instances are deleted.
- Allowed when vpls-group is in no shutdown state and has completed successfully instantiating services.
- In both cases, only the services that do not have user configured SAPs will be deleted. Otherwise the above commands are rejected. Existing declarations or registrations do not prevent service deletion.

• start-vlan-id change can be achieved in two ways:

- Allowed when vpls-group is in shutdown state
- At the time of range decrease by increasing the start-vlan-id which can be done when vpls-group is in no shutdown state and has completed successfully instantiating services

The **no** form of this command removes the specified ranges and deletes the pre-provisioned VPLS instances and related SAPs. The command will fail if any of the VPLS instances in the affected ranges have a provisioned SAP.

Default no service-range

Parameters *startid-endid* — Specifies the range of service IDs

Values 1 to 2147483647

startvid — Specifies the starting VLAN ID; it provides a way to set aside a service ID range that is not the same as the VLAN range and allows for multiple MVRP control-VPLSs to control same VLAN range on different ports.

Values 1 to 4094

vpls-template-binding

Syntax **vpls-template-binding** *name/id*
no vpls-template-binding

Context config>service>vpls>vpls-group

Description This command configures the binding to a VPLS template to be used to instantiate pre-provisioned data VPLS using as input variables the service IDs generated by the vid-range command.

The **no** form of this command removes the binding and deletes the related VPLS instances. The command will fail if any of the affected VPLS instances have either a provisioned SAP or an active MVRP declaration/registration or if the related vpls-group id is in no shutdown state. Any changes to the vpls-template-binding require the vpls-group to be in shutdown state.

Default no vpls-template-binding

Parameters *name/id* — Specifies the name or the ID of the VPLS template

Values 1 to 1024

sap-template-binding

Syntax **sap-template-binding** *name/id*
no sap-template-binding

Context config>service>vpls>vpls-group

Description This command configures the binding to a SAP template to be used to instantiate SAPs in the data VPLS using as input variables the VLAN IDs generated by the vid-range command.

The **no** form of this command removes the binding and deletes the related SAP instances. The command will fail if any of the affected VPLS instances have either a provisioned SAP or an active MVRP declaration/registration registration or if the related vpls-group is in no shutdown state. Any changes to the **sap-template-binding** require the **vpls-group** to be in **shutdown** state. New control SAP additions to the management VPLS are allowed as long as data VPLS instantiations/removals for vpls-groups are not in progress. Control SAPs can be removed at any time generating the removal of related data SAPs from the data VPLS. The **shutdown** or **no shutdown** state for the control SAPs does not have any effect on data SAPs instantiated with this command.

Default no sap-template-binding

Parameters *name* — Specifies the name of the VPLS template

Values ASCII character string

id — Specifies the ID of the VPLS template

Values 1 to 8196

mvrp-control

Syntax [**no**] **mvrp-control**

Context config>service>vpls>vpls-group

Description This command enables MVRP control in the VPLS instances instantiated using the templates for the specified vpls-group. That means the flooding FDB will be created empty and will be populated with endpoints whenever MVRP receives a declaration and a registration on a specific endpoint. Also the VLAN ID associated by the control VPLS with the instantiated VPLS will be declared on service activation by MVRP on all virtual MVRP ports in the control VPLS. Service activation takes place when at least one other SAP is provisioned and brought up under the data VPLS. This is usually a customer facing SAP or a SAP leading outside of the MVRP controlled domain.

The **no** form of this command disallows MVRP control over this VPLS. The VPLS will be created with a regular FDB and will become as a result active upon creation time. Command change is allowed only when the related vpls-group is in shutdown state.

Default no mvrp-control

mvrp

Syntax **mvrp**

Context config>service>vpls>mrp
config>service>vpls>sap>mrp

Description This object consolidates the MVRP attributes. MVRP is only supported initially in the management VPLS so the object is not supported under BVPLS, IVPLS or regular VPLS not marked with the m-vpls tag.

hold-time

Syntax **hold-time** *value*
no hold-time

Context config>service>vpls>mrp>mvrp

Description This command enables the dampening timer and applies to both types of provisioned SAPs – end-station and UNI. When a value is configured for the timer, it controls the delay between detecting that the last provisioned SAP in VPLS goes down and reporting it to the MVRP module. The CPM will wait for the time specified in the value parameter before reporting it to the MVRP module. If the SAP comes up before the hold-timer expires, the event will not be reported to MVRP module.

The non-zero hold-time does not apply for SAP transition from down to up, This kind of transition is reported immediately to MVRP module without waiting for hold-time expiration. Also this parameter applies only to the provisioned SAPs. It does not apply to the SAPs configured with the **vpls-sap-template** command. Also when endstation QinQ SAPs are present only the “no hold-time” configuration is allowed.

The **no** form of this command disables tracking of the operational status for the last active SAP in the VPLS. MVRP will stop declaring the VLAN only when the last provisioned customer (UNI) SAP associated locally with the service is deleted. Also MVRP will declare the associated VLAN attribute as soon as the first provisioned SAP is created in the associated VPLS instance, regardless of the operational state of the SAP.

Default no hold-time

Parameters *value* — Specifies the hold time in minutes

Values 1 to 30

endstation-vid-group

Syntax **endstation-vid-group** *id* **vlan-id** *startvid-endvid*
no endstation-vid-group *id*

Context config>service>vpls>mrp>mvrp

Description This command specifies the range of VLAN IDs that are controlled by MVRP on the port associated with the parent SAP. When the command is present under a certain SAP, the MVRP will treat the associated virtual port as an endstation.

MVRP endstation behavior means that configuration of a new data SAP with the outer tag in the configured endstation-vid-group will generate down that virtual port a MVRP declaration for the new [outer] VLAN attribute. Also registration received for the VLAN attribute in the range will be accepted but not propagated in the rest of MVRP context.

VPLS-groups are not allowed under the associated Management VPLS (M-VPLS) when the endstation is configured under one SAP. VPLS-groups can be supported in the chassis using a different M-VPLS.

The **no** form of the command removes the specified group id.

Default no endstation-vid-group

Parameters *id* — Specifies the range index

Values 1 to 4094

startvid-endvid — Specifies the range of VLANs to be controlled by MVRP

Values 1 to 4094

root-guard

Syntax [**no**] **root-guard**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.

Default no root-guard

static-host

Syntax **static-host ip** *ip-address* [**mac** *ieee-address*] [**create**]
static-host mac *ieee-address* [**create**]

no static-host [**ip** *ip-address*] **mac** *ieee-address*
no static-host all [**force**]
no static-host ip *ip-address*

Context config>service>vpls>sap

Description This command configures a static host on this SAP.

Parameters *ip-address* — Specifies the IPv4 unicast address

ieee-address — Specify this optional parameter when defining a static host. Every static host definition must have at least one address defined, IP or MAC.

force — Specifies the forced removal of the static host addresses.

create — This keyword is mandatory while configuring a static host.

ancp-string

Syntax **ancp-string** *ancp-string*
no ancp-string

Context config>service>vpls>sap>static-host

Description This command specifies the ANCP string associated to this SAP host.

Parameters *ancp-string* — Specifies the ANCP string up to 63 characters in length

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>vpls>sap>static-host

Description This command specifies an application profile name.

Parameters *app-profile-name* — Specifies the application profile name up to 32 characters in length

inter-dest-id

Syntax **inter-dest-id** *intermediate-destination-id*
no inter-dest-id

Context config>service>vpls>sap>static-host

Description Specifies to which intermediate destination (for example a DSLAM) this host belongs.

Parameters *intermediate-destination-id* — Specifies the intermediate destination ID. 32 characters maximum

sla-profile

Syntax **sla-profile** *sla-profile-name*
no sla-profile

Context config>service>vpls>sap>static-host

Description This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

Parameters *sla-profile-name* — Specifies the SLA profile name

sub-profile

Syntax **sub-profile** *sub-profile-name*
no sub-profile

Context config>service>vpls>sap>static-host

Description This command specifies an existing subscriber profile name to be associated with the static subscriber host.

Parameters *sub-profile-name* — Specifies the sub-profile name

subscriber

Syntax **subscriber** *sub-ident*
no subscriber

Context config>service>vpls>sap>static-host

Description This command specifies an existing subscriber identification profile to be associated with the static subscriber host.

Parameters *sub-ident* — Specifies the subscriber identification

subscriber-sap-id

Syntax [**no**] **subscriber-sap-id**

Context config>service>vpls>sap>static-host

Description This command enables the use of the SAP ID as the subscriber ID.

The **no** form of the command disables the use of SAP ID as the subscriber ID.

trigger-packet

Syntax **trigger-packet** [**dhcp**] [**pppoe**] [**arp**] [**dhcp6**] [**ppp**] [**data**]
no trigger-packet

Context config>service>vpls>sap

Description This command enables triggering packet to initiate RADIUS authentication that provides a service context. The authentication, together with the service context for this request, creates a managed SAP. The VLAN is the same as the triggering packet. This SAP behaves as a regular SAP but the configuration is not user-editable and not maintained in the configuration file. The managed SAP remains active as long as the session is active.

Default none

Parameters **dhcp** — Specifies whether the receipt of DHCP trigger packets on this VPLS SAP when the keyword **capture-sap** is specified in the **sap** command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of **managed**

pppoe — Specifies whether the receipt of PPPoE trigger packets on this VPLS SAP when the keyword **capture-sap** is specified in the **sap** command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of **managed**

arp — Indicates that ARP is the type of trigger packets for this entry

dhcp6 — Indicates that DHCP6 is the type of trigger packets for this entry

ppp — Indicates that PPP is the type of trigger packets for this entry

data — Indicates that subscriber data packets are the type of trigger packets for this entry

vxlan

Syntax **vxlan vni** *vni-id* [**create**]
no vxlan vni *vni-id*

Context config>service>vpls

Description This command enables the use of vxlan in the VPLS service.

Default none

Parameters *vni-id* — Specifies the vxlan network identifier (VNI) configured in the VPLS service. All EVPN advertisements (mac routes and inclusive multicast routes) for this services will encode the configured VNI in the Ethernet Tag field of the NLRI.

Values 1 to 16777215

create — This keyword is mandatory when creating a vxlan instance

network

Syntax **network**

Context config>service>vpls>vxlan

Description This command enters the context to configure network parameters for the VPLS VXLAN service.

ingress

Syntax **ingress**

Context config>service>vpls>vxlan>network

Description This command enters the context to configure network ingress parameters for the VPLS VXLAN service.

qos

Syntax **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
no qos

Context config>service>vpls>vxlan>network>ingress

Description This command is used to redirect traffic arriving on VXLAN tunnels in an EVPN VXLAN service as a single entity (per forwarding class) to policers in an ingress forwarding plane queue group for the purpose of rate-limiting.

For the policer to be used, the following must be true:

- The configured queue group template name must be applied to the forwarding plane on which the ingress traffic arrives using the instance id specified.
- The policer referenced in the FC-to-policer mappings in the ingress context of a network QoS policy must be present in the specified queue group template.

The command will fail if the queue group template name does not exist or if the policer specified in the network QoS policy does not exist in the queue group template. If the queue group template name with the specified instance is not applied to the forwarding plane on which the VXLAN traffic arrives, then this traffic will use the ingress network queues related to the network interface; however, the ingress classification is still based on the applied network QoS policy.

The unicast traffic can be redirected to a policer under the forwarding class **fp-redirect-group** command in the ingress section of a network QoS policy. Similarly, broadcast, unknown and multicast traffic can be redirected to a **broadcast-policer**, **unknown-policer** or **mcast-policer**, respectively, also under the forwarding class **fp-redirect-group** command in the ingress section of a network QoS policy.

Ingress classification is based on the configuration of the ingress section of the specified network QoS policy, noting that the dot1p and DSCP classification is based on the outer Ethernet header and IP header, and the use of **ler-use-dscp**, **ip-criteria** and **ipv6-criteria** statements are ignored.

When this command is applied, it overrides the QoS applied to the related network interfaces for traffic arriving on VXLAN tunnels in that service but does not affect traffic received on a spoke-SDP in the same service.

The **no** version of this command removes the redirection of VXLAN tunnel traffic from the queue group policers.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

queue-group-name — Specifies the name of the queue group template up to 32 characters in length

instance-id — Specifies the identification of a specific instance of the queue-group

Values 1 to 65535

restrict-protected-src

Syntax **restrict-protected-src discard-frame**
no restrict-protected-src

Context config>service>vpls>vxlan

Description This command enables protected SRS MAC restrictions.

3.7.2.3 VPLS Interface Commands

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>service>vpls

Description This command creates a logical IP routing interface for a VPLS service. Once created, attributes such as IP address and service access points (SAP) can be associated with the IP interface.

The interface command, under the context of services, is used to create and maintain IP routing interfaces within the VPLS service IDs. The IP interface created is associated with the VPLS management routing instance.This instance does not support routing.

Interface names are case-sensitive and must be unique within the group of defined IP interfaces defined for the network core router instance. Interface names in the dotted decimal notation of an IP address are not allowed. For example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. Duplicate interface names can exist in different router instances.

Enter a new name to create a logical router interface. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, no default IP interface names are defined within the system. All VPLS IP interfaces must be explicitly defined in an enabled state.

The no form of this command removes the IP interface and the entire associated configuration. The interface must be administratively shutdown before issuing the no interface command.

For VPLS services, the IP interface must be shutdown before the SAP on that interface is removed.

For VPLS service, ping and traceroute are the only applications supported.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP.

An interface name:

- Should not be in the form of an IP address.
- Can be from 1 to 32 alphanumeric characters.
- If the string contains special characters (such as #,$,spaces), the entire string must be enclosed within double quotes.

If ip-int-name already exists within the service ID, the context changes to maintain that IP interface. If ip-int-name already exists within another service ID, an error occurs and the context does not change to that IP interface. If ip-int-name does not exist, the interface is created and the context is changed to that interface for further command processing.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*}
address *ip-address mask*

Context config>service>vpls>interface

Description This command assigns an IP address and an IP subnet, to a VPLS IP router interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each VPLS IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created. Use the no form of this command to remove the IP address assignment from the IP interface. When the no address command is entered, the interface becomes operationally down.

Address	Admin State	Oper State
No Address	Up	Down
No Address	Down	Down
1.1.1.1	Up	Up
1.1.1.1	Down	Down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up.

Parameters *ip-address* — The IP address of the IP interface. The ip-address portion of the address command specifies the IP host address that will be used by the IP interface within the subnet.

This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — The forward slash is a parameter delimiter and separates the ip-address portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the ip-address, the "/" and the mask-length parameter. If a forward slash is not immediately following the ip-address, a dotted decimal mask must follow the prefix.

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the ip-address from the mask-length parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. The values allowed are integers in the range 0 – 30. A mask length of 32 is reserved for system IP addresses.

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the ip-address from a traditional dotted decimal mask. The mask parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses.

3.7.2.4 General Switch Management Protocol Commands

gsmp

Syntax **gsmp**

Context config>service>vpls

Description This command enters the context to configure General Switch Management Protocol (GSMP) connections maintained in this service.

Default not enabled

group

Syntax [**no**] **group** ***name***

Context config>service>vpls>gsmp

Description This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.

ancp

Syntax **ancp**

Context config>service>vpls>gsmp>group

Description This command configures Access Node Control Protocol (ANCP) parameters for this GSMP group.

dynamic-topology-discover

Syntax [**no**] **dynamic-topology-discover**

Context config>service>vpls>gsmp>group>ancp

Description This command enables the ANCP dynamic topology discovery capability.

The **no** form of this command disables the feature.

idle-filter

Syntax [**no**] **idle-filter**

Context config>service>vpls>gsmp

Description This command when applied will filter out new subscriber's ANCP messages from subscriber with "DSL-line-state" IDLE.

Default no idle-filter

line-configuration

Syntax [**no**] **line-configuration**

Context config>service>vpls>gsmp>group>ancp

Description This command enables the ANCP line-configuration capability.

The **no** form of this command disables the feature.

oam

Syntax [**no**] **oam**

Context config>service>vpls>gsmp>group>ancp

Description This command specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection.

The **no** form of this command disables the feature.

hold-multiplier

Syntax **hold-multiplier** *multiplier*
no hold-multiplier

Context config>service>vpls>gsmp>group

Description This command configures the hold-multiplier for the GSMP connections in this group.

Parameters *multiplier* — Specifies the GSMP hold multiplier value

Values 1 to 100

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vpls>gsmp>group

Description This command configures keepalive values for the GSMP connections in this group.

Parameters *seconds* — Specifies the GSMP keepalive timer value in seconds

Values 1 to 25

neighbor

Syntax [**no**] **neighbo**r *ip-address*

Context config>service>vpls>gsmp>group

Description This command configures a GSMP ANCP neighbor.

Parameters *ip-address* — Specifies the IP address of the GSMP ANCP neighbor

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>service>vpls>gsmp>group>neighbor

Description This command configures the source ip-address used in the connection toward the neighbor. The local address is optional. If specified the node will accept connections only for that address in the service running ANCP. The address may be created after the reference but connections will not be accepted until it is created. If the local address is not used, the system accepts connections on any interface within the routing context.

Parameters *ip-address* — Specifies the source IP address to be used in the connection toward the neighbor

priority-marking

Syntax **priority-marking dscp** *dscp-name*
priority-marking prec *ip-prec-value*
no priority-marking

Context config>service>vpls>gsmp>group>neighbor

Description This command configures the type of priority marking to be used.

Parameters **dscp** *dscp-name* — Specifies the DSCP code-point to be used

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ed, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

prec ip-prec-value — Specifies the precedence value to be used

Values 0 to 7

persistency-database

Syntax **persistency-database**
no persistency-database

Context config>service>vpls>gsmp

Description This command enables the system to store DSL line information in memory. If the GSMP connection terminates, the DSL line information will remain in memory and accessible for Radius authentication and accounting.

Default no persistency-database

3.7.2.4.1 VPLS DHCP Commands

dhcp

Syntax **dhcp**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command enters the context to configure DHCP parameters.

lease-populate

Syntax **lease-populate** [*nmb-of-entries*]
no lease-populate

Context config>service>vpls>sap>dhcp

Description This command enables and disables dynamic host lease state management for VPLS SAPs. For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.

The optional number-of-entries parameter is used to define the number of lease state table entries allowed for this SAP or IP interface. If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

The retained lease state information representing dynamic hosts may be used to:

- populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding a new lease state entry or updating an existing lease state entry.
- generate dynamic ARP replies if **arp-reply-agent** is enabled.

Default no lease-populate

Parameters *nbr-of-entries* — Specifies the number of DHCP leases allowed

Values 1 to 8000

option

Syntax [**no**] **option**

Context config>service>vpls>sap>dhcp
config>service>vpls>sap>dhcp6

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command returns the system to the default.

Default no option

action

Syntax **action** [*dhcp-action*]
no action

Context config>service>vpls>sap>dhcp>option

Description This command configures the Relay Agent Information Option (Option 82) processing.

The **no** form of this command returns the system to the default value.

Default The default is to keep the existing information intact.

Parameters *dhcp-action* — Specifies the DHCP option action

replace — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (toward the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on toward the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **vlan-ascii-tuple**]
circuit-id hex [*0x0..0xFFFFFFFF...*(*64 hex nibbles*)]
no circuit-id

Context config>service>vpls>sap>dhcp>option

Description When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|". If no keyword is configured, then the circuit-id sub-option will not be part of the information option (Option 82).

When the command is configured without any parameters, it equals to **circuit-id ascii-tuple**.

If disabled, the **circuit-id** sub-option of the DHCP packet will be left empty.

Default no circuit-id

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used

hex — Specifies the circuit-id hex string

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits as well as what is included in ascii-tuple already. The format is supported on dot1q and qinq encapsulated ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

Syntax **remote-id** [**mac** | **string** *string*]
remote-id hex [*0x0..0xFFFFFFFF...*(*64 hex nibbles*)]
no remote-id

Context config>service>vpls>sap>dhcp>option
config>service>vpls>sap>dhcp6>option

Description This command specifies what information goes into the remote-id sub-option in the DHCP Relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

When the command is configured without any parameters, it equals to the **remote-id mac** option.

The **no** form of this command returns the system to the default.

Default no remote-id

Parameters **hex** — Specifies the remote-id hex string

mac — This keyword specifies the MAC address of the remote end is encoded in the sub-option

string *string* — Specifies the remote-id.

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>vpls>sap>dhcp>option

Description This command configures the vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

sap-id

Syntax [**no**] **sap-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies the string in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command returns the default value.

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

proxy-server

Syntax **proxy-server**

Context config>service>vpls>sap>dhcp

Description This command configures the DHCP proxy server.

emulated-server

Syntax **emulated-server** *ip-address*
no emulated-server

Context config>service>vpls>sap>dhcp>proxy

Description This command configures the IP address which will be used as the DHCP server address in the context of this VPLS SAP. Typically, the configured address should be in the context of the subnet represented by the VPLS.

The **no** form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.

Parameters *ip-address* — Specifies the emulated server address

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**radius-override**]
no lease-time

Context config>service>vpls>sap>dhcp>proxy

Description This command defines the length of lease time that will be provided to DHCP clients. By default, the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The **no** form of this command disables the use of the lease-time command. The local proxy server will use the lease time offered by either a RADIUS or DHCP server.

Default 7 days 0 hours 0 seconds

Parameters *days* — Specifies the number of days that the specified IP address is valid

Values 0 to 3650

hours — Specifies the number of hours that the specified IP address is valid

Values 0 to 23

minutes — Specifies the number of minutes that the specified IP address is valid

Values 0 to 59

seconds — Specifies the number of seconds that the specified IP address is valid.

Values 0 to 59

snoop

Syntax [**no**] **snoop**

Context config>service>vpls>sap>dhcp6
config>service>vpls>sap>dhcp
config>service>vpls>spoke-sdp>dhcp
config>service>vpls>mesh-sdp>dhcp

Description This command enables DHCP snooping of DHCP messages on the SAP or SDP. Enabling DHCP snooping on VPLS interfaces (SAPs and SDP bindings) is required where DHCP messages important to lease state table population are received, or where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.

Use the **no** form of the command to disable DHCP snooping on the specified VPLS SAP or SDP binding.

Default no snoop

3.7.2.4.2 VPLS STP Commands

stp

Syntax **stp**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>template>vpls-template

Description This command enters the context to configure the Spanning Tree Protocol (STP) parameters. Nokia's STP is simply the Spanning Tree Protocol (STP) with a few modifications to better suit the operational characteristics of VPLS services. The most evident change is to the root bridge election. Since the core network operating between Nokia's service routers should not be blocked, the root path is calculated from the core perspective.

auto-edge

Syntax **auto-edge**
no auto-edge

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures automatic detection of the edge port characteristics of the SAP or spoke-SDP.

If auto-edge is enabled, and STP concludes there is no bridge behind the spoke-SDP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see edge-port).

The **no** form of this command returns the auto-detection setting to the default value.

Default auto-edge

edge-port

Syntax [**no**] **edge-port**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the SAP or SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP, this value will be used only as the initial value.

Note: The function of the **edge-port** command is similar to the **rapid-start** command. It tells RSTP that it is on the edge of the network (for example, there are no other bridges connected to that port) and, as a consequence, it can immediately transition to a forwarding state if the port becomes available.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of the SAP or spoke-SDP parameter is set to edge-port. This value will change if:

- A BPDU is received on that port. This means that after all there is another bridge connected to this port. Then the edge-port becomes disabled.
- If auto-edge is configured and no BPDU is received within a certain period of time, RSTP concludes that it is on an edge and enables the edge-port.

The **no** form of this command returns the edge port setting to the default value.

Default no edge-port

forward-delay

Syntax **forward-delay** *seconds*
no forward-delay

Context config>service>vpls>stp
config>service>template>vpls-template>stp

Description RSTP, as defined in the IEEE 802.1D-2004 standards, will normally transition to the forwarding state via a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (e.g. on shared links, see below), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two nodes (for example, a shared 10/100BaseT segment). The port-type command is used to configure a link as point-to-point or shared.

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP or spoke-SDP spends in the discarding and learning states when transitioning to the forwarding state.

The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in rstp or mstp mode, but only when the SAP or spoke-SDP has not fallen back to legacy STP operation, the value configured by the hello-time command is used;
- in all other situations, the value configured by the forward-delay command is used.

Default 15 seconds

Parameters *seconds* — The forward delay timer for the STP instance in seconds

Values 4 to 30

hello-time

Syntax **hello-time** *hello-time*
no hello-time

Context config>service>vpls>stp
config>service>template>vpls-template>stp

Description This command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The hello time parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time can also be used to calculate the forward delay. See auto-edge.

The **no** form of this command returns the hello time to the default value.

Default 2 seconds

Parameters *hello-time* — The hello time for the STP instance in seconds

Values 1 to 10

hold-count

Syntax **hold-count** *BDPU tx hold count*
no hold-count

Context config>service>vpls>stp
config>service>template>vpls-template>stp

Description This command configures the peak number of BPDUs that can be transmitted in a period of one second.

The **no** form of this command returns the hold count to the default value

Default 6

Parameters *BDPU tx hold count* — The hold count for the STP instance in seconds

Values 1 to 10

link-type

Syntax **link-type** {**pt-pt** | **shared**}
no link-type

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command instructs STP on the maximum number of bridges behind this SAP or spoke-SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP or spoke-SDPs should all be configured as shared, and timer-based transitions are used.

The **no** form of this command returns the link type to the default value.

Default pt-pt

mst-instance

Syntax **mst-instance** *mst-inst-number*

Context config>service>vpls>sap>stp

Description This command enters the context to configure MSTI related parameters at SAP level. This context can be open only for existing mst-instances defined at the service level (see mst-instance).

Default none

Parameters *mst-inst-number* — Specifies an existing Multiple Spanning Tree Instance number

Values 1 to 4094

mst-path-cost

Syntax **mst-path-cost** *inst-path-cost*
no mst-path-cost

Context config>service>vpls>sap>stp>mst-instance

Description This commands specifies path-cost within a specified instance, expressing probability that a specified port will be put into the forwarding state in case a loop occurs (the highest value expresses lowest priority).

The **no** form of this command sets port-priority to its default value.

Default The path-cost is proportional to link speed.

Parameters *inst-path-cost* — Specifies the contribution of this port to the MSTI path cost of paths toward the spanning tree regional root which include this port

Values 1 to 200000000

mst-priority

Syntax **mst-priority** *stp-priority*
no mst-priority

Context config>service>vpls>sap>stp>mst-instance

Description This commands specifies the port priority within a specified instance, expressing probability that a specified port will be put into the forwarding state if a loop occurs.

The **no** form of this command sets port-priority to its default value.

Default 128

Parameters *stp-priority* — Specifies the value of the port priority field.

max-age

Syntax **max-age** *seconds*
no max-age

Context config>service>vpls>stp
config>service>template>vpls-template>stp

Description This command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge will take the message_age value from BPDUs received on their root port and increment this value by 1. The message_age therefore reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.

The **no** form of this command returns the max age to the default value.

Default 20 seconds

Parameters *seconds* — The max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40.

mode

Syntax **mode** {*rstp* | *comp-dot1w* | *dot1w* | *mstp* | *pmstp*}
no mode

Context config>service>vpls>stp
config>service>template>vpls-template>stp

Description This command specifies the version of Spanning Tree Protocol the bridge is currently running.

See section Spanning Tree Operating Modes for details on these modes.

The **no** form of this command returns the STP variant to the default.

Default rstp

Parameters **rstp** — Corresponds to the Rapid Spanning Tree Protocol specified in IEEE 802.1D/D4-2003

dot1w — Corresponds to the mode where the Rapid Spanning Tree is backward compatible with IEEE 802.1w

compdot1w — Corresponds to the Rapid Spanning Tree Protocol in conformance with IEEE 802.1w

mstp — Sets MSTP as the STP mode of operation. Corresponds to the Multiple Spanning Tree Protocol specified in 802.1Q REV/D5.0-09/200

pmstp — The PMSTP mode is only supported in VPLS services where the mVPLS flag is configured

mst-instance

Syntax [**no**] **mst-instance** *mst-inst-number*

Context config>service>vpls>stp

Description This command creates the context to configure MST instance (MSTI) related parameters. Up to 16 instances will be supported by MSTP. The instance 0 is mandatory by protocol and therefore, it cannot be created by the CLI. The software will maintain this instance automatically.

Default none

Parameters *mst-inst-number* — Specifies the Multiple Spanning Tree instance

Values 1 to 4094

mst-priority

Syntax **mst-priority** *bridge-priority*
no mst-priority

Context config>service>vpls>stp>mst-instance

Description This command specifies the bridge priority for this specific Multiple Spanning Tree Instance for this service. The *bridge-priority* value reflects likelihood that the switch will be chosen as the regional root switch (65535 represents the least likely). It is used as the highest 4 bits of the Bridge ID included in the MSTP BPDUs generated by this bridge.

The priority can only take on values that are multiples of 4096 (4k). If a value is specified that is not a multiple of 4K, then the value will be replaced by the closest multiple of 4K, which is lower than the value entered.

The **no** form of this command sets the bridge-priority to its default value.

Default 32768 — All instances created by vlan-range command and not having explicit definition of bridge-priority will inherit default value

Parameters *bridge-priority* — Specifies the priority of this specific Multiple Spanning Tree Instance for this service

Values 0 to 65535

vlan-range

Syntax [**no**] **vlan-range** [*vlan-range*]

Context config>service>vpls>stp>mst-instance

Description This command specifies a range of VLANs associated with a certain MST-instance. This range applies to all SAPs of the mVPLS.

Every VLAN range that is not assigned within any of the created mst-instance is automatically assigned to mst-instance 0. This instance is automatically maintained by the software and cannot be modified. Changing the VLAN range value can be performed only when the specified mst-instance is shutdown.

The **no** form of this command removes the **vlan-range** from the specified mst-instance.

Parameters *vlan-range* — The first VLAN range specifies the left-bound (i.e., minimum value) of a range of VLANs that are associated with the mVPLS SAP. This value must be smaller than (or equal to) the second VLAN range value. The second VLAN range specifies the right-bound (i.e., maximum value) of a range of VLANs that are associated with the mVPLS SAP.

Values 1 to 4094 to 1 to 4094

mst-max-hops

Syntax **mst-max-hops** *hops-count*
no mst-max-hops

Context config>service>vpls>stp

Description This command specifies the number of hops in the region before BPDU is discarded and the information held for the port is aged out. The root bridge of the instance sends a BPDU (or M-record) with remaining-hop-count set to configured <*max-hops*>. When a bridge receives the BPDU (or M-record), it decrements the received remaining-hop-count by 1 and propagates it in BPDU (or M-record) it generates.

The **no** form of this command sets the *hops-count* to its default value.

Default 20

Parameters *hops-count* — Specifies the maximum number of hops

Values 1 to 40

mst-name

Syntax **mst-name** *region-name*
no mst-name

Context config>service>vpls>stp

Description This command defines an MST region name. Two bridges are considered as a part of the same MST region as soon as their configuration of the MST region name, the MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command removes *region-name* from the configuration.

Default no mst-name

Parameters *region-name* — Specifies an MST-region name up to 32 characters in length

mst-revision

Syntax **mst-revision** *revision-number*

Context config>service>vpls>stp

Description This command defines the MST configuration revision number. Two bridges are considered as a part of the same MST region as soon as their configuration of MST-region name, MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command returns MST configuration revision to its default value.

Default 0

Parameters *revision-number* — Specifies the MSTP region revision number to define the MSTP region

Values 0 to 65535

path-cost

Syntax **path-cost** *sap-path-cost*
no path-cost

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the Spanning Tree Protocol (STP) path cost for the SAP or spoke-SDP.

The path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP or spoke-SDP. When BPDUs are sent out other egress SAPs or spoke-SDPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs and spoke-SDPs are controlled by complex queuing dynamics, in the 7450 ESS, 7750 SR, and 7950 XRS the STP path cost is a purely static configuration.

The **no** form of this command returns the path cost to the default value.

Parameters *path-cost* — The path cost for the SAP or spoke-SDP

Values 1 to 200000000 (1 is the lowest cost)

Default 10

port-num

Syntax [**no**] **port-num** *virtual-port-number*

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the virtual port number which uniquely identifies a SAP within configuration bridge protocol data units (BPDUs). The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with it's own virtual port number that is unique to every other SAP defined on the TLS. The virtual port number is assigned at the time that the SAP is added to the TLS. Since the order that the SAP was added to the TLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance.

The virtual port number cannot be administratively modified.

priority

Syntax **priority** *bridge-priority*
no priority

Context config>service>vpls>stp
config>service>template>vpls-template>stp

Description The bridge-priority command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

The **no** form of this command returns the bridge priority to the default value.

Default By default, the bridge priority is configured to 4096 which is the highest priority.

Parameters *bridge-priority* — The bridge priority for the STP instance

Values Allowed values are integers in the range of 4096 to 65535 with 4096 being the highest priority. The actual bridge priority value stored/ used is the number entered with the lowest 12 bits masked off which means the actual range of values is 4096 to 61440 in increments of 4096.

priority

Syntax **priority** *stp-priority*
no priority

Context config>service>vpls>spoke-sdp
config>service>vpls>sap>stp

Description This command configures the Nokia Spanning Tree Protocol (STP) priority for the SAP or spoke-SDP.

STP priority is a configurable parameter associated with a SAP or spoke-SDP. When configuration BPDUs are received, the priority is used in some circumstances as a tie breaking mechanism to determine whether the SAP or spoke-SDP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP or spoke-SDP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP or spoke-SDP within the STP instance.

STP computes the actual priority by taking the input value and masking out the lower four bits.The result is the value that is stored in the SDP priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

The **no** form of this command returns the STP priority to the default value.

Default 128

Parameters *stp-priority* — The STP priority value for the SAP or spoke-SDP. Allowed values are integer in the range of 0 to 255, 0 being the highest priority. The actual value used for STP priority (and stored in the configuration) will be the result of masking out the lower 4 bits, therefore the actual value range is 0 to 240 in increments of 16.

Default 128

3.7.2.4.3 VPLS SAP Commands

sap

Syntax **sap** *sap-id* [**split-horizon-group** *group-name*] [**capture-sap**] [**create**] [**eth-ring** *ring-index*] [**root-leaf-tag** | **leaf-ac]**
no sap *sap-id*

Context config>service>vpls

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7450 ESS, 7750 SR, and 7950 XRS. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Enhanced Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Default No SAPs are defined.

Special Cases **VPLS service SAP limits —** A VPLS SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. The limits of the number of SAPs and SDPs supported in a VPLS service depends on the hardware used. Each SDP must have a unique destination or an error will be generated. Split horizon groups can only be created in the scope of a VPLS service.

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS). This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0).

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

create — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

eth-ring — When used with Ethernet Rings control the split horizon group accepts the major ring instance "value". This parameter applies to the 7450 ESS or 7750 SR only. The split horizon group prevents loops in the cases where a Ethernet Virtual Ring is miss configured on the main ring. Each path a and path b major ring are configured in the group and associated with the sub-ring control instance in the VPLS service.

ring-index — Specifies the ring index of the Ethernet ring

root-leaf-tag — Specifies a SAP as a root leaf tag SAP. Only SAPs of the form dot1q (for example, 1/1/1:X) or qinq (for example, 1/1/1:X.Y, 1/1/1:X.*) are supported. The default E-Tree SAP type is a root AC, if *root-leaf-tag* **(**or *leaf-ac***)** is not specified at SAP creation. This option is only available when the VPLS is designated as an E-Tree VPLS; it is not available on BGP EVPN-enabled E-Tree VPLS services.

leaf-tag-vid — Specified after *root-leaf-tag* to replace the outer SAP-ID for leaf traffic. The leaf tag VID is only significant between peering VPLS but the values must be consistent on each end. This option is not available on BGP EVPN-enabled E-Tree VPLS services.

leaf-ac — Specifies a SAP as a leaf access (AC) SAP. The default E-Tree SAP type is root AC if *leaf-ac* (or *root-leaf-tag*) is not specified at SAP creation. This option is available when the VPLS is designated as an E-Tree VPLS. BGP EVPN-enabled E-Tree VPLS services also support the leaf-ac option.

split-horizon-group *group-name* — Specifies the name of the split horizon group to which the SAP belongs. This parameter applies to the 7450 ESS or 7750 SR only.

capture-sap — Specifies a capturing SAP in which triggering packets will be sent to the CPM. Non-triggering packets captured by the capture SAP will be dropped. This parameter applies to the 7450 ESS or 7750 SR only.

cflowd

Syntax [**no**] **cflowd**

Context config>service>vpls>sap

Description This command enables cflowd to collect traffic flow samples through a service interface (SAP) for analysis. When cflowd is enabled on an Ethernet service SAP, the Ethernet traffic can be sampled and processed by the system's cflowd engine and exported to IPFIX collectors with the l2-ip template enabled.

cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When cflowd is enabled at the SAP level, all packets forwarded by the interface are subjected to analysis according to the cflowd configuration.

For Layer 2 services, only ingress sampling is supported.

Default no cflowd

discard-unknown-source

Syntax [**no**] **discard-unknown-source**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>template>vpls-sap-template

Description When this command is enabled, packets received on a SAP or a spoke-SDP with an unknown source MAC address will be dropped only if the maximum number of MAC addresses for that SAP or spoke-SDP (see max-nbr-mac-addr) has been reached. If **max-nbr-mac-addr** has not been set for the SAP or spoke-SDP, enabling **discard-unknown-source** has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.

Default no discard-unknown-source

3.7.2.5 ETH-CFM Service Commands

eth-cfm

Syntax **eth-cfm**

Context config>service>vpls
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command enters the context to configure ETH-CFM parameters.

eth-tunnel

Syntax **eth-tunnel**

Context config>service>vpls>sap

Description The command enables the context to configure Ethernet tunnel SAP parameters.

eth-ring

Syntax **eth-ring** *ring-id*
no eth-ring

Context config>service>vpls

Description This command configures a VPLS SAP to be associated with an Ethernet ring. The SAP port ID is associated with the corresponding Ethernet ring path configured on the same port ID. The encapsulation type must be compatible with the Ethernet ring path encapsulation.

The **no** form of this command removes the Ethernet ring association from this SAP.

Default no eth-ring

Parameters *ring-id* — Specifies the ring ID.

Values 1 to 128

path

Syntax **path** *path-index* **tag** *qtag*[.*qtag*]
no path *path-index*

Context config>service>vpls>sap>eth-tunnel

Description This command configures Ethernet tunnel SAP path parameters.

The **no** form of the command removes the values from the configuration.

Default none

Parameters *path-index* — Specifies the path index value.

Values 1 to 16

tag *qtag*[.*qtag*] — Specifies the qtag value.

Values 0 to 4094, * (wildcard)

collect-lmm-fc-stats

Syntax **collect-lmm-fc-stats**

Context config>service>vpls>mesh-sdp>eth-cfm
config>service>vpls>sap>eth-cfm
config>service>vpls>spoke-sdp>eth-cfm

Description This command enters the context to configure per-forwarding class (FC) LMM information collection.

This command is mutually exclusive with the **collect-lmm-stats** command when there is entity resource contention.

fc

Syntax **fc** *fc-name* [*fc-name* ... (up to 8 max)]
no fc

Context config>service>vpls>mesh-sdp>eth-cfm>collect-lmm-fc-stats
config>service>vpls>sap>eth-cfm>collect-lmm-fc-stats
config>service>vpls>spoke-sdp>eth-cfm>collect-lmm-fc-stats

Description This command creates individual counters for the specified FCs without regard for profile. All countable packets that match a configured FC, regardless of profile, will be included in this counter.

A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.

Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-aware FC using the **fc-in-profile** command under the same context.

The **no** form of the command removes all previously defined FCs and stops counting for those FCs.

Default no fc

Parameters *fc-name* — Specifies the name of the FC for which to create an individual profile-unaware counter. In order for the counter to be used, the **config**>**oam-pm**>**session**>**ethernet**>**priority** command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the **config**>**oam-pm**>**session**>**ethernet**>**lmm**>**enable-fc-collection** command must be enabled.

Values nc, h1, ef, h2, l1, af, l2, be

fc-in-profile

Syntax **fc-in-profile** *fc-name* [*fc-name* ... (up to 8 max)]
no fc-in-profile

Context config>service>vpls>mesh-sdp>eth-cfm>collect-lmm-fc-stats
config>service>vpls>sap>eth-cfm>collect-lmm-fc-stats
config>service>vpls>spoke-sdp>eth-cfm>collect-lmm-fc-stats

Description This command creates individual counters for the specified FCs with regard for profile. All countable packets that match a configured FC and are deemed to be in-profile will be included in this counter.

A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.

Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-unaware FC using the **fc** command under the same context.

The **no** form of the command removes all previously defined FCs and stops counting for those FCs.

Default no fc-in-profile

Parameters *fc-name* — Specifies the name of the FC for which to create an individual profile-aware counter. In order for the counter to be used, the **config**>**oam-pm**>**session**>**ethernet**>**priority** command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the **config**>**oam-pm**>**session**>**ethernet**>**lmm**>**enable-fc-collection** command must be enabled.

Values nc, h1, ef, h2, l1, af, l2, be

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}] [**primary-vlan-enable**]
no mep *mep-id* **domain** *md-index* **association** *ma-index*

Context config>service>vpls>mesh-sdp>eth-cfm
config>service>vpls>spoke-sdp>eth-cfm
config>service>vpls>eth-cfm
config>service>vpls>sap>eth-cfm

Description This command configures the ETH-CFM maintenance endpoint (MEP). A MEP created at the VPLS service level **vpls>eth-cfm** creates a virtual MEP.

The **no** version of the command will remove the MEP.

Parameters *mep-id* — Specifies the MEP identifier.

Values 1 to 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 to 4294967295

ma-index — Specifies the maintenance association (MA) index value.

Values 1 to 4294967295

direction up | down — Indicates the direction in which the MEP faces on the bridge port. Direction is not supported when a MEP is created directly under the vpls>eth-cfm construct (vMEP).

down — Sends ETH-CFM messages away from the MAC relay entity

up — Sends ETH-CFM messages toward the MAC relay entity

primary-vlan-enable — Provides a method for linking the MEP with the primary VLAN configured under the bridge-identifier for the MA. MEPs cannot be changed from or to primary VLAN functions. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Layer 2 Epipe and VPLS services.

mip

Syntax **mip** [**mac** *mac-address*] [**primary-vlan-enable** *vlan-id*]
mip default-mac
no mip

Context config>service>vpls>sap>eth-cfm
config>service>vpls>spoke-sdp>eth-cfm
config>service>vpls>mesh-sdp>eth-cfm

Description This command allows Maintenance Intermediate Points (MIPs). The creation rules of the MIP are dependent on the mhf-creation configuration for the MA.

The **no** form of the command removes the MIP creation request.

Default no mip

Parameters *mac-address* — Specifies the MAC address of the MEP.

Values 6-byte MAC address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MIP. The MAC must be unicast. Using the all-zeros address is equivalent to the **no** form of this command.

default-mac — Using the **no** command deletes the MIP. If the operator wants to change the MAC back to the default MAC without having to delete the MIP and reconfiguring, this command is useful.

primary-vlan-enable — Provides a method for linking the MIP with the primary VLAN configured under the bridge-identifier for the MA. This is only allowed if the mhf-creation method is static. MIPs cannot be changed from or to primary VLAN functions without first being deleted. Primary VLANs are only supported under Layer 2 Epipe and VPLS services.

vlan-id — Must match the vlan-id under the bridge-identifier for the MA that is appropriate for this service.

Values 0 to 4094

mip

Syntax **mip primary-vlan-enable** [**vlan** *vlan-id*]
no mip

Context config>service>template>vpls-sap-template>eth-cfm

Description This command allows Maintenance Intermediate Points (MIPs). The creation rules of the MIP are dependent on the mhf-creation configuration for the MA. This MIP option is only available for default and static mhf-creation methods.

Parameters **primary-vlan-enable** — Provides a method for linking the MIP with the primary VLAN configured under the bridge-identifier for the MA. This is only allowed if the mhf-creation method is static. MIPs cannot be changed from or to primary VLAN functions without first being deleted. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs.

vlan — A required parameter when including primary-vlan-enable. Provides a method for associating the VLAN under the bride-identifier under the MA with the MIP.

vlan-id — Must match the vlan-id under the bridge-identifier for the MA that is appropriate for this service.

Values 0 to 4094

ais-enable

Syntax [**no**] **ais-enable**

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command enables the generation and the reception of AIS messages.

interface-support-enable

Syntax [**no**] **interface-support-enable**

Context config>service>vpls>sap>eth-cfm>mep>ais
config>service>vpls>spoke-sdp>eth-cfm>mep>ais
config>service>vpls>mesh-sdp>eth-cfm>mep>ais

Description This command enables the AIS function to consider the operational state of the entity on which it is configured. With this command, ETH-AIS on operationally down MEPs will be triggered and cleared based on the operational status of the entity on which it is configured. If CCM is also enabled then transmission of the AIS PDU will be based on either the non-operational state of the entity or on any CCM defect condition. AIS generation will cease if BOTH operational state is UP and CCM has no defect conditions. If the MEP is not CCM enabled then the operational state of the entity is the only consideration assuming this command is present for the MEP.

Default no interface-support-enabled (AIS will not be generated or stopped based on the state of the entity on) which the operationally down MEP is configured.

client-meg-level

Syntax **client-meg-level** [[*level* [*level* ...]]

no client-meg-level

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>ais-enable

Description This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters *level* — Specifies the client MEG level

Values 1 to 7

Default 1

ccm-padding-size

Syntax **ccm-padding-size** *ccm-padding*
no ccm-padding-size *ccm-padding*

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description Set the byte size of the optional Data TLV to be included in the ETH-CC PDU. This will increase the size of the ETH-CC PDU by the configured value. The base size of the ETH-CC PDU, including the Interface Status TLV and Port Status TLV, is 83 bytes not including the Layer Two encapsulation. CCM padding is not supported when the CCM-Interval is less than one second.

Default no ccm-padding-size

Parameters *ccm-padding* — Specifies the byte size of the Optional Data TLV.

Values 3 to 1500

csf-enable

Syntax [**no**] **csf-enable**

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command enables the reception and local processing of ETH-CSF frames.

multiplier

Syntax **multiplier** *multiplier-value*
no multiplier

Context config>service>vpls>mesh-sdp>eth-cfm>mep>cfs-enable
config>service>vpls>sap>eth-cfm>mep>cfs-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>cfs-enable

Description This command enables the multiplication factor applied to the receive time used to clear the CSF condition in increments of .5.

Default 3.5

Parameters *multiplier-value* — Specifies the multiplier used for timing out CSF

Values 0.0, 2.0 to 30.0

interval

Syntax **interval** {**1** | **60**}
no interval

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>ais-enable

Description This command specifies the transmission interval of AIS messages in seconds.

Parameters **1 | 60** — The transmission interval of AIS messages in seconds

Default 1

priority

Syntax **priority** *priority-value*
no priority

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>ais-enable

Description This command specifies the priority of AIS messages originated by the node.

Parameters *priority-value* — Specifies the priority value of the AIS messages originated by the node

Values 0 to 7

Default 1

ccm-enable

Syntax [**no**] **ccm-enable**

Context config>service>vpls>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

Context config>service>vpls>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command specifies the priority value for CCMs and LTMs transmitted by the MEP.

The **no** form of the command removes the priority value from the configuration.

Default The highest priority on the bridge-port.

Parameters *priority* — Specifies the priority of CCM and LTM messages

Values 0 to 7

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep

Description For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**data-length** *data-length*]

A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
no test-pattern

Context config>service>vpls>sap>eth-cfm>mep>eth-test-enable
config>service>vpls>spoke-sdp>eth-cfm>mep>eth-test-enable
config>service>vpls>mesh-sdp>eth-cfm>mep>eth-test-enable

Description This command configures the test pattern for eth-test frames.

The **no** form of the command removes the values from the configuration.

Parameters **all-zeros** — Specifies to use all zeros in the test pattern

all-ones — Specifies to use all ones in the test pattern

crc-enable — Generates a CRC checksum

Default all-zeros

bit-error-threshold

Syntax **bit-error-threshold** *bit-errors*

Context config>service>vpls>mesh-sdp

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default 1

Parameters *bit-errors* — Specifies the lowest priority defect.

Values 0 to 11840

fault-propagation-enable

Syntax **fault-propagation-enable** {**use-if-tlv** | **suspend-ccm**}
no fault-propagation-enable

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command configures the fault propagation for the MEP.

Parameters **use-if-tlv** — Specifies to use the interface TLV.

suspend-ccm — Specifies to suspend the continuity check messages.

grace

Syntax **grace**

Context config>service>vpls>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep

Description This command enters the context to configure Nokia ETH-CFM Grace and ITU-T Y.1731 ETH-ED expected defect functional parameters.

eth-ed

Syntax **eth-ed**

Context config>service>vpls>eth-cfm>mep>grace
config>service>vpls>mesh-sdp>eth-cfm>mep>grace
config>service>vpls>spoke-sdp>eth-cfm>mep>grace
config>service>vpls>sap>eth-cfm>mep>grace

Description This command enters the context to configure ITU-T Y.1731 ETH-ED expected defect functional parameters.

max-rx-defect-window

Syntax **max-rx-defect-window** *seconds*
no max-rx-defect-window

Context config>service>vpls>eth-cfm>mep>grace>eth-ed
config>service>vpls>mesh-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>sap>eth-cfm>mep>grace>eth-ed

Description This command limits the duration of the received ETH-ED expected defect window to the lower value of either the received value from the peer or this parameter.

The **no** form of the command removes the limitation, and any valid defect window value received from a peer MEP in the ETH-ED PDU will be used.

Default no max-rx-defect-window

Parameters *seconds* — Specifies the duration, in seconds, of the maximum expected defect window.

Values 1 to 86400

priority

Syntax **priority** *priority*
no priority

Context config>service>vpls>eth-cfm>mep>grace>eth-ed
config>service>vpls>mesh-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>sap>eth-cfm>mep>grace>eth-ed

Description This command sets the priority bits and determines the forwarding class based on the mapping of priority to FC.

The **no** form of the command disables the local priority configuration and sets the priority to the **ccm-ltm-priority** associated with this MEP.

Default no priority

Parameters *priority* — Specifies the priority bit.

Values 0 to 7

rx-eth-ed

Syntax [**no**] **rx-eth-ed**

Context config>service>vpls>eth-cfm>mep>grace>eth-ed
config>service>vpls>mesh-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>sap>eth-cfm>mep>grace>eth-ed

Description This command enables the reception and processing of the ITU-T Y.1731 ETH-ED PDU on the MEP.

The **no** form of the command disables the reception of the ITU-T Y.1731 ETH-ED PDU on the MEP.

Default rx-eth-ed

tx-eth-ed

Syntax [**no**] **tx-eth-ed**

Context config>service>vpls>eth-cfm>mep>grace>eth-ed

config>service>vpls>mesh-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>vpls>sap>eth-cfm>mep>grace>eth-ed

Description This command enables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP when a system soft reset notification is received for one or more cards.

The **config**>**eth-cfm**>**system**>**grace-tx-enable** command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP.

Default no tx-eth-ed

eth-vsm-grace

Syntax **eth-vsm-grace**

Context config>service>vpls>eth-cfm>mep>grace
config>service>vpls>mesh-sdp>eth-cfm>mep>grace
config>service>vpls>spoke-sdp>eth-cfm>mep>grace
config>service>vpls>sap>eth-cfm>mep>grace

Description This command enters the context to configure Nokia ETH-CFM Grace functional parameters.

rx-eth-vsm-grace

Syntax [**no**] **rx-eth-vsm-grace**

Context config>service>vpls>eth-cfm>mep>grace>eth-vsm-grace
config>service>vpls>mesh-sdp>eth-cfm>mep>grace>eth-vsm-grace
config>service>vpls>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace
config>service>vpls>sap>eth-cfm>mep>grace>eth-vsm-grace

Description This command enables the reception and processing of the Nokia ETH-CFM Grace PDU on the MEP.

The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.

The **no** form of the command disables the reception of the Nokia ETH-CFM Grace PDU on the MEP.

Default rx-eth-vsm-grace

tx-eth-vsm-grace

Syntax [**no**] **tx-eth-vsm-grace**

Context config>service>vpls>eth-cfm>mep>grace>eth-vsm-grace
config>service>vpls>mesh-sdp>eth-cfm>mep>grace>eth-vsm-grace
config>service>vpls>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace
config>service>vpls>sap>eth-cfm>mep>grace>eth-vsm-grace

Description This command enables the transmission of the Nokia ETH-CFM Grace PDU from the MEP when a system soft reset notification is received for one or more cards.

The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.

The **config**>**eth-cfm**>**system**>**grace-tx-enable** command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the Nokia ETH-CFM Grace PDU from the MEP.

Default tx-eth-vsm-grace

lbm-svc-act-responder

Syntax [**no**] **lbm-svc-act-responder**

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command enables the MEP to process service activation streams encapsulated in ETH-CFM LBM frames that are directed to the MEP. The MEP will be allocated additional resources to rapidly respond to a high-speed stream of LBM messages. A MEP created with this option will not validate any TLVs, will not validate the ETH-LBM MAC Address, and will not increment or compute any loopback statistics. Statistical computation and reporting is the responsibility of the test head-end. The ETH-CFM level of the high speed ETH-LBM stream must match the level of a MEP configured with this command. It must not target any lower ETH-CFM level the MEP will terminate. When the service activation test is complete, the MEP may be returned to standard processing by removing this command. If there is available bandwidth, the MEP will respond to other ETH-CFM PDUs, such as ETH-DMM marker packets, using standard processing.

The interaction between this command and the **tools perform service id** *service-id* **loopback eth** command must be carefully considered. It is recommended that either the **lbm-svc-act-responder** or the **tools perform service id** *service-id* **loopback eth** command be used at any given time within a service. If both commands must be configured, and the target reflection point is the MAC Swap Loopback function, the inbound stream of data must not include ETH-CFM traffic that is equal to or lower than the domain level of any configured MEP which would otherwise extract and process the ETH-CFM message. If the reflection target is a MEP configured with the **lbm-svc-act-responder** option, the mode (ingress or egress) of the SAP or SDP specified with this tools command and the MEP **direction** (up or down) must match when the functions are enabled on the same reflection point, and the domain level of the inbound ETH-LBM must be the same as that of the MEP configured with the **lbm-svc-act-responder** option. At no time should the two functions be conflicting with each other along the path of the stream. This conflict would lead to unpredictable and possibly destabilizing situations.

The **no** form of the command reverts to MEP LBM standard processing.

Default no lbm-svc-act-responder

low-priority-defect

Syntax **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default macRemErrXcon

Parameters **low-priority-defect** — The low priority defect values are defined below.

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
errXcon	Only DefErrorCCM and DefXconCCM
xcon	Only DefXconCCM; or
noXcon	No defects DefXcon or lower are to be reported

mac-address

Syntax **mac-address** *mac-address*

no mac-address

Context config>service>vpls>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep

Description This command specifies the MAC address of the MEP.

The **no** form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke).

Parameters *mac-address* — Specifies the MAC address of the MEP

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MEP. Must be unicast. Using the all zeros address is equivalent to the no form of this command.

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>service>vpls>sap>eth-cfm>mep

Description This command enables/disables eth-test functionality on MEP.

Parameters *seconds* — Specifies the one way delay threshold, in seconds

Values 0 to 600

Default 3

tunnel-fault

Syntax **tunnel-fault** {**accept** | **ignore**}

Context config>service>vpls>eth-cfm
config>service>vpls>sap>eth-cfm

Description Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, Epipe will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an Epipe service as well as setting the operational flag. If AIS generation is the requirement for the Epipe services this command is not required. See the command ais-enable under epipe>sap>eth-cfm>ais-enable for more details. This works in

conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is “ignore” and the sap level command is “accept”. This means simply changing the service level command to “accept” will enable the feature for all SAPs. This is not required for Epipe services that only wish to generate AIS on failure.

Default ignore (Service Level)

accept (SAP Level for Epipe and VPLS)

Parameters *accept* — Specifies to share fate with the facility tunnel MEP

ignore — Does not share fate with the facility tunnel MEP

vmep-extensions

Syntax [**no**] **vmep-extensions**

Context config>service>vpls>eth-cfm

Description This command enables and disables enhanced Virtual Maintenance Endpoints functionality. This must manually be configured for a B-VPLS to change the legacy behavior and cannot be disable for VPLS contexts that are not B-VPLS based.

The **no** form of the command reverts to the default values. This is not applicable to a VPLS contexts that is not B-VPLS based.

Default no vmep-extensions (for B-VPLS)

vmep-extensions (for VPLS contexts not B-VPLS based)

vmep-filter

Syntax [no] **vmep-**filter

Context config>service>vpls>eth-cfm>sap
config>service>vpls>eth-cfm>spoke-sdp
config>service>vpls>eth-cfm>mesh-sdp

Description Suppress eth-cfm PDUs based on level lower than or equal to configured Virtual MEP. This command is not supported under a B-VPLS context. This will also delete any MIP configured on the SAP or Spoke-SDP.

The **no** form of the command reverts to the default values.

Default no vmep-filter

limit-mac-move

Syntax **limit-mac-move** [**blockable** | **non-blockable**]
no limit-mac-move

Context config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command indicates whether or not the mac-move agent, when enabled using **config>service>vpls>mac-move** or **config>service>epipe>mac-move**, will limit the MAC re-learn (move) rate on this SAP.

Default blockable

Parameters **blockable** — Specifies the agent will monitor the MAC re-learn rate on the SAP, and it will block it when the re-learn rate is exceeded

non-blockable — Specifies the that this SAP will not be blocked, and another blockable SAP will be blocked instead

mac-pinning

Syntax [**no**] **mac-pinning**

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>sap
config>service>vpls>endpoint

Description Enabling this command will disable re-learning of MAC addresses on other SAPs within the VPLS. The MAC address will remain attached to a specified SAP for duration of its age-timer.

The age of the MAC address entry in the FDB is set by the age timer. If mac-aging is disabled on a specified VPLS service, any MAC address learned on a SAP/SDP with mac-pinning enabled will remain in the FDB on this SAP/SDP forever. Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).

MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

Default When a SAP or spoke-SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise, MAC pinning is not enabled by default.

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>endpoint
config>service>template>vpls-sap-template

Description This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP, spoke-SDP, or endpoint.

When the configured limit has been reached, and **discard-unknown-source** has been enabled for this SAP or spoke-SDP (see discard-unknown-source), packets with unknown source MAC addresses will be discarded.

The **no** form of the command restores the global MAC learning limitations for the SAP or spoke-SDP.

Default no max-nbr-mac-addr

Parameters *table-size* — Specifies the maximum number of learned and static entries allowed in the FDB of this service

Values 1 to 511999

mc-endpoint

Syntax **mc-endpoint** *mc-ep-id*
mc-endpoint

Context config>service>vpls>endpoint

Description This command specifies the identifier associated with the multi-chassis endpoint. This value should be the same on both MC-EP peers for the pseudowires that must be part of the same group.

The **no** form of this command removes the endpoint from the MC-EP. Single chassis behavior applies.

Default no mc-endpoint

Parameters *mc-ep-id* — Specifies a multi-chassis endpoint ID

Values 1 to 4294967295

mc-ep-peer

Syntax **mc-ep-peer** *name*
mc-ep-peer *ip-address*
no mc-ep-peer

Context config>service>vpls>endpoint>mc-ep

Description This command adds multi-chassis endpoint object.

The **no** form of this command removes the MC-Endpoint object.

Default mc-endpoint is not provisioned.

Parameters *name* — Specifies the name of the multi-chassis endpoint peer

ip-address — Specifies the IP address of multi-chassis endpoint peer

msap-defaults

Syntax **msap-defaults**

Context config>service>vpls>sap

Description This command configures the msap-defaults.

service

Syntax [**no**] **service** *service-id*

Context config>service>vpls>sap>msap-defaults

Description This command sets default service for all subscribers created based on trigger packets received on the specified capture SAP in case the corresponding VSA is not included in RADIUS authentication response. This command is applicable to capture SAP only.

Default no service

policy

Syntax **policy** *msap-policy-name*
no policy

Context config>service>vpls>sap>msap-defaults

Description This command sets default msap-policy for all subscribers created based on trigger packets received on the specified capture-sap in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only.

Default no policy

multi-service-site

Syntax **multi-service-site** *customer-site-name*

no multi-service-site

Context config>service>vpls>sap

Description This command associates the SAP with a *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command will not execute. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within *customer-site-name* as parent schedulers.

This command is mutually exclusive with the SAP ingress and egress **scheduler-policy** commands. If a **scheduler-policy** has been applied to either the ingress or egress nodes on the SAP, the **multi-service-site** command will fail without executing. The locally applied scheduler policies must be removed prior to executing the **multi-service-site** command.

The **no** form of the command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future queues to enter an orphaned state.

Default None

Parameters *customer-site-name* — Specifies the customer name site. The customer-site-name must exist in the context of the customer-id defined as the service owner. If customer-site-name exists and local scheduler policies have not been applied to the SAP, the current and future queues defined on the SAP will look for their parent schedulers within the scheduler hierarchies defined on customer-site-name.

Values Any valid customer-site-name created within the context of the customer-id

precedence

Syntax **precedence** [*precedence-value* | **primary**]
no precedence

Context config>service>vpls>spoke-sdp

Description This command configures the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint. When an SDP bind goes down, the next highest precedence SDP bind begins forwarding traffic.

Parameters *precedence-value* — Specifies the precedence of this SDP bind

Values 1 to 4

primary — Assigns this as the primary spoke-SDP

static-isid

Syntax [**no**] **static-isid range** *entry-id isid* [**to** *isid*] [**create**]

Context config>service>vpls>sap
config>service>vpls>vpls>spokesdp

Description This 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 an altered by an isid-policy. This allows communication from I-VPLS services external to SPBM through this node. The SAP may be a regular SAP or MC-LAG SAP. The spoke-SDP may be a active/standby spoke. When used with MC-Lag or active/stand-by PWs the conditional static-mac must be configured. ISIDs declared this way become part of the ISID multicast and consume MFIBs. Multiple SPBM static-isid ranges are allowed under a SAP/spoke-SDP.

The static-isids are associated with a remote BMAC that must be declared as a static-mac for unicast traffic. 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 group MACs are then associated with the active SAP or spoke-SDP. An ISID policy may be defined to suppress the advertisement of an ISID if the ISID is primary used for unicast services. The following rules govern the usage of multiple ISID statements:

- overlapping values are allowed:
  - isid from 301 to 310
  - isid from 305 to 315
  - isid 316
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “ISID from 301 to 316” statement.
- there is no consistency check with the content of ISID statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation t o execute the associated action for that entry.

The **no** form of the command removes all the previous statements under one interface

no isid value | from value to higher-value - removes a specific ISID value or range. Must match a previously used positive statement: for example if the command “isid 316 to 400” was used using “no isid 316 to 350” will not work but “no isid 316 to 400 will be successful.

Parameters *entry-id* — Sets context for specified entry ID for the static-isids

Values 1— 65535

isid — Configures the ISID or the start of an ISID range. Specifies the ISID value in 24 bits. When just one present identifies a particular ISID to be used for matching.

Values 0 to 16777215

to isid — Identifies upper value in a range of ISIDs to be used as matching criteria

Values 0 to 16777215

static-mac

Syntax [**no**] **static-mac** *ieee-mac-address* [*create*]

Context config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command creates a local static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Access Point (SAP).

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Local static MAC entries create a permanent MAC address to SAP association in the forwarding database for the VPLS instance so that MAC address will not be learned on the edge device.

Static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SAP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SAP from the VPLS forwarding database.

Parameters *ieee-mac-address* — Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

create — This keyword is mandatory when specifying a static MAC address

managed-vlan-list

Syntax **managed-vlan-list**

Context config>service>vpls>sap

Description This command enters the context to configure VLAN ranges to be managed by a management VPLS. The list indicates, for each SAP, the ranges of associated VLANs that will be affected when the SAP changes state. This managed-vlan-list is not used when STP mode is MSTP in which case the vlan-range is taken from the **config>service>vpls>stp>msti** configuration.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

default-sap

Syntax [**no**] **default-sap**

Context config>service>vpls>sap>managed-vlan-list

Description This command adds a default SAP to the managed VLAN list.

The **no** form of the command removes the default SAP to the managed VLAN list.

range

Syntax [**no**] **range** *vlan-range*

Context config>service>vpls>sap>managed-vlan-list

Description This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.

This command is only valid when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or qinq, or on a SONET/SDH port with encapsulation type of bcp-dot1q.

To modify the range of VLANs, first the new range should be entered and afterwards the old range removed. See Modifying VPLS Service Parameters.

Default None

Parameters *vlan-range* — Specify the VLAN start value and VLAN end value. The end-vlan must be greater than start-vlan. The format is <start-vlan>-<end-vlan>.

Values start-vlan: 0 to 4094
end-vlan: 0 to 4094

3.7.2.5.1 VPLS SAP ATM Commands

atm

Syntax **atm**

Context config>service>vpls>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a specified context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a specified context, the command returns an error.

egress

Syntax **egress**

Context config>service>vpls>sap>atm

Description This command enters the context to configure egress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>vpls>sap>atm

Description This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684, *Multiprotocol Encapsulation over ATM AAL5,* and to the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation will be dropped.

Default aal5snap-bridged

Parameters *atm-encap-type* — Specifies the encapsulation type

Values **aal5snap-bridged** — Bridged encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.

aal5mux-bridged-eth-nofcs — Bridged IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

ingress

Syntax **ingress**

Context config>service>vpls>sap>atm

Description This command enters the context to configure ingress ATM attributes for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>vpls>sap>atm>ingress
config>service>vpls>sap>atm>egress

Description This command assigns an ATM traffic descriptor profile to a specified context (for example, a SAP).

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.

When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId.=1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specifies a defined traffic descriptor profile (see the QoS atm-td-profile command)

oam

Syntax **oam**

Context config>service>vpls>sap>atm

Description This command enters the context to configure OAM functionality for a PVCC delimiting a SAP.

The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>vpls>sap>atm

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.

When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes operationally down, or enters a fault state and becomes operationally up, or exits that fault state). RDI cells are generated when PVCC is operationally down. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The **no** command disables **alarm-cells** functionality for a PVCC. When **alarm-cells** functionality is disabled, the PVCC's operational status is no longer affected by the PVCC's OAM state changes due to AIS/RDI processing. When **alarm-cells** is disabled, a PVCC will change operational status to operationally up from operationally down due to alarm-cell processing). RDI cells are not generated as result of PVCC going into an AIS or RDI state, however, the PVCC's OAM status will record OAM faults as described above.

Default Enabled for PVCCs delimiting VPLS SAPs

3.7.2.5.2 VPLS Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>vpls>sap

Description This command enters the context to configure egress filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If **no** egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>vpls>sap

Description This command enters the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

agg-rate

Syntax [**no**] **agg-rate**

Context config>service>vpls>sap>egress>
config>service>template>vpls-sap-template>egress
config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

rate

Syntax **rate** {**max** | **rate**}
no rate

Context config>service>vpls>sap>egress>agg-rate
config>service>template>vpls-sap-template>egress>agg-rate
config>service>vpls>sap>egress>encap-defined-qos>encap-group>agg-rate

Description This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered active on the context's object (SAP, subscriber, Vport, and so on.).

limit-unused-bandwidth

Syntax [**no**] **limit-unused-bandwidth**

Context config>service>vpls>sap>egress>agg-rate
config>service>template>vpls-sap-template>egress>agg-rate
config>service>vpls>sap>egress>encap-defined-qos>encap-group>agg-rate

Description This command is used to enable aggregate rate overrun protection on the agg-rate context.

The **no** form of the command disables the overrun protection.

queue-frame-based-accounting

Syntax [**no**] **queue-frame-based-accounting**

Context config>service>vpls>sap>egress>agg-rate

config>service>template>vpls-sap-template>egress>agg-rate

Description This command is used to enabled frame-based accounting on all policers and queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured; however the offsets are applied to the statistics.

The **no** form of the command disables the-frame based accounting.

dest-mac-rewrite

Syntax **dest-mac-rewrite** *ieee-address*
no dest-mac-rewrite

Context config>service>vpls>sap>egress>agg-rate

Description This commands enables the overwriting of a destination MAC address to an operator-configured value for all unicast packets egressing the specified SAP. The command is intended to be deployed with L2 PBF SAP redirect when a remote end of the SAP interface is an L3 interface with a MAC address different from the MAC address of the non-PBF-ed L3 interface. See Filter Policy in the *7450 ESS, 7750 SR, and 7950 XRS Router Configuration Guide* for more details.

The **no** form disables the option.

Default no dest-mac-rewrite

Parameters *ieee-address* — Specifies the MAC address

Values 1xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
Cannot be all zeros

encap-defined-qos

Syntax **encap-defined-qos**

Context config>service>vpls>sap>egress

Description This command creates a new QoS sub-context in B-VPLS SAP egress context. The user can define encapsulation groups, referred to as encap-group, based on the ISID value in the packet's encapsulation and assign a QoS policy and a scheduler policy or aggregate rate limit to the group.

encap-group

Syntax **encap-group** *group-name* [**type** *group-type*] [**qos-per-member**] [**create**]

no encap-group *group-name*

Context config>service>vpls>sap>egress>encap-defined-qos

Description This command defines an encapsulation group which consists of a group of ISID values. All packets forwarded on the egress of a B-VPLS SAP which payload header matches one of the ISID value in the encap-group will use the same QoS policy instance and scheduler policy or aggregate rate limit instance.

The user adds or removes members to the encap-group one at a time or as a range of contiguous values using the member command. However, when the qos-per-member option is enabled, members must be added or removed one at a time. These members are also referred to as ISID contexts.

The user can configure one or more encap-groups in the egress context of the same B-SAP, thus defining different ISID values and applying each a different SAP egress QoS policy, and optionally a different scheduler policy/**agg-rate**. ISID values are unique within the context of a B-SAP. The same ISID value cannot be re-used in another encap-group under the same B-SAP but can be re-used in an encap-group under a different B-SAP. Finally, if the user adds to an encap-group an ISID value which is already a member of this encap-group, the command causes no effect. The same if the user attempts to remove an ISID value which is not a member of this encap-group.

Once a group is created, the user will assign a SAP egress QoS policy, and optionally a scheduler policy or aggregate rate limit, using the following commands:

config>service> vpls>sap>egress>encap-defined-qos>encap-group>qos *sap-egress-policy-id*

config>service> vpls>sap>egress>encap-defined-qos>encap-group>scheduler-policy *scheduler-policy-name*

config>service> vpls>sap>egress>encap-defined-qos>encap-group>agg-rate *kilobits-per-second*

A SAP egress QoS policy must first be assigned to the created encap group before the user can add members to this group. Conversely, the user cannot perform the **no qos** command until all members are deleted from the encap-group.

An explicit or the default SAP egress QoS policy will continue to be applied to the entire B-SAP but this will serve to create the set of egress queues which will be used to store and forward a packet which does not match any of the defined ISID values in any of the encap-groups for this SAP.

Only the queue definition and fc-to-queue mapping from the encap-group SAP egress QoS policy is applied to the ISID members. All other parameters configurable in a SAP egress QoS policy must be inherited from egress QoS policy applied to the B-SAP.

Furthermore, any other CLI option configured in the egress context of the B-SAP will continue to apply to packets matching a member of any encap-group defined in this B-SAP.

The keyword **qos-per-member** allows the user to specify that a separate queue set instance and scheduler/**agg-rate** instance will be created for each ISID value in the encap-group. By default, shared instances will be created for the entire encap-group.

When the B-SAP is configured on a LAG port, the ISID queue instances defined by all the encap-groups applied to the egress context of the SAP will be replicated on each member link of the LAG. The set of scheduler/**agg-rate** instances will be replicated per link or per IOM or XMA depending if the adapt-qos option is set to link/port-fair mode or distribute mode. This is the same behavior as that applied to the entire B-SAP in the current implementation.

The **no** form of this command deletes the encap-group.

Parameters *group-name* — Specifies the name of the encap-group and can be up to 32 ASCII characters in length

type — Specifies the type of the encapsulation ID used by this encap-group

Values isid

Default None

qos-per-member — Specifies that a separate queue set instance and scheduler/**agg-rate** instance will be created for each ISID value in the encap-group

member

Syntax [**no**] **member** *encap-id* [**to** *encap-id*]

Context config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command adds or removes a member ISID or a range of contiguous ISID members to an encap-group. The user can add or remove members to the encap-group one at a time or as a range of contiguous values using the member command. However, when the **qos-per-member** option is enabled, members must be added or removed one at a time.

The **no** form of this command removes the single or range of ISID values from the encap-group.

Parameters *encap-id* — The value of the single encap-id or the start encap-id of the range. ISID is the only encap-id supported.

to *encap-id* — The value of the end encap-id of the range. ISID is the only encap-id supported.

qos

Syntax **qos** *policy-id*
no qos

Context config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command configures the QoS ID.

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>service>vpls>sap>egress>encap-defined-qos>encap-group

Description This command configures the scheduler policy.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*
no filter [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>vpls>mesh-sdp>egress
config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>egress
config>service>vpls>spoke-sdp>ingress

Description This command associates an IP filter policy or MAC filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *filter ID* with an ingress or egress SAP. The *filter ID* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Special Cases **VPLS —** Both MAC and IP filters are supported on a VPLS service SAP.

Parameters **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6 *ipv6-filter-id* — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac *mac-filter-id* — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

hsmda-queue-override

Syntax [**no**] **hsmda-queue-override**

Context config>service>vpls>sap>egress

Description This command enters the context to configure HSMDA queue overrides.

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>service>vpls>sap>egress>hsmda-queue-override

Description This command configures overrides for a HSMDA queue. The actual valid values are those defined in the specified SAP QoS policy.

Parameters *queue-id* — Specifies the queue ID to override

Values 1 to 8

create — This keyword is mandatory while creating a new queue override

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
no packet-byte-offset

Context config>service>vpls>sap>egress>hsmda-queue-over

Description This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4 byte CRC (everything except the preamble and inter-frame gap). As an example, the packet-byte-offset command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 31 bytes may be added to the packet and up to 32 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. The packet-byte-offset, when set, applies to all queues in the queue group. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscribers packets on an Ethernet aggregation network.

The packet-byte-offset value may be overridden at the queue-group level.

Parameters **add** *add-bytes* — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The **add** keyword is mutually exclusive with the **subtract** keyword.

Values 0 to 31

subtract *sub-bytes* — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The **subtract** keyword is mutually exclusive with the **add** keyword. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 1 to 64

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context config>service>vpls>sap>egress>hsmda-queue-over>queue

Description This command specifies an existing slope policy name.

rate

Syntax **rate** *pir-rate*
no rate

Context config>service>vpls>sap>egress>hsmda-queue-over>queue

Description This command specifies the administrative PIR by the user.

Parameters *pir-rate* — Configures the administrative PIR specified by the user.

Values 1 to 40000000

wrr-weight

Syntax **wrr-weight value**
no wrr-weight

Context config>service>vpls>sap>egress>hsmda-queue-over>queue

Description This command assigns the weight value to the HSMDA queue.

The **no** form of the command returns the weight value for the queue to the default value.

Parameters *percentage* — Specifies the weight for the HSMDA queue

Values 1 to 32

wrr-policy

Syntax **wrr-policy** *hsmda-wrr-policy-name*
no wrr-policy

Context config>service>vpls>sap>egress>hsmda-queue-over

Description This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

Parameters *hsmda-wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue

secondary-shaper

Syntax **secondary-shaper** *secondary-shaper-name*
no secondary-shaper

Context config>service>vpls>sap>egress>hsmda-queue-over

Description This command configures an HSMDA secondary shaper. A shaper override can only be configured on an HSMDA SAP.

Parameters *secondary-shaper-name* — Specifies a secondary shaper name up to 32 characters in length

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>vpls>sap>egress

Description When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the **qinq-mark-top-only** command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When enabled, only the P-bits/DEI bit in the top Q-tag are marked.

The **no** form of this command disables the command.

Default no qinq-mark-top-only

policer-control-override

Syntax **policer-control-override** [**create**]
no policer-control-override

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command, within the SAP ingress or egress contexts, creates a CLI node for specific overrides to the applied policer-control-policy. A policy must be applied for a policer-control-overrides node to be created. If the policer-control-policy is removed or changed, the policer-control-overrides node is automatically deleted from the SAP.

The **no** form of the command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.

Default no policer-control-override

Parameters **create** — The create keyword is required when the policer-control-overrides node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

max-rate

Syntax **max-rate** {*rate* | **max**}

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command, within the SAP ingress and egress contexts, overrides the root arbiter parent policer max-rate that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy max-rate parameter have no effect on the SAP's parent policer until the override is removed using the **no max-rate** command within the SAP.

Parameters *rate* **| max** — Specifies the max rate override in kilobits per second or use the maximum

Values 1 to 20000000 kb/s

priority-mbs-thresholds

Syntax **priority-mbs-thresholds**

Context config>service>vpls>sap>egress

Description This command overrides the CLI node contains the configured min-thresh-separation and the various priority level mbs-contribution override commands.

min-thresh-separation

Syntax **min-thresh-separation** *size* [**bytes** | **kilobytes**]

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command within the SAP ingress and egress contexts is used to override the root arbiter's parent policer min-thresh-separation parameter that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy min-thresh-separation parameter have no effect on the SAP's parent policer until the override is removed using the no min-thresh-separation command within the SAP.

The no form of the command removes the override and allows the min-thresh-separation setting from the policer-control-policy to control the root arbiter's parent policer's minimum discard threshold separation size.

Default no min-thresh-separation

Parameters **bytes** — Signifies that size is expressed in bytes. The bytes and kilobytes keywords are mutually exclusive and are optionally used to qualify whether size is expressed in bytes or kilobytes. The default is kilobytes.

kilobytes — The size parameter is required when specifying the min-thresh-separation override. It is specified as an integer representing either a number of bytes or kilobytes that are the minimum separation between the parent policer's priority level discard thresholds.

Values 0 to 16777216

Default kilobytes

priority

Syntax [**no**] **priority** *level*

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.

This node does not need to be created and will not be output in show or save configurations unless an mbs-contribution override exist for level.

Parameters *level* — Specifies that the level parameter is required when specifying priority-level and identifies which of the parent policer instances priority level's the mbs-contribution is overriding

Values 1 to 8

mbs-contribution

Syntax **mbs-contribution** *size* [**bytes** | **kilobytes**]

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description The mbs-contribution override command within the SAP ingress and egress contexts is used to override a parent policer's priority level's mbs-contribution parameter that is defined within the policer-control-policy applied to the SAP. This override allow the priority level's burst tolerance to be tuned based on the needs of the SAP's child policers attached to the priority level.

When the override is defined, modifications to the policer-control-policy priority level's mbs-contribution parameter have no effect on the SAP's parent policer priority level until the override is removed using the no mbs-contribution command within the SAP.

The **no** form of the command removes the override and allows the mbs-contribution setting from the policer-control-policy to control the parent policer's priority level's burst tolerance.

Default no mbs-contribution

Parameters **bytes** — Signifies that size is expressed in bytes

kilobytes — Signifies that the size is expressed in kilobytes

Values 0 – 16777216 or default

policer-control-policy

Syntax **policer-control-policy** *policy-name*
no policer-control-policy

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>template>vpls-sap-template>egress
config>service>template>vpls-sap-template>ingress

Description This command, within the qos CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy may also be applied to the ingress or egress context of a sub-profile.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and therefore the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair-based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

To derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of the command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP or subscriber management sub-profile context.

Default none

Parameters *policy-name* — Specifies the policy name. Each policer-control-policy must be created with a unique policy name. The name must given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create — The keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

policer-override

Syntax [**no**] **policer-override**

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of the command is used to remove any existing policer overrides.

Default no policer-overrides

policer

Syntax **policer** *policer-id* [**create**]
no policer *policer-id*

Context config>service>vpls>sap>egress>policer-override
config>service>vpls>sap>ingress>policer-override

Description This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy.

The **no** form of the command is used to remove any existing overrides for the specified policer-id.

Parameters *policer-id* — The policer-id parameter is required when executing the policer command within the policer-overrides context. The specified policer-id must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the policer-id.

create — The create keyword is required when a policer policer-id override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

cbs

Syntax **cbs** *size* [**bytes** | *kilobytes*]
no cbs

Context config>service>vpls>sap>egress>policer-override
config>service>vpls>sap>ingress>policer-override

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured CBS parameter for the specified policer-id.

The **no** form of this command returns the CBS size to the default value.

Default no cbs

Parameters *size-in-kbytes* — This parameter is required when specifying mbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 or default

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>service>vpls>sap>egress>policer-override>policer
config>service>vpls>sap>ingress>policer-override>policer

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured mbs parameter for the specified policer-id.

The **no** form of the command restores the policer's mbs setting to the policy defined value.

Default no mbs

Parameters **size** — The size parameter is required when specifying an mbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 | default

bytes — When **bytes** is defined, the value given for size is interpreted as the policer's MBS value in bytes.

kilobytes — When **kilobytes** is defined, the value given for size is interpreted as the policer's MBS value in kilobytes.

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}

Context config>service>vpls>sap>egress>policer-override>policer
config>service>vpls>sap>ingress>policer-override>policer

Description This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured packet-byte-offset parameter for the specified policer-id. Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured, however the offsets are applied to the statistics.

The **no** form of the command restores the policer's packet-byte-offset setting to the policy defined value.

Default no packet-byte-offset

Parameters **add** *add-bytes* — The add keyword is mutually exclusive to the subtract keyword. Either add or subtract must be specified. When add is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 1 to 31

subtract *sub-bytes* — The subtract keyword is mutually exclusive to the add keyword. Either add or subtract must be specified. When subtract is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 to 64

rate

Syntax **rate** {*rate* | **max**} [**cir** {**max** | *rate*}]

Context config>service>vpls>sap>egress>policer-override>policer
config>service>vpls>sap>ingress>policer-override>policer

Description This command within the SAP ingress and egress policer-overrides contexts is used to override the sap-ingress and sap-egress QoS policy configured rate parameters for the specified policer-id.

The **no** form of the command restores the policy defined metering and profiling rate to a policer.

Parameters {*rate* | **max**} — Specifying the keyword **max** or an explicit kilobits per second parameter directly following the rate override command is required and identifies the policer instance's metering rate for the PIR leaky bucket. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to **max**.

Values 1 to 2000000000, **max**

cir {**max** | *rate*} — The optional **cir** keyword is used to override the policy derived profiling rate of the policer. Specifying the keyword **max** or an explicit kilobits per second parameter directly following the **cir** keyword is required. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to **max**.

Values 0 to 2000000000, **max**

stat-mode

Syntax **stat-mode** st*at-mode*
no stat-mode

Context config>service>vpls>sap>egress>policer-override>policer
config>service>vpls>sap>ingress>policer-override>policer

Description The SAP-egress QoS policy's policer stat-mode command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. A policer has multiple types of offered packets (for example, soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overrides) and each of these offered types is interacting with the policers metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The stat-mode command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires that the policer's stat-mode to be set at least to the minimal setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated.

Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources.You can view the total/allocated/free stats by using the **tools dump system-resources** command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer.

The default stat-mode when a policer is created within the policy is minimal.

The stat-mode setting defined for the policer in the QoS policy may be overridden on a SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The previous stat-mode setting active for the policer will continue to be used by the policer.

The **no** form of the command returns the policer's stat-mode setting to minimal.

Refer to the *7450 ESS, 7750 SR, and 7950 XRS Quality of Service Guide* for detailed information about the **policer stat-mode** command parameters.

qos

Syntax **qos** *policy-id* [*shared-queuing* | *multipoint-shared*] [*fp-redirect-group queue-group-name instance instance-id*]
no qos

Context config>service>vpls>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the policy-id does not exist, an error will be returned.

The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

When an ingress QoS policy is defined on IES ingress IP interface that is bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPLS SAP; packets that are routed will be processed using the policy defined in the IES IP interface-binding context.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, the default QoS policy is used.

The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Default none

Parameters *policy-id* — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.

Values 1 to 65535

shared-queuing — This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared — This keyword can only be specified on SAP ingress. Multipoint shared queuing is a superset of shared queuing. When multipoint shared queuing keyword is set, as well as the unicast packets, multipoint packets also used shared queues.

Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice, similar to the shared-queuing option. In addition, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

Values Multipoint or not present

Default Present (the queue is created as non-multipoint)

fp-redirect-group — Creates an instance of a named queue group template on the ingress forwarding plane of a specified IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command will fail.

queue-group-name — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid ingress queue group template name, configured under config>qos>queue- group-templates.

instance-id — Specifies the instance of the named queue group to be created on the IOM/IMMXMA ingress forwarding plane.

qos

Syntax **qos policy-id** [**port-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>vpls>sap>egress

Description This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy-id does not exist, an error will be returned.

The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

When an egress QoS policy is associated with an IES IP interface that has been bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the egress QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPLS SAP; packets that are routed will be processed using the policy defined in the IES IP interface- binding context.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Default none

Parameters *port-redirect-group* — Associates a SAP egress with an instance of a named queue group template on the egress port of a specified IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command.

queue-group-name — Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under *config>port>ethernet>access>egress*.

instance *instance-id* — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA

Values 1 to 40960

Default 1

queue-override

Syntax [**no**] **queue-override**

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

config>service>vpls>sap>egress>hsmda-queue-over>queue
config>service>vpls>sap>ingress>hsmda-queue-over>queue

Description This command enters the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy.

queue

Syntax [**no**] **queue** *queue-id*

Context config>service>vpls>sap>egress>queue-override
config>service>vpls>sap>ingress>queue-override

Description This command specifies the ID of the queue whose parameters are to be overridden.

The **no** form of the command removes the queue-id from the configuration.

Parameters *queue-id* — The queue ID whose parameters are to be overridden

Values 1 to 32

adaptation-rule

Syntax **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
no adaptation-rule

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default no adaptation-rule

Parameters *pir* — The **pir** parameter defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>service>vpls>sap>egress>queue-override>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a SONET or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 0.1 or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50 x 20 or 1000 octets.

- Frame based offered-load — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- Packet to frame factor — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be 1000 / 10000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- Frame based CIR — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to calculate the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0 to 100

burst-limit

Syntax **burst-limit** {**default** | *size* [**bytes** | **kilobytes**]}
no burst-limit

Context config>service>vpls>sap>egress>queue-override>queue

Description The queue **burst-limit** command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override queue context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy.

Default no burst-limit

Parameters **default** — Reverts the queue's burst limit to the system default value.

size — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** qualifier must be added following size.

Values 1 to 13671 kilobytes
1 to 14000000 bytes

Default No default for *size*; use the **default** keyword to specify default burst limit.

bytes — Specifies that the value given for size must be interpreted as the burst limit in bytes.

kilobytes — Specifies that the value given for size must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is **kilobytes**.

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's CBS parameters.

It is permissible to oversubscribe the total CBS reserved buffers for a specified access port egress buffer pool. Over-subscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

The **no** form of this command returns the CBS to the default value.

Default no cbs

Parameters *size-in-kbytes* — Specifies the number of kilobytes reserved for the queue. For a value of 10 kbytes, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimum reserved size can be applied for scheduling purposes).

Values 0 to 131072 or default

drop-tail

Syntax **drop-tail**

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command enters the context to configure queue drop tail parameters.

low

Syntax **low**

Context config>service>vpls>sap>egress>queue-override>queue>drop-tail
config>service>vpls>sap>ingress>queue-override>queue>drop-tail

Description This command enters the context to configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and will be discarded.

percent-reduction-from-mbs

Syntax **percent-reduction-from-mbs** *percent*
no percent-reduction-from-mbs

Context config>service>vpls>sap>egress>queue-override>queue>drop-tail>low
config>service>vpls>sap>ingress>queue-override>queue>drop-tail>low

Description This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets are not accepted into the queue if the queue depth is greater than the low drop tail value, and so will be discarded.

Parameters *percent* — Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting correct CBS parameters and controlling CBS over-subscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS assigned to the queue to the default setting.

Default default

Parameters *size* — The size parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether the size represents bytes or kilobytes.

Values 0 to 1073741824
default

bytes — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

default — Specifying the keyword **default** sets the MBS to its default value.

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
no mbs

Context config>service>vpls>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting correct CBS parameters and controlling CBS over-subscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS assigned to the queue to the default value.

Default default

Parameters *size* — The size parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether the size represents bytes or kilobytes.

Values 0 to 1073741824
default

bytes — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

default — Specifying the keyword **default** sets the MBS to its default value.

parent

Syntax **parent** {[**weight** *weight*] [**cir-weight** *cir-weight*]}
no parent

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier** *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state mentioned above and automatically return to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters *weight* — These optional keywords are mutually exclusive to the **level** keyword. The weight defines the relative weight of this queue in comparison to other child schedulers, policers, and queues while vying for bandwidth on the parent *scheduler-name*. Any policers, queues, or schedulers defined as weighted receive no parental bandwidth until all policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active policers, queues, and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the policer, queue, or scheduler. A weight is considered to be active when the pertaining policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit. All child policers, queues, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted policers, queues, and schedulers at that level are operating at the maximum bandwidth or are idle.

Values 0 to 100

Default 1

cir-weight — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the policer, queue, or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description The **percent-rate** command supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gb and a 10-Gb Ethernet port, the queue's rates will be 10 times greater on the 10-Gb port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

Parameters *pir-percent* — Specifies the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.01 to 100.00

Default 100.00

cir-percent — Specifies the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.00 to 100.00

Default 100.00

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>hsmda-queue-over>queue

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by over-subscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile and then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default rate max cir 0

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

For **egress>queue-override>queue** and **ingress>queue-override>queue**:

Values 1 to 2000000000, **max** in kb/s

Default **max**

For **egress>hsmda-queue-over>queue:**

Values 1 to 100000000, **max** in kb/s

Default **max**

cir cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues.

For **egress>queue-override>queue** and **ingress>queue-override>queue**:

Values 0 to 2000000000, **max** in kb/s

Default 0

queue-override

Syntax [**no**] **queue-override**

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>vpls>sap>egress>hsmda-queue-over>queue
config>service>vpls>sap>ingress>hsmda-queue-over>queue

Description This command enters the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy.

queue

Syntax [**no**] **queue** *queue-id*

Context config>service>vpls>sap>egress>queue-override
config>service>vpls>sap>ingress>queue-override

Description This command specifies the ID of the queue whose parameters are to be overridden.

Parameters *queue-id* — Specifies the queue ID whose parameters are to be overridden

Values 1 to 32

adaptation-rule

Syntax **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
no adaptation-rule

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default no adaptation-rule

Parameters *pir* — Defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — The Defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*

no avg-frame-overhead

Context config>service>vpls>sap>egress>queue-override>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a SONET or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 0.1 or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50 x 20 or 1000 octets.

- Frame based offered-load — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- Packet to frame factor — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be 1000 / 10000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- Frame based CIR — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500 x 1.1 or 550 octets.

- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to calculate the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — Sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0 to 100

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a specified access port egress buffer pool. Over-subscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

The **no** form of this command returns the CBS size to the default value.

Default no cbs

Parameters *size-in-kbytes* — Specifies an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is needed, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 131072 or default

mbs

Syntax **mbs** {*size* [**byte**s | **kilobytes**] | **default**}
no mbs

Context config>service>vpls>sap>egress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting correct CBS parameters and controlling CBS over-subscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size assigned to the queue.

Default default

Parameters *size* — The size parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 1073741824
default

bytes — When **byte**s is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

default — Specifying the keyword **default** sets the MBS to its default value.

mbs

Syntax **mbs** {*size* [**byte**s | **kilobytes**] | **default**}
no mbs

Context config>service>vpls>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting correct CBS parameters and controlling CBS over-subscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size assigned to the queue to the default value.

Default default

Parameters *size* — The size parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 1073741824
default

bytes — When **bytes** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

default — Specifying the keyword **default** sets the MBS to its default value.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>hsmda-queue-over>queue

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by over-subscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile and then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default rate max cir 0

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 100000000

Default max

cir cir-rate — Overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues.

Values 0 to 100000000, **max**, **sum**

Default 0

scheduler-override

Syntax [**no**] **scheduler-override**

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command specifies the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's scheduler policy.

scheduler

Syntax **scheduler** *scheduler-name*
no scheduler *scheduler-name*

Context config>service>vpls>sap>egress>sched-override

Description This command can be used to override specific attributes of the specified scheduler name. A scheduler defines a bandwidth controls that limit each child (other schedulers, policers, and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause policers, queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

Step 1. The maximum number of schedulers has not been configured.

Step 2. The provided *scheduler-name* is valid.

Step 3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

Parameters *scheduler-name* — Specifies the name of the scheduler.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

Default None. Each scheduler must be explicitly created.

create — Specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

parent

Syntax **parent** [**weight** *weight*] [**cir-weight** *cir-weight*]
no parent

Context config>service>vpls>sap>ingress>sched-override>scheduler
config>service>vpls>sap>egress>sched-override>scheduler

Description This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers.

The **no** form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default no parent

Parameters *weight* — Defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue, or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.
A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

cir-weight — Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is

considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.
A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>vpls>sap>egress>sched-override>scheduler
config>service>vpls>sap>ingress>sched-override>scheduler

Description This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child policers, queues, or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other then its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child policers, queues, and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters *pir-rate* — Specifies a value of 1 to 3200000000, or the keyword **max**. Any other value will result in an error without modifying the current PIR rate.

Values 1 to 3200000000, **max**

cir cir-rate — Specifies a value of 0 to 3200000000, or the keyword **max** or **sum**. Any other value will result in an error without modifying the current CIR rate.

If **cir** is set to max, then the CIR rate is set to infinity, but bounded by the PIR rate.

The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers or queues.

Values 0 to 3200000000, **max, sum**

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>service>vpls>sap>ingress
config>service>vpls>sap>egress

Description This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues and, at egress only, policers associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created after the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the SAP policers and queues associated with the customer site. Policers and queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have policers or queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers or queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

Parameters *scheduler-policy-name:* — Applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.

Values Any existing valid scheduler policy name.

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p de

Context config>service>vpls>sap>ingress

Description This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The setting also applies to classification based on the DE indicator bit.

The **no** form of this command reverts the dot1p and de bits matching to the default tag.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 40 defines the default behavior for Dot1P evaluation.

Table 40 **Default QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Default no match-qinq-dot1p (no filtering based on p-bits)
(top or bottom must be specified to override the default QinQ dot1p behavior)

Parameters **top** — The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. The following table defines the dot1p evaluation behavior when the top parameter is specified.

Table 41

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits

Table 41 **(Continued)**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	TopQ PBits

bottom — The bottom parameter is mutually exclusive to the top parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. The following table defines the dot1p evaluation behavior when the bottom parameter is specified; see Table 42.

Table 42 **Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Table 43

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
Null	No preserved Dot1P bits	None
Null	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
Dot1Q	No preserved Dot1P bits	New PBits marked using dot1p-value
Dot1Q	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
TopQ	No preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	No preserved Dot1P bits	TopQ PBits and BottomQ PBits marked using dot1p-value
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits and BottomQ PBits marked using dot1p-value

The QinQ and TopQ SAP PBit/DEI bit marking follows the default behavior defined in the table above when **qinq-mark-top-only** is not specified.

The dot1p *dot1p-value* command must be configured without the qinq-mark-top-only parameter to remove the TopQ PBits only marking restriction.

A QinQ-encapsulated Ethernet port can have two different sap types:

For a TopQ SAP type, only the outer (top) tag is explicitly specified. For example, **sap 1/1/1:10.***

For QinQ SAP type, both inner (bottom) and outer (top) tags are explicitly specified. For example, **sap 1/1/1:10.100**.

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>vpls>sap

Description This command defines which subscriber authentication policy must be applied when a DHCP message is received on the interface. The authentication policies must already be defined. The policy will only be applied when DHCP snooping is enabled on the SAP.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>sap

Description This command creates the accounting policy context that can be applied to a SAP or SDP. An accounting policy must be defined before it can be associated with a SAP or SDP. If the *policy-id* does not exist, an error message is generated.
A maximum of one accounting policy can be associated with a SAP or SDP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SAP or SDP, and the accounting policy reverts to the default.

Default Default accounting policy.

Parameters *acct-policy-id* — Specifies the accounting *policy-id* as configured in the **config>log>accounting-policy** context

Values 1 to 99

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>vpls>spoke-sdp

Description This command configures the application profile name.

Parameters *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>sap

Description This command enables accounting and statistical data collection for either the SAP or SDP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM or XCM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default no collect-stats

3.7.2.5.3 VPLS Template Commands

template

Syntax **template**

Context config>service

Description This is the node for service templates.

vpls-template

Syntax **vpls-template** *name/id* **create**
[**no**] **vpls-template** *name/id*

Context config>service>template

Description This command is used to create a vpls-template to be used to auto-instantiate a range of VPLS services. Only certain existing VPLS attributes specified in the command reference section can be changed in the vpls-template, not in the instantiated VPLS. The following attributes will be automatically set in the instantiated VPLSs (no template configuration necessary) and the operator cannot change these values.

vpn-id: none

description: “Service <svc id> auto-generated by control VPLS <svc-id>”

service-name: “Service <svc id>” (Auto-generated)

shutdown: no shutdown

Following existing attributes can be set by the user in the instantiated VPLSs:

[no] sap

All the other VPLS attributes are not supported.

Parameters *name/id* — Specifies the name in ASCII or the template ID

Values name: ASCII string

Values ID: [1 to 2147483647]

vpls-sap-template

Syntax **vpls-sap-template** *name/id* **create**
[**no**] **vpls-sap-template** *name/id*

Context config>service>template

Description This is the command used to create a SAP template to be used in a vpls-template. Only certain existing VPLS SAP attributes can be changed in the vpls-sap-template, not in the instantiated VPLS SAP

Following SAP attributes will be set in the instantiated saps (no configuration allowed):

description:"Sap <sap-id> controlled by MVRP service <svc id>" – auto generated

shutdown: no shutdown

Parameters *name/id* — Specifies the name in ASCII or the template ID

Values 1 to 2147483647

mac-move-level

Syntax **mac-move-level** {**primary** | **secondary**}
no mac-move-level

Context config>service>template>vpls-sap-template

Description When a sap is instantiated using vpls-sap-template, if the MAC move feature is enabled at VPLS level, the command mac-move-level indicates whether the sap should be populated as primary-port, secondary-port or tertiary-port in the instantiated VPLS.

Default no mac-move-level; SAP is populated as a tertiary-port

temp-flooding

Syntax **temp-flooding flood-time**
no temp-flooding

Context config>service>vpls
config>service>template>vpls-template

Description The temporary flooding is designed to minimize failover times by eliminating the time it takes to flush the MAC tables and if MVRP is enabled the time it takes for MVRP registration. Temporary flooding is initiated only upon xSTP TCN reception. During this procedure while the MAC flush takes place the frames received on one of the VPLS SAPs/pseudowires are flooded in a VPLS context which for MVRP case includes also the unregistered MVRP trunk ports. The MAC Flush action is initiated by the STP TCN reception or if MVRP is enabled for the data VPLS, by the reception of a MVRP New message for the SVLAN ID associated with the data VPLS. As soon as the MAC Flush is done, regardless of whether the temp-flooding timer expired or not, traffic will be delivered according to the regular FDB content which may be built from MAC Learning or based on MVRP registrations. This command provides a flood-time value that configures a fixed amount of time, in seconds, during which all traffic is flooded (BUM or known unicast) as a safety mechanism. Once the flood-time expires, traffic will be delivered according to the regular FDB content which may be built from MAC Learning or based on MVRP registrations. The temporary flooding timer should be configured in such a way to allow auxiliary processes like MAC Flush, MMRP and/or MVRP to complete/converge. The temporary flooding behavior applies to regular VPLS, VPLS instantiated with VPLS-template, IVPLS and BVPLS when MMRP is disabled.

The **no** form of the command disables the temporary flooding behavior.

Default no temp-flooding

Parameters *flood-time* — Specifies the flood time, in seconds

Values 3 to 600

3.7.2.5.4 Provider Tunnel Commands

provider-tunnel

Syntax **provider-tunnel**

Context config>service>vpls

Description This command creates the context to configure the use of a P2MP LSP for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLS instance. The P2MP LSP is referred to as the Provider Multicast Service Interface (PMSI).

inclusive

Syntax **inclusive**

Context config>service>vpls>provider-tunnel

Description his command creates the context to configure the use of a P2MP LSP as the default tree for forwarding Broadcast, Unicast unknown, and Multicast (BUM) packets of a VPLS or B-VPLSs instance. The P2MP LSP is referred to, in this case, as the Inclusive Provider Multicast Service Interface (I-PMSI).

When enabled, this feature relies on BGP Auto-Discovery (BGP-AD) or BGP-VPLS to discover the PE nodes participating in a specified VPLS/B-VPLS instance. The AD route contains the information required to signal both the point-to-point (P2P) PWs used for forwarding unicast known Ethernet frames and the RSVP or mLDP P2MP LSP used to forward the BUM frames.

The root node signals the RSVP P2MP LSP based on an LSP template associated with the I-PMSI at configuration time. The leaf node will join automatically the P2MP LSP, which matches the I-PMSI tunnel information discovered via BGP.

With a mLDP I-PMSI, each leaf node will initiate the signaling of the mLDP P2MP LSP upstream using the P2MP FEC information in the I-PMSI tunnel information discovered via BGP-AD.

If IGMP or PIM snooping are configured on the VPLS instance, multicast packets matching a L2 multicast Forwarding Information Base (FIB) record will also be forwarded over the P2MP LSP.

The user enables the use of an RSVP P2MP LSP as the I-PMSI for forwarding Ethernet BUM and IP multicast packets in a VPLS/B-VPLS instance using the following commands:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>rsvp>lsp-template *p2mp-lsp-template-name*

The user enables the use of an LDP P2MP LSP as the I-PMSI for forwarding Ethernet BUM and IP multicast packets in a VPLS instance using the following command:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>mldp

After the user performs a **no shutdown** under the context of the inclusive node and the expiration of a delay timer, BUM packets will be forwarded over an automatically signaled mLDP P2MP LSP or over an automatically signaled instance of the RSVP P2MP LSP specified in the LSP template.

The user can specify if the node is both root and leaf in the VPLS instance:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>root-and-leaf

The **root-and-leaf** command is required otherwise this node will behave as a leaf only node by default. When the node is leaf only for the I-PMSI of type P2MP RSVP LSP, no PMSI Tunnel Attribute is included in BGP-AD route update messages and therefore no RSVP P2MP LSP is signaled but the node can join RSVP P2MP LSP rooted at other PE nodes participating in this VPLS/B-VPLS service. The user must still configure a LSP template even if the node is a leaf only. For the I-PMSI of type mLDP, the leaf-only node will join I-PMSI rooted at other nodes it discovered but will not include a PMSI Tunnel Attribute in BGP-AD route update messages. This way a leaf only node will forward packets to other nodes in the VPLS/B-VPLS using the point-to-point spoke-SDPs.

BGP-AD must have been enabled in this VPLS/B-VPLS instance or the execution of the **no shutdown** command under the context of the inclusive node is failed and the I-PMSI will not come up.

Any change to the parameters of the I-PMSI, such as disabling the P2MP LSP type or changing the LSP template requires that the inclusive node be first shutdown. The LSP template is configured in MPLS.

If the P2MP LSP instance goes down, VPLS/B-VPLS immediately reverts the forwarding of BUM packets to the P2P PWs. The user can however restore at any time the forwarding of BUM packets over the P2P PWs by performing a **shutdown** under the context of the inclusive node.

This feature is supported with VPLS, H-VPLS, and B-VPLS. It is not supported with I-VPLS and Routed VPLS.

data-delay-interval

Syntax **data-delay-interval** *seconds*
no data-delay-interval

Context config>service>vpls>provider-tunnel>inclusive

Description This command configures the I-PMSI data delay timer.

This delay timer is intended to allow time for the RSVP control plane to signal and bring up the S2L sub-LSP to each destination PE participating in the VPLS/B-VPLS service. The delay timer is started as soon as the P2MP LSP instance becomes operationally up after the user performed a ‘no shutdown’ under the inclusive node, i.e., as soon as the first S2L sub-LSP is up. In general, it is started when the P2MP LSP instance transitions from the operationally down state to the up state.

For a mLDP P2MP LSP, the delay timer is started as soon as the P2MP FEC corresponding to the I-PMSI is resolved and installed at the root node. The user must factor in the value configured in the data-delay-interval at the root node any delay configured in IGP-LDP sync timer (**config**>**router**>**if**>**ldp-sync-timer**) on interfaces over the network. This is because the mLDP P2MP LSP may move to a different interface at the expiry of this timer since the routing upstream of the LDP Label Mapping message may change when this timer expires and the interface metric is restored.

At the expiry of this timer, the VPLS/B-VPLS will begin forwarding of BUM packets over the P2MP LSP instance even if not all the S2L paths are up.

The **no** version of this command re-instates the default value for this delay timer.

Parameters *seconds* — Specifies the delay time value in seconds

Values 3 to 180

Default 15

mldp

Syntax [**no**] **mldp**

Context config>service>vpls>provider-tunnel>inclusive

Description This command creates the context to configure the parameters of an LDP P2MP LSP used for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLSs instance.

root-and-leaf

Syntax [**no**] **root-and-leaf**

Context config>service>vpls>provider-tunnel>inclusive

Description This command configures the node to operate as both root and leaf of the I-PMSI in a specified VPLS/B-VPLS instance.

By default, a node will behave as a leaf only node. When the node is leaf only for the I-PMSI of type P2MP RSVP LSP, no PMSI Tunnel Attribute is included in BGP-AD route update messages and therefore no RSVP P2MP LSP is signaled but the node can join RSVP P2MP LSP rooted at other PE nodes participating in this VPLS/B-VPLS service. The user must still configure a LSP template even if the node is a leaf only.

For the I-PMSI of type mLDP, the leaf-only node will join I-PMSI rooted at other nodes it discovered but will not include a PMSI Tunnel Attribute in BGP-AD route update messages. This way a leaf only node will forward packets to other nodes in the VPLS/B-VPLS using the point-to-point spoke-SDPs.

The **no** version of this command re-instates the default value.

rsvp

Syntax [**no**] **rsvp**

Context config>service>vpls>provider-tunnel>inclusive

Description This command creates the context to configure the parameters of an RSVP P2MP LSP used for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLS instance.

lsp-template

Syntax **lsp-template** *p2mp-lsp-template-name*
no lsp-template

Context config>service>vpls>provider-tunnel>inclusive>rsvp

Description This command specifies the template name of the RSVP P2MP LSP instance to be used by the leaf node or the root-and-leaf node that participates in BGP-AD VPLS. The P2MP LSP is referred to as the Inclusive Provider Multicast Service Interface (I-PMSI).

After the user performs a **no shutdown** under the context of the inclusive node and the delay timer expires, BUM packets will be forwarded over an automatically signaled instance of the RSVP P2MP LSP specified in the LSP template.

The **no** version of this command removes the P2MP LSP template from the I-PMSI configuration.

Parameters *p2mp-lsp-template-name* — Specifies the name of the P2MP LSP template up to 32 characters in length.

Default None

3.7.2.5.5 VPLS SDP Commands

mesh-sdp

Syntax **mesh-sdp** ***sdp-id***[**:***vc-id***] [**vc-type** {**ether** | **vlan**}] [root-leaf-tag | leaf-ac]
no mesh-sdp ***sdp-id***[**:***vc-id***]

Context config>service>vpls

Description This command binds a VPLS service to an existing Service Distribution Point (SDP). Mesh SDPs bound to a service are logically treated like a single bridge “port” for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke-SDPs and SAPs) and not transmitted on any mesh SDPs.

This command creates a binding between a service and an SDP. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate the SDP with a valid service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default No *sdp-id* is bound to a service.

Special Cases **VPLS —** Several SDPs can be bound to a VPLS. Each SDP must be destined for a different router. If two *sdp-id* bindings terminate on the same router, an error occurs and the second SDP is binding is rejected.

Parameters *sdp-id* — Specifies the SDP identifier

Values 1 to 17407

vc-id — Specifies the virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

Values 1 to 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.
VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

ether — Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke-SDP binding. (hex 5)

vlan — Defines the VC type as VLAN. The top VLAN tag, if a VLAN tag is present, is stripped from traffic received on the pseudowire, and a vlan-tag is inserted when forwarding into the pseudowire. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for mesh SDP bindings.

Note: The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

root-leaf-tag — Specifies a tagging mesh SDP under an E-Tree VPLS. When a tag SDP binding is required, it is created with a root-leaf-tag flag. Only VLAN tag SDP bindings are supported. The VLAN type must be set to VC VLAN type. The root-leaf-tag parameter indicates this SDP binding is a tag SDP that will use a default VID 1 for root and 2 for leaf. The SDP binding tags egress E-Tree traffic with root and leaf VIDs as appropriate. Root and leaf VIDs are only significant between peering VPLS but the values must be consistent on each end. On ingress a tag SDP binding removes the VID tag on the interface between VPLS in the same E-Tree service. The tag SDP receives root tagged traffic and marks the traffic with a root indication internally. This option is not available on BGP EVPN-enabled E-Tree services.

leaf-ac — Specifies an access (AC) mesh SDP binding under a E-Tree VPLS as a leaf access (AC) SDP. The default E-Tree SDP type is a root AC if *leaf-ac* or *root-leaf-tag* is not specified at SDP binding creation. This option is only available when the VPLS is designated as an E-Tree VPLS. BGP EVPN-enabled E-Tree VPLS services support the **leaf-ac** option.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] endpoint [**no-endpoint**] [**root-leaf-tag** | **leaf-ac**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>vpls

Description This command binds a service to an existing Service Distribution Point (SDP). A spoke-SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default No *sdp-id* is bound to a service.

Special Cases **VPLS —** Several SDPs can be bound to a VPLS service. Each SDP must use unique *vc-ids*. An error message is generated if two SDP bindings with identical *vc-ids* terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.

Parameters *sdp-id* — Specifies the SDP identifier

Values 1 to 17407

vc-id — Specifies the virtual circuit identifier

Values 1 to 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.
VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

Values ether, vlan

ether — Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke-SDP binding. (hex 5)

vlan — Defines the VC type as VLAN. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings.
The VLAN VC-type inserts one dot1q tag within each encapsulated Ethernet packet transmitted to the far end and strips one dotQ tag, if a tag is present, from traffic received on the pseudowire.

Note: The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

group-name — Specifies the name of the split horizon group to which the SDP belongs

endpoint — Specifies the service endpoint to which this SDP bind is attached. The service ID of the SDP binding must match the service ID of the service endpoint.

no endpoint — Removes the association of a spoke-SDP with an explicit endpoint name

root-leaf-tag — Specifies a tagging spoke-SDP under an E-Tree VPLS. When a tag SDP binding is required, it is created with a root-leaf-tag flag. Only VLAN tag SDP bindings are supported. The VLAN type must be set to VC VLAN type. The root-leaf-tag parameter indicates this SDP binding is a tag SDP that will use a default VID tag of 1 for root and 2 for leaf. The SDP binding tags egress E-Tree traffic with root and leaf VIDs as appropriate. Root and leaf VIDs are only significant between peering VPLS but the values must be consistent on each end. On ingress a tag SDP binding removes the VID tag on the interface between VPLS in the same E-Tree service. The tag SDP receives root tagged traffic and marks the traffic with a root indication internally. This option is not available on BGP EVPN-enabled E-Tree services.

leaf-ac — Specifies an access (AC) spoke-SDP binding under a E-Tree VPLS as a leaf access (AC) SDP. The default E-Tree SDP binding type is a root AC if *leaf-ac* or *root-leaf-tag* is not specified at SDP creation. This option is only available when the VPLS is designated as an E-Tree VPLS. BGP EVPN-enabled E-Tree VPLS services support the **leaf-ac** option.

control-word

Syntax [**no**] **control word**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command enables the use of the control word on pseudowire packets in VPLS and enables the use of the control word individually on each mesh SDP or spoke-SDP. By default, the control word is disabled. When the control word is enabled, all VPLS packets, including the BPDU frames, are encapsulated with the control word when sent over the pseudowire. The T-LDP control plane behavior is the same as in the implementation of control word for VLL services. The configuration for the two directions of the Ethernet pseudowire should match. The **no** form of the command reverts the mesh SDP or spoke-SDP to the default behavior of not using the control word. The control word must be enabled to use MPLS-TP OAM on a static spoke-sdp terminating in a VPLS.

Default no control word

egress

Syntax **egress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command enters the context to configure egress SDP parameters.

qos

Syntax **qos** *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*
no qos [*network-policy-id*]

Context config>service>apipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>egress
config>service>ipipe>spoke-sdp>egress
config>service>vpls>spoke-sdp>egress
config>service>vpls>mesh-sdp>egress
config>service>pw-template>egress

Description This command is used to redirect pseudowire packets to an egress port queue-group for the purpose of shaping.

The egress pseudowire shaping provisioning model allows the mapping of one ore more pseudowires to the same instance of queues, or policers and queues, which are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface on which the pseudowire packets can be forwarded. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-SDP inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. When a pseudowire FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface on which the pseudowire packet is forwarded. This queue can be a queue-group queue, or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a pseudowire packet.

2. When a pseudowire FC is redirected to use a queue or a policer, and a queue in a queue-group and the queue-group name exists, but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the pseudowire packet is forwarded on.
3. When a pseudowire FC is redirected to use a queue, or a policer and a queue in a queue-group, and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:

When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.

When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the pseudowire packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

1. If a network QoS policy is applied to the egress context of a pseudowire, any pseudowire FC, which is not explicitly redirected in the network QoS policy, will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the pseudowire is redirected to exists and the redirection succeeds, the marking of the packet DEI/dot1-p/DSCP and the tunnel DEI/dot1-p/DSCP/EXP is performed; according to the relevant mappings of the (FC, profile) in the egress context of the network QoS policy applied to the pseudowire. This is true regardless of whether an instance of the queue-group exists or not on the egress port to which the pseudowire packet is forwarded. If the packet profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet DEI/dot1-p and the tunnel DEI/dot1-p/EXP, and the DSCP/prec will be remarked if **enable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the marking of the packet DEI/dot1-p/DSCP and the tunnel DEI/dot1-p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the pseudowire packet is forwarded.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

queue-redirect-group *queue-group-name* — This optional parameter specifies that the *queue-group-name* will be used for all egress forwarding class redirections within the network QoS policy ID. The specified *queue-group-name* must exist as a port egress queue group on the port associated with the IP interface.

egress-instance *instance-id* — Specifies the identification of a specific instance of the queue-group

Values 1 to 16384

ingress

Syntax **ingress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command configures the ingress SDP context.

qos

Syntax **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>apipe>spoke-sdp>ingress
config>service>cpipe>spoke-sdp>ingress
config>service>epipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>ingress
config>service>vpls>spoke-sdp>ingress
config>service>vpls>mesh-sdp>ingress
config>service>pw-template>ingress

Description This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers, which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

Step 1. Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally, for each traffic type (unicast or multicast).

Step 2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface to which the pseudowire packets can be received. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.

Step 3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.

Step 4. Apply this network QoS policy to the ingress context of a spoke-SDP inside a service, or to the ingress context of a pseudowire template, and specify the redirect queue-group name.

Step 5. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

Step 1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

Step 2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

Step 3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as *policer-output-queues*.

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

- If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

- If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:
  - the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.
  - a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group. The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from an R-VPLS spoke-SDP, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.

When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless of whether an instance of the named policer queue-group exists on the ingress FP on which the pseudowire packet is received. The user can apply a QoS filter matching the dot1-p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload IP header if the user enabled the **ler-use-dscp** option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface on which the pseudowire packet is received.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

queue-group-name — Specifies the name of the queue group template up to 32 characters in length

instance-id — Specifies the identification of a specific instance of the queue-group

Values 1 to 16384

mfib-allowed-mda-destinations

Syntax **mfib-allowed-mda-destinations**

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress

Description This command enters the context to configure MFIB-allowed MDA destinations.

The allowed-mda-destinations node and the corresponding **mda** command are used on spoke and mesh SDP bindings to provide a list of MDA destinations in the chassis that are allowed as destinations for multicast streams represented by [*,g] and [s,g] multicast flooding records on the VPLS service. The MDA list only applies to IP multicast forwarding when IGMP snooping is enabled on the VPLS service. The MDA list has no effect on normal VPLS flooding such as broadcast, L2 multicast, unknown destinations or non-snooped IP multicast.

At the IGMP snooping level, a spoke or mesh SDP binding is included in the flooding domain for an IP multicast stream when it has either been defined as a multicast router port, received a IGMP query through the binding or has been associated with the multicast stream through an IGMP request by a host over the binding. Due to the dynamic nature of the way that a spoke or mesh SDP binding is associated with one or more egress network IP interfaces, the system treats the binding as appearing on all network ports. This causes all possible network destinations in the switch fabric to be included in the multicast streams flooding domain. The MDA destination list provides a simple mechanism that narrows the IP multicast switch fabric destinations for the spoke or mesh SDP binding.

If no MDAs are defined within the allowed-mda-destinations node, the system operates normally and will forward IP multicast flooded packets associated with the spoke or mesh SDP binding to all switch fabric taps containing network IP interfaces.

The MDA inclusion list should include all MDAs that the SDP binding may attempt to forward through. A simple way to ensure that an MDA that is not included in the list is not being used by the binding is to define the SDP the binding is associated with as MPLS and use an RSVP-TE LSP with a strict egress hop. The MDA associated with the IP interface defined as the strict egress hop should be present in the inclusion list. If the inclusion list does not currently contain the MDA that the binding is forwarding through, the multicast packets will not reach the destination represented by the binding.

By default, the MDA inclusion list is empty.

If an MDA is removed from the list, the MDA is automatically removed from the flooding domain of any snooped IP multicast streams associated with a destination on the MDA unless the MDA was the last MDA on the inclusion list. Once the inclusion list is empty, all MDAs are eligible for snooped IP multicast flooding for streams associated with the SDP binding.

mda

Syntax [**no**] **mda** *mda-id*

Context config>service>vpls>mesh-sdp>egress>mfib-allowed-mda-destinations
config>service>vpls>spoke-sdp>egress>mfib-allowed-mda-destinations

Description This command specifies an MFIB-allowed MDA destination for an SDP binding configured in the system.

Parameters *mda-id* — Specifies an MFIB-allowed MDA destination

Values slot/mda slot:1 to 10
mda:1 to 2

vc-label

Syntax **vc-label** *egress-vc-label*
no vc-label [*egress-vc-label*]

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress

Description This command configures the egress VC label.

Parameters *vc-label* — Specifies the VC egress value that indicates a specific connection

Values 16 to 1048575

vc-label

Syntax **vc-label** *ingress-vc-label*
no vc-label [*ingress-vc-label*]

Context config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>ingress

Description This command configures the ingress VC label.

Parameters *vc-label* — Specifies the VC ingress value that indicates a specific connection.

Values 2048 to 18431

entropy-label

Syntax [**no**] **entropy-label**

Context config>service>epipe>spoke-sdp
config>service>ipipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>pw-template
config>service>vpls>bgp-evpn>mpls
config>service>epipe>bgp-evpn>mpls

Description This command enables or disables the use of entropy labels for spoke-SDPs.

If **entropy-label** is configured, the entropy label and ELI are inserted in packets for which at least one LSP in the stack for the far-end of the tunnel used by the service has advertised entropy-label-capability. If the tunnel is RSVP type, then **entropy-label** must not have been disabled under the **config>router>mpls** or **config>router>mpls>lsp** contexts.

The entropy label and hash label features are mutually exclusive. The entropy label cannot be configured on a spoke-SDP or service where the hash label feature has already been configured.

Default no entropy-label

static-mac

Syntax [**no**] **static-mac** *ieee-mac-address*

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command creates a remote static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Distribution Point (SDP).

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Remote static MAC entries create a permanent MAC address to SDP association in the forwarding database for the VPLS instance so that MAC address will not be learned on the edge device.

Static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SDP from the VPLS forwarding database.

Default none

Parameters *ieee-mac-address* — Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa*-*bb*-*cc*-*dd*-*ee*-*ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

transit-policy

Syntax **transit-policy prefix** *prefix-aasub-policy-id*
no transit-policy

Context config>service>vpls>spoke-sdp

Description This command assigns a transit policy id.

The **no** form of the command removes the transit policy ID from the spoke-SDP configuration.

Default no transit-policy

Parameters *prefix-aasub-policy-id* — Specifies the transit policy ID

Values 1 to 65535

vlan-vc-tag

Syntax **vlan-vc-tag** *0 to 4094*
no vlan-vc-tag [*0 to 4094*]

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command.

Default no vlan-vc-tag

Parameters *0 to 4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

3.7.2.5.6 SAP Subscriber Management Commands

cpu-protection

Syntax **cpu-protection** *policy-id* [**mac-monitoring**]
no cpu-protection

Context config>service>vpls>sap
config>service>template>vpls-sap-template

Description This command assigns an existing CPU protection policy to the associated service SAP. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

If no CPU protection policy is assigned to a service SAP, then a the default policy is used to limit the overall-rate.

Default cpu-protection 254 (for access interfaces)

cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.

Parameters *policy-id* — Specifies an existing CPU protection policy

Values 1 to 255

mac-monitoring — When specified, the per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy

default-msap-policy

Syntax **default-msap-policy** *policy-name*
no default-msap-policy

Context config>service>vpls>sap

Description This command specifies an existing managed SAP policy. Managed SAPs allow the use of policies and a SAP template for the creation of a SAP. Managed SAP policies are created in the **config>subscr-mgmt** context. This command is only applicable to SAPs created as a capture-sap.

Default none

Parameters *msap-policy-name* — Specifies an existing managed SAP policy name up to 32 characters in length

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

Context config>service>vpls>sap

Description This command enters the context to configure subscriber management parameters for this SAP.

Default no sub-sla-mgmt

def-inter-dest-id

Syntax **def-inter-dest-id** {**string** *string* | **use-top-q**}
no def-inter-dest-id

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies a default destination string for all subscribers associated with the SAP. The command also accepts the **use-top-q** flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.

The **no** form of the command removes the default subscriber identification string from the configuration.

no def-sub-id

Default no def-inter-dest-id

Parameters **use-top-q** — Specifies to derive the string based on the top most delineating Dot1Q tag from the SAP's encapsulation

string *string* — Specifies the subscriber identification applicable for a subscriber host.

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sla-profile** context.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Default no def-sla-profile

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

mac-da-hashing

Syntax [**no**] **mac-da-hashing**

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.

This command is only meaningful if subscriber management is enabled and can be configured for this VPLS service.

multi-sub-sap

Syntax **multi-sub-sap** [*subscriber-limit*]
no multi-sub-sap

Context config>service>vpls>sap>sub-sla-mgmt

Description This command configures the maximum number of subscribers for this SAP.

The **no** form of this command returns the default value.

Default 1

Parameters *number-of-sub* — Specifies the maximum number of subscribers for this SAP

Values 2 to 8000

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]
no non-sub-traffic

Context config>service>vpls>sap>sub-sla-mgmt>single-sub

Description This command configures non-subscriber traffic profiles. It is used in conjunction with the **profiled-traffic-only** command on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes the profiles and disables the feature.

Parameters *sub-profile-name* — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile-name — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

subscriber sub-ident-string — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the SAP arp-reply-agent to determine the correct handling of received ARP requests from subscribers.

For SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the service destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>service>vpls>sap>sub-sla-mgmt>single-sub

Description This command enables profiled traffic only for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).

The **no** form of the command disables the command.

single-sub-parameters

Syntax **single-sub-parameters**

Context config>service>vpls>sap>sub-sla-mgmt

Description This command enters the context to configure single subscriber parameters for this SAP.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vpls>sap>sub-sla-mgmt

Description This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of the command removes the default subscriber identification policy from the SAP configuration.

Default no sub-ident-policy

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

3.7.2.6 VPLS Multicast Commands

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command enables fast leave. When IGMP or MLD fast leave processing is enabled, the SR OS will immediately remove a SAP or SDP from the multicast group when it detects an IGMP or MLD “leave” on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP or SDP, and therefore speeds up the process of changing channels ('zapping').

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP. When fast leave is enabled, the configured last-member-query-interval value is ignored.

Default no fast-leave

from-vpls

Syntax **from-vpls** *vpls-id*
no from-vpls

Context config>service>vpls>sap>igmp-snooping>mvr
config>service>vpls>sap>mld-snooping>mvr

Description This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request. IGMP snooping must be enabled on the MVR VPLS.

Default no from-vpls

Parameters *vpls-id* — Specifies the MVR VPLS from which multicast channels should be copied into this SAP

Values *service-id*:1 to 2147483648

group

Syntax [**no**] **group** *grp-ip-address*
[**no**] **group** *grp-ipv6-address*

Context config>service>vpls>igmp-snooping>static
config>service>vpls>mesh-sdp>igmp-snooping>static
config>service>vpls>mesh-sdp>mld-snooping>static
config>service>vpls>sap>igmp-snooping>static
config>service>vpls>sap>mld-snooping>static
config>service>vpls>spoke-sdp>igmp-snooping>static
config>service>vpls>spoke-sdp>mld-snooping>static

Description This command enters the context to add a static multicast group as a (*, G) or as one or more (S,G) records. When a static MLD or IGMP group is added, multicast data for that (*,G) or (S,G) is forwarded to the specific SAP or SDP without receiving any membership report from a host.

Default none

Parameters *grp-ip-address* — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

grp-ipv6-address — Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

group-policy

Syntax **group-policy** *policy-name*
no group-policy

Context config>service>vpls>pim-snooping

config>service>vpls>sap>igmp-snooping>mvr
config>service>vpls>mld-snooping>mvr

Description This command identifies filter policy of multicast groups to be applied to this VPLS entity. The sources of the multicast traffic must be a member of the VPLS.

The **no** form of the command removes the policy association from the VPLS configuration.

Default No group policy is specified.

Parameters *policy-name* — Specifies the group policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported. For details on IGMP policies, refer to "Enabling IGMP Group Membership Report Filtering" in the SR OS *Triple Play Service Delivery Architecture Guide*.

fault-propagation-bmac

Syntax **fault-propagation-bmac** [*mac-name* | *ieee-address*] [**create**]
no fault-propagation-bmac [*mac-name* | *ieee-address*]

Context config>service>vpls>mesh-sdp
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command configures associated BMAC addresses for fault propagation on a B-VPLS SAP or SDP binding. The statement can appear up to four times in the configuration to support four remote BMAC addresses in the same remote B-VPLS. The configured VPLS must be a B-VPLS.

The **no** form of the command removes the specified MAC name or MAC address from the list of Fault Propagation BMAC addresses associated with the SAP (or SDP).

Parameters *mac-name* — Specifies a (predefined) MAC name to associate with the SAP or SDP, indirectly specifying a Fault Propagation BMAC address. Up to 32 characters in length

ieee-address — Specifies a MAC address to associate with the SAP or SDP, directly specifying a Fault Propagation BMAC address. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

force-qinq-vc-forwarding

Syntax [**no**] **force-qinq-vc-forwarding**

Context config>service>epipe>spoke-sdp

config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>pw-template

Description This command forces two VLAN tags to be inserted and removed for spoke and mesh SDPs that have either **vc-type ether** or **vc-type vlan**. The use of this command is mutually exclusive with the **force-vlan-vc-forwarding** command.

The VLAN identifiers and dot 1p/DE bits inserted in the two VLAN tags are taken from the inner tag received on a qinq SAP or qinq mesh/spoke-SDP, or from the VLAN tag received on a dot1q SAP or mesh/spoke-SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**), or taken from the outer tag received on a qtag.* SAP or 0 if there is no service delimiting VLAN tag at the ingress SAP or mesh/spoke-SDP. The VLAN identifiers in both VLAN tags can be set to the value configured in the **vlan-vc-tag** parameter in the **pw-template** or under the mesh/spoke-SDP configuration. In the received direction, the VLAN identifiers are ignored and the dot1p/DE bits are not used for ingress classification. However, the inner dot1p/DE bits are propagated to the egress QoS processing.

The Ethertype inserted and used to determine the presence of a received VLAN tag for both VLAN tags is 0x8100. A different Ethertype can be used for the outer VLAN tag by configuring the pseudowire template with the **use-provisioned-sdp** or **prefer-provisioned-sdp** options and setting the Ethertype using the **sdp vlan-vc-etype** parameter (this Ethertype value is then used for all mesh and spoke-SDPs using that SDP).

The **no** form of this command sets the default behavior.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

Context config>service>epipe>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command forces vc-vlan-type forwarding in the data path for spoke or mesh SDPs which have **ether** vc-type. This command is not allowed on vlan-vc-type SDPs.

The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

The **no** form of this command sets default behavior.

Default disabled

hash-label

Syntax **hash-labe**l [**signal-capability**]
no hash-label

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command enables the use of the hash label on a VLL, VPRN, or VPLS service bound to any MPLS type encapsulated SDP, as well as to a VPRN service using the **auto-bind-tunnel** with the **resolution-filter** set to any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option. This feature is also not supported on multicast packets forwarded using RSVP P2MP LSP or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance. It is, however, supported when forwarding multicast packets using an IES/VPRN spoke-interface.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).

To allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note, however, that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets generated in CPM and that are forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh-sdp, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The 7450 ESS, 7750 SR, and 7950 XRS local PE will insert the flow label interface parameters sub-TLV with F=1 in the pseudowire ID FEC element in the label mapping message for that spoke-sdp or mesh-sdp.

- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the pseudowire but must not insert the hash label in the user and control packets over that spoke-sdp or mesh-sdp. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the 7450 ESS, 7750 SR, and 7950 XRS must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the pseudowire ID FEC element.

The **no** form of this command disables the use of the hash label.

Default no hash-label

Parameters **signal-capability** — Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

igmp-snooping

Syntax **igmp-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>bind

Description This command enables the Internet Group Management Protocol (IGMP) snooping context.

Default none

igmp-host-tracking

Syntax **igmp-host-tracking**

Context config>service>vpls

config>service>vpls>sap

Description This command enters the context to configure IGMP host tracking parameters.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>igmp-host-tracking
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command enables the IGMP or MLD router alert check option.

The **no** form of the command disables the router alert check.

expiry-time

Syntax **expiry-time** *expiry-time*
no expiry-time

Context config>service>vpls>igmp-host-tracking
config>service>vpls>sap>igmp-host-tracking

Description This command configures the time that the system continues to track inactive hosts.

The **no** form of the command removes the values from the configuration.

Default no expiry-time

Parameters *expiry-time* — Specifies the time, in seconds, that this system continues to track an inactive host

Values 1 to 65535

import

Syntax **import** *policy-name*
no import

Context config>service>vpls>sap>igmp-host-tracking

Description This command associates an import policy to filter IGMP packets.

The **no** form of the command removes the values from the configuration.

Default no import

Parameters *policy-name* — Specifies the import policy name

max-num-groups

Syntax **max-num-groups** *max-num-groups*
no max-num-groups

Context config>service>vpls>sap>igmp-host-tracking

Description This command configures the maximum number of multicast groups allowed to be tracked.

The **no** form of the command removes the values from the configuration.

Default no max-num-groups

Parameters *max-num-groups* — Specifies the maximum number of multicast groups allowed to be tracked

Values 1 to 196607

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>igmp-host-tracking
config>service>vpls>sap>igmp-snooping
cconfig>service>vpls>spoke-sdp>igmp-snooping

Description This command configures the maximum number of multicast sources allowed per group.

The **no** form of the command removes the value from the configuration.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources allowed per group

Values 1 to 1000

max-num-grp-sources

Syntax **max-num-grp-sources** [1 to 32000]
no max-num-grp-sources

Context config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>igmp-host-tracking
config>service>vpls>sap>igmp-snooping
cconfig>service>vpls>spoke-sdp>igmp-snooping

Description This command defines the maximum number of multicast (S,G)s that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of (S,G)s, the request is ignored.

The **no** form of this command disables the check.

Default no max-num-grp-sources

Parameters **1 to 32000** — Specifies the maximum number of multicast sources allowed to be tracked per group

import

Syntax **import** *policy-name*
no import

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config> service>vpls> mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP or SDP at any time.

The **no** form of the command removes the policy association from the SAP or SDP.

Default no import

Parameters *policy-name* — Specifies the import policy name. Values can be string up to 32 characters long of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. These policies are configured in the **config>router> policy-options** context The router policy must be defined before it can be imported.

last-member-query-interval

Syntax **last-member-query-interval** *tenths-of-seconds*
no last-member-query-interval

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping

config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP or SDP.

Default 10

Parameters *seconds* — Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 to 50

mcac

Syntax **mcac**

Context config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command configures multicast CAC policy and constraints for this interface.

Default none

if-policy

Syntax **if-policy** *mcac-if-policy-name*
no if-policy

Context config>service>vpls>mesh-sdp>igmp-snooping>mcac
config>service>vpls>mesh-sdp>mld-snooping>mcac
config>service>vpls>sap>igmp-snooping>mcac
config>service>vpls>sap>mld-snooping>mcac
config>service>vpls>spoke-sdp>igmp-snooping>mcac
config>service>vpls>spoke-sdp>mld-snooping>mcac

Description This command assigns existing MCAC interface policy to this interface. MCAC interface policy is not supported with MLD-snooping, hence executing the command in the mld-snooping contexts will return an error.

The **no** form of the command removes the MCAC interface policy association.

Default no if-policy

Parameters *mcac-if-policy-name* — Specifies an existing MCAC interface policy

policy

Syntax **policy** *policy-name*
no policy

Context config>service>vpls>mesh-sdp>igmp-snooping>mcac
config>service>vpls>mesh-sdp>mld-snooping>mcac
config>service>vpls>sap>mld-snooping>mcac
config>service>vpls>sap>igmp-snooping>mcac
config>service>vpls>spoke-sdp>igmp-snooping>mcac
config>service>vpls>spoke-sdp>mld-snooping>mcac

Description This command configures the multicast CAC policy name. MCAC policy is not supported with MLD-snooping, hence executing the command in the mld-snooping contexts will return an error.

Parameters *policy-name* — Specifies the multicast CAC policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

unconstrained-bw

Syntax **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*
no unconstrained-bw

Context config>service>vpls>mesh-sdp>igmp-snooping>mcac
config>service>vpls>mesh-sdp>mld-snooping>mcac
config>service>vpls>sap>igmp-snooping>mcac
config>service>vpls>sap>mld-snooping>mcac
config>service>vpls>spoke-sdp>igmp-snooping>mcac
config>service>vpls>spoke-sdp>mld-snooping>mcac

Description This command configures the bandwidth for the interface's multicast CAC policy traffic. When disabled (**no unconstrained-bw**) there will be no checking of bandwidth constraints on the interface level. When enabled and a policy is defined, enforcement is performed. The allocated bandwidth for optional channels should not exceed the **unconstrained-bw** minus the **mandatory-bw** and the mandatory channels have to stay below the specified value for the **mandatory-bw**. After this interface check, the bundle checks are performed.

Parameters *bandwidth* — The bandwidth assigned for interface's MCAC policy traffic, in kilobits per second (kb/s)

Values 0 to 2147483647

mandatory-bw *mandatory-bw* — Specifies the bandwidth pre-reserved for all the mandatory channels on a specified interface in kilobits per second (kb/s)

If the *bandwidth* value is 0, no mandatory channels are allowed. If *bandwidth* is not configured, then all mandatory and optional channels are allowed.

If the value of *mandatory-bw* is equal to the value of *bandwidth*, then all the unconstrained bandwidth on a specified interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed.

The value of *mandatory-bw* should always be less than or equal to that of *bandwidth*, An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, will result in inconsistent value error.

Values 0 to 2147483647

use-lag-port-weight

Syntax **use-lag-port-weight**
no use-lag-port-weight

Context config>service>vpls>sap>igmp-snooping>mcac>mc-constraints
config>service>vpls>sap>igmp-snooping>mcac>mc-constraints

Description This command enables port weight to be used when determining available bandwidth per level when LAG ports go down/come up. The command is required for correct operation on mixed port-speed LAGs and can be used for non-mixed port-speed LAGs as well.

Default no use-lag-port-weight — The port number is used when determining available BW per level when LAG ports go down/come up.

mc-constraints

Syntax **mc-constraints**

Context config>service>vpls>sap>igmp-snooping>mcac
config>service>vpls>sap>mld-snooping>mcac

Description This command enters the context to configure multicast CAC constraints.

Default none

level

Syntax **level** *level-id* **bw** *bandwidth*
no level *level-id*

Context config>service>vpls>sap>igmp-snooping>mcac>mc-constraints
config>service>vpls>sap>mld-snooping>mcac>mc-constraints

Description This command configures levels and their associated bandwidth for multicast cac policy on this interface.

Parameters *level-id* — Specifies has an entry for each multicast CAC policy constraint level configured on this system

Values 1 to 8

bandwidth — Specifies the bandwidth in kilobits per second (kb/s) for the level.

Values 1 to 2147483647

number-down

Syntax **number-down** *number-lag-port-down*
no number-down

Context config>service>vpls>sap>igmp-snooping>mcac>mc-constraints
config>service>vpls>sap>mld-snooping>mcac>mc-constraints

Description This command configure the number of ports down along with level for multicast cac policy on this interface.

Default not enabled

Parameters *number-lag-port-down* — Specifies that the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface will be as per the levels configured in this context.

Values 1 to 64 (for 64-link LAG)
1 to 32 (for other LAGs)

max-num-groups

Syntax **max-num-groups** *count*
no max-num-groups

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

The **no** form of this command disables the check.

Default no max-num-groups

Parameters *count* — Specifies the maximum number of groups that can be joined on this SAP or SDP

Values 1 to 1000

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>service>vpls>sap>igmp-snooping

Description This command defines the maximum number of multicast sources that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of sources, the request is ignored.

The **no** form of this command disables the check.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources allowed per group

Values 1 to 1000

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>bind>igmp-snpg
config>service>vpls>bind>mld-snpg
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping

Description This command specifies whether a multicast router is attached behind this SAP, SDP, or routed VPLS IP interface.

Configuring these objects as an mrouter-port will have a double effect. Firstly, all multicast traffic received on another SAP, SDP, or routed VPLS IP interface will be copied to this SAP, SDP, or routed VPLS IP interface. Secondly, IGMP/MLD reports generated by the system as a result of a router joining or leaving a multicast group, will be sent to this SAP, SDP, or routed VPLS IP interface.

If two multicast routers exist in the network, one of them will become the active querier. While the other multicast router (non-querier) stops sending IGMP queries, it should still receive reports to keep its multicast trees up-to-date. To support this, the mrouter-port should be enabled on all SAPs, SDPs, or routed VPLS IP interfaces connecting to a multicast router.

The IGMP version to be used for the reports (v1, v2, or v3) can only be determined after an initial query has been received. Until such time, no reports are sent on the SAP, spoke-SDP, or routed VPLS IP interface, even if **mrouter-port** is enabled.

If the **send-queries** command is enabled on this SAP or spoke-SDP, the **mrouter-port** parameter cannot be set.

When PIM-snooping is enabled within a VPLS service, all IP multicast traffic and PIM messages will be sent to any SAP or SDP binding configured with an IGMP-snooping mrouter port. This will occur even without IGMP-snooping enabled, but is not supported in a BGP-VPLS or M-VPLS service.

Default no mrouter-port

mvr

Syntax **mvr**

Context config>service>vpls>igmp-snooping
config>service>vpls>mld-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping

Description This command enters the context to configure Multicast VPLS Registration (MVR) parameters.

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>service>vpls>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping

config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mld-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command configures the IGMP query interval. If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on this SAP or SDP. The configured query-interval must be greater than the configured query-response-interval. If send-queries is not enabled on this SAP or SDP, the configured query-interval value is ignored.

Default 125

Parameters *seconds* — Specifies the time interval, in seconds, that the router transmits general host-query messages

Values 2 to 1024

Values **config>service>vpls>igmp-snooping**: 1 - 65535
config>service>vpls>sap>igmp-snooping: 2 - 1024

query-src-ip

Syntax **query-src-ip** *ip-address*
no query-src-ip

Context config>service>vpls>igmp-snooping

Description This command configures the IP source address used in IGMP queries.

query-response-interval

Syntax **query-response-interval** *seconds*

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.

The configured query-response-interval must be smaller than the configured query-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.

Default 10

Parameters *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host

Values 1 to 1023

query-src-ip

Syntax **query-src-ip** *ipv6-address*
no query-src-ip

Context config>service>vpls>mld-snooping

Description This command configures the IP source address used in MLD queries.

report-src-ip

Syntax **report-src-ip**-*address*
no report-src-ip

Context config>service>vpls>igmp-snooping

Description This parameter specifies the source IP address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

Default 0.0.0.0

Parameters *ip-address* — Specifies the source IP source address in transmitted IGMP reports

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>service>vpls>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If this SAP or SDP is expected to be 'lossy', this parameter may be increased. IGMP snooping on this SAP or SDP is robust to (robust-count-1) packet losses.

If send-queries is not enabled, this parameter will be ignored.

Default 2

Parameters *robust-count* — Specifies the robust count for the SAP or SDP

Values **config>service>vpls>sap>igmp-snooping**: 2 to 7
config>service>vpls>igmp-snooping: 1 to 255

mrp

Syntax **mrp**

Context config>service>vpls
config>service>vpls>mesh-sdp
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command configures Multiple Registration Protocol (MRP) parameters.

mvrp

Syntax **mvrp**

Context config>service>vpls

Description This command configures MVRP parameters.

attribute-table-size

Syntax [**no**] **attribute-table-size** *value*

Context config>service>vpls>mvrp

Description This command controls the number of attributes accepted on a per BVPLS basis. When the limit is reached, no new attributes will be registered.

If a new lower limit (smaller than the current number of attributes) from a local or dynamic IVPLS is being provisioned, a CLI warning will be issued stating that the system is currently beyond the new limit. The value will be accepted, but any creation of new attributes will be blocked under the attribute count drops below the new limit; the software will then start enforcing the new limit.

Default maximum number of attributes

Parameters *value* — Specifies the number of attributes accepted on a per BVPLS basis

Values 1 to 4095 for MVRP

attribute-table-size

Syntax [**no**] **attribute-table-size** *value*

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command controls the number of attributes accepted on a per BVPLS basis. When the limit is reached, no new attributes will be registered.

If a new lower limit (smaller than the current number of attributes) from a local or dynamic IVPLS is being provisioned, a CLI warning will be issued stating that the system is currently beyond the new limit. The value will be accepted, but any creation of new attributes will be blocked under the attribute count drops below the new limit; the software will then start enforcing the new limit.

Default maximum number of attributes

Parameters *value* — Specifies the number of attributes accepted on a per BVPLS basis

Values ESS-7/12: 1 to 2047

Values SR-7/SR-12: 1 to 2047

attribute-table-high-wmark

Syntax [**no**] **attribute-table-high-wmark** *high-water-mark*

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command specifies the percentage filling level of the MMRP attribute table where logs and traps are sent.

Default 95%

Parameters *high-water-mark* — Specifies the utilization of the MRP attribute table of this service at which a table full alarm will be raised by the agent

Values 1% to 100%

attribute-table-low-wmark

Syntax [**no**] **attribute-table-low-wmark** *low-water-mark*

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command specifies the MMRP attribute table low watermark as a percentage. When the percentage filling level of the MMRP attribute table drops below the configured value, the corresponding trap is cleared and/or a log entry is added.

Default 90%

Parameters *low-water-mark* — Specifies utilization of the MRP attribute table of this service at which a table full alarm will be cleared by the agent

Values 1% to 100%

flood-time

Syntax **flood-time** *flood-time*
no flood-time

Context config>service>vpls>mrp
config>service>vpls>mvrp

Description This command configures the amount of time, in seconds, after a status change in the VPLS service during which traffic is flooded. Once that time expires, traffic will be delivered according to the MMRP registrations that exist in the VPLS.

Default 3 seconds

Parameters *flood-time* — Specifies the MRP flood time, in seconds

Values 3 to 600

join-time

Syntax [**no**] **join-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the interval between transmit opportunities that are applied to the Applicant state machine. An instance of this Join Period Timer is required on a per-Port, per-MRP Participant basis. For additional information, refer to IEEE 802.1ak-2007 section 10.7.4.1.

Default 2

Parameters *value* — Specifies the timer value in 10th of seconds for sending join-messages

Values 1 to 10 tenths of a second

leave-time

Syntax [**no**] **leave-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the period of time that the Registrar state machine will wait in the leave state before transitioning to the MT state when it is removed. An instance of the timer is required for each state machine that is in the leave state. The Leave Period Timer is set to the value leave-time when it is started.

A registration is normally in “in” state where there is an MFIB entry and traffic is being forwarded. When a “leave all” is performed (periodically around every 10-15 seconds per SAP/SDP binding - see leave-all-time-below), a node sends a message to its peer indicating a leave all is occurring and puts all of its registrations in leave state.

The peer refreshes its registrations based on the leave all PDU it receives and sends a PDU back to the originating node with the state of all its declarations.

Refer to IEEE 802.1ak-2007 section 10.7.4.2.

Default 30

Parameters *value* — Specifies the timer value in 10th of seconds to hold attributes in a leave-state,

Values 30 to 60

leave-all-time

Syntax [**no**] **leave-all-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the frequency with which the LeaveAll state machine generates LeaveAll PDUs. The timer is required on a per-Port, per-MRP Participant basis. The Leave All Period Timer is set to a random value, T, in the range LeaveAllTime<T<1.5*leave-all-time when it is started. Refer to IEEE 802.1ak-2007 section 10.7.4.3.

Default 100

Parameters *value* — Specifies the timer value in 10th of seconds for refreshing all attributes

Values 60 to 300

mrp-policy

Syntax [**no**] **mrp-policy**-*name*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sd>mrp
config>service>vpls>mesh-sdp>mrp

Description This command instructs MMRP to use the mrp-policy specified in the command to control which Group BMAC attributes will be declared and registered on the egress SAP/Mesh-SDP/Spoke-SDP. The Group BMACs will be derived from the ISIDs using the procedure used in the PBB solution. The Group MAC=standard OUI with the last 24 bits being the ISID value. If the policy-name refers to a non-existing mrp-policy the command should return error. Changes to a mrp-policy are allowed and applied to the SAP/SDPs under which the policy is referenced.

Default no mrp-policy is defined

Parameters *policy-name* — Specifies the redirect policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

periodic-time

Syntax [**no**] **periodic-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the frequency the PeriodicTransmission state machine generates periodic events if the Periodic Transmission Timer is enabled. The timer is required on a per-Port basis. The Periodic Transmission Timer is set to one second when it is started.

Default 10

Parameters *value* — Specifies the timer value in 10th of seconds for re-transmission of attribute declarations

Values 10 to 100

periodic-timer

Syntax [**no**] **periodic-timer**

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command enables or disables the Periodic Transmission Timer.

Default disabled

multicast-info-policy

Syntax **multicast-info-policy** *policy-name*
no multicast-info-policy

Context config>service>vpls

Description This command specifies the multicast policy name configured on this service.

pim-snooping

Syntax [**no**] **pim-snooping**

Context config>service>vpls
config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command enables PIM snooping for the VPLS service. When enabled, it is enabled for all SAPs except default SAPs. A default SAP is a SAP that has a wild card VLAN ID, such as sap 1/1/1:*.

The **no** form of the command removes the PIM snooping configuration.

max-num-groups

Syntax **max-num-groups** *num-groups*
no max-num-groups

Context config>service>vpls>spoke-sdp>pim-snooping
config>service>vpls>sap>pim-snooping

Description This command configures the maximum groups for PIM snooping.

Parameters *num-groups* — Specifies the maximum groups for PIM snooping

Values 1 to 16000

oper-group

Syntax [**no**] **oper-group** *name*

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>bgp>pw-template-binding

Description This command associates the context to which it is configured to the operational group specified in the *nam*e. The oper-group *name* must be already configured under config>service before its name is referenced in this command.

The **no** form of the command removes the association.

Default no oper-group

Parameters *name* — Specifies a character string of maximum 32 ASCII characters identifying the group instance.

monitor-oper-group

Syntax [**no**] **monitor-oper-group** *name*

Context config>service>vpls>site
config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command specifies the operational group to be monitored by the object under which it is configured. The **oper-group** *name* must be already configured under **config>service** before its name is referenced in this command.

The **no** form of the command removes the association.

Default no oper-group

Parameters *name* — Specifies a character string of maximum 32 ASCII characters identifying the group instance

hold-time

Syntax **hold-time** *seconds*
no hold-time

Context config>service>vpls>pim-snooping

Description This command configures the duration that allows the PIM-snooping switch to snoop all the PIM states in the VPLS. During this duration, multicast traffic is flooded in the VPLS. At the end of this duration, multicast traffic is forwarded using the snooped states.

When PIM snooping is enabled in VPLS, there is a period of time when the PIM snooping switch may not have built complete snooping state. The switch cannot build states until the routers connected to the VPLS refresh their PIM messages.

This parameter is applicable only if PIM snooping is enabled.

Parameters *seconds* — Specifies the PIM snooping hold time, in seconds.

Values 0 to 300

Default 90

ipv4-multicast-disable

Syntax [**no**] **ipv4-multicast-disable**

Context config>service>vpls>pim-snooping

Description This command disables PIM snooping for IPv4 multicast traffic within a VPLS service.

The **no** form of the command enables PIM snooping for IPv4 multicast traffic within a VPLS service. To fully remove PIM snooping from a VPLS service it is necessary to issue the no pim-snooping command.

Default no ipv4-multicast-disable

ipv6-multicast-disable

Syntax [**no**] **ipv6-multicast-disable**

Context config>service>vpls>pim-snooping

Description This command disables PIM snooping for IPv6 multicast traffic within a VPLS service.

The **no** form of the command enables PIM snooping for IPv6 multicast traffic within a VPLS service. To fully remove PIM snooping from a VPLS service it is necessary to issue the no pim-snooping command.

Default ipv6-multicast-disable

mode

Syntax **mode** *mode*

Context config>service>vpls>pim-snooping

Description This command sets the PIM snooping mode to proxy or plain snooping.

Parameters *mode* — Specifies PIM snooping mode

Values snooping, proxy

Default proxy

precedence

Syntax **precedence** *precedence-value*| **primary**
no precedence

Context config>service>vpls>spoke-sdp

Description This command configures the spoke-SDP precedence.

Default 4

Parameters *precedence-value* — Specifies the spoke-SDP precedence

Values 0 to 4

primary — Specifies that the precedence is primary

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>vpls>spoke-sdp

Description This command specifies the type of signaling used by this multi-segment pseudowire provider-edge for this service.

When no pw-status-signaling is enabled, a 7450 ESS, 7750 SR, and 7950 XRS will not include the pseudowire status TLV in the initial label mapping message of the pseudowire used for a spoke-SDP. This will force both 7450 ESS, 7750 SR, and 7950 XRS PEs to use the pseudowire label withdrawal method for signaling pseudowire status.

If pw-status-signaling is configured, the node will include the use of the pseudowire status TLV in the initial label mapping message for the pseudowire.

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description This command specifies whether MAC flush messages received from the specified LDP are propagated to all spoke and mesh SDPs within the context of this VPLS service. The propagation will follow the split-horizon principle and any data-path blocking in order to avoid the looping of these messages.

Default no propagate-mac-flush

send-queries

Syntax [**no**] **send-queries**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command specifies whether to send IGMP general query messages on the SAP or SDP.

When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented. If send-queries is not configured, the version command has no effect. The version used will be the version of the querier. This implies that, for example, when we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

Default no send-queries

source

Syntax [**no**] **source** *ip-address*
[**no**] **source** *src-ipv6-address*

Context config>service>vpls>sap>igmp-snooping>static>group
config>service>vpls>spoke-sdp>igmp-snooping>static>group
config>service>vpls>mesh-sdp>igmp-snooping>static>group
config>service>vpls>sap>mld-snooping>static>group
config>service>vpls>spoke-sdp>mld-snooping>static>group
config>service>vpls>mesh-sdp>mld-snooping>static>group

Description This command specifies a IPv4 or IPv6 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the sources that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command in combination with the group is used to create a specific (S,G) static group entry.

Use the **no** form of the command to remove the source from the configuration.

Default none

Parameters *ip-address* — Specifies the IPv4 unicast address

src-ipv6-address — Specifies the IPv6 unicast address.

starg

Syntax [**no**] **starg**

Context config>service>vpls>sap>igmp-snooping>static>group
config>service>vpls>spoke-sdp>igmp-snooping>static>group
config>service>vpls>mesh-sdp>igmp-snooping>static>group
config>service>vpls>sap>mld-snooping>static>group
config>service>vpls>spoke-sdp>mld-snooping>static>group
config>service>vpls>mesh-sdp>mld-snooping>static>group

Description This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

Use the **no** form of the command to remove the starg entry from the configuration.

Default no starg

static

Syntax **static**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command enables access to the context to configure static group addresses. Static group addresses can be configured on a SAP or SDP. When present either as a (*, g) or a (s,g) entry, multicast packets matching the configuration will be forwarded even if no join message was registered for the specific group.

Default none

version

Syntax **version** *version*
no version

Context config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command specifies the version of IGMP or MLD which is running on this SAP or SDP. This object can be used to configure a router capable of running either value. For IGMP or MLD to function correctly, all routers on a LAN must be configured to run the same version of IGMP or MLD on that LAN.

When the **send-query** command is configured, all type of queries generate ourselves are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new “wrong version” counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP or SDP will be the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.

Parameters *version* — Specifies the IGMP or MLD version

Values 1, 2, 3

to-sap

Syntax **to-sap** *sap-id*
no to-sap

Context config>service>vpls>sap>igmp-snooping>mvr

Description In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP.

This command configures the SAP to which the multicast data needs to be copied.

Default no to-sap

Parameters *sap-id* — Specifies the SAP to which multicast channels should be copied

3.7.2.6.1 VPLS DHCP and Anti-Spoofing Commands

anti-spoof

Syntax **anti-spoof** {**ip** | **mac** | **ip-mac**}
no anti-spoof

Context config>service>vpls>sap

Description This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **mac**, **ip-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

The **no** form of the command disables anti-spoof filtering on the SAP.

Default no anti-spoof

Parameters **ip** — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the **anti-spoof ip** command will fail.

mac — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the **anti-spoof mac** command will fail.

ip-mac — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the **anti-spoof ip-mac** command will fail.

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>vpls>sap

Description This command configures the application profile name.

Parameters *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

arp-host

Syntax **arp-host**

Context config>service>vpls>sap

Description This command enters the context to configure ARP host parameters.

host-limit

Syntax **host-limit** *max-num-hosts*
no host-limit

Context config>service>vpls>sap>arp-host

Description This command configures the maximum number of ARP hosts.

The **no** form of the command returns the value to the default.

Default 1

Parameters *max-num-hosts* — Specifies the maximum number of ARP hosts allowed on this SAP

Values 1 to 32767

min-auth-interval

Syntax **min-auth-interval** *min-auth-interval*
no min-auth-interval

Context config>service>vpls>sap>arp-host

Description This command configures the minimum authentication interval.

The **no** form of the command returns the value to the default.

Default 15

Parameters *min-auth-interval* — Specifies the minimum authenticated interval, in minutes

Values 1 to 6000

arp-reply-agent

Syntax **arp-reply-agent** [**sub-ident**]
no arp-reply-agent

Context config>service>vpls>sap

Description This command enables a special ARP response mechanism in the system for ARP requests destined for static or dynamic hosts associated with the SAP. The system responds to each ARP request using the hosts MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.

ARP replies and requests received on a SAP with **arp-reply-agent** enabled will be evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP will be discarded when anti-spoof filtering is enabled.

The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the SAP.

A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.

Static hosts can be defined on the SAP using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the SAP's **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

The **arp-reply-agent** command will fail if an existing static host on the SAP does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the SAP without both an IP address and MAC address will fail.

The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.

The **no** form of the command disables ARP-reply-agent functions for static and dynamic hosts on the SAP.

Default none

Parameters *sub-ident* — Configures the arp-reply-agent to discard ARP requests received on the SAP that are targeted for a known host on the same SAP with the same subscriber identification.

Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.

When arp-reply-agent is enabled with **sub-ident**:

- If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same SAP as the source, the ARP request is silently discarded.
- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the SAP's Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.

force-l2pt-boundary

Syntax [**no**] **force-l2pt-boundary**

Context config>service>vpls>sap

Description Enabling force-l2pt-boundary will force that all SAPs managed by the specified m-vpls instance on the corresponding port will have to have l2pt-termination enabled. This command is applicable only to SAPs created under m-vpls and this regardless the flavor of STP currently being active. It is not applicable to spoke-SDPs.

The execution of this command will fail as soon as at least one of the currently managed SAPs (all SAPs falling within the specified managed-vlan-range) does not have l2pt-termination enabled, and this regardless its admin/operational status.

If force-l2pt-boundary is enabled on a specified m-vpls SAP, all newly created SAPs falling into the specified managed-vlan-range will have l2pt-termination enabled per default.

Extending or adding new range into a managed-vlan-range declaration will fail as soon as there is at least one SAPs falling into the specified vlan-range does not have l2pt-termination enabled.

Disabling l2pt-termination on currently managed SAPs will fail as soon as the force-l2pt-boundary is enabled under corresponding m-vpls SAP.

frame-relay

Syntax **frame-relay**

Context config>service>vpls>sap

Description This command enters the context to configure frame-relay parameters.

frf-12

Syntax [**no**] **frf-12**

Context config>service>vpls>sap>fr

Description This command enables FRF12 headers. This must be set to disabled for this entry to be added to an MLFR bundle.

The **no** form of the command disables FRF12 headers.

ete-fragment-threshold

Syntax **ete-fragment-threshold** *threshold*
no ete-fragment-threshold

Context config>service>vpls>sap>fr>frf-12

Description This command configures the FRF.12 fragmentation threshold.

The **no** form of the command removes the value.

Default 128

Parameters *threshold* — Specifies the maximum length of a fragment to be transmitted.

Values 128 to 512

interleave

Syntax **interleave**
no interleave

Context config>service>vpls>sap>frame-relay>frf.12

Description This command enables interleaving of high priority frames and low-priority frame fragments within a FR SAP using FRF.12 end-to-end fragmentation.

When this option is enabled, only frames of the FR SAP non-expedited forwarding class queues are subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI).

When this option is disabled, frames of all the FR SAP forwarding class queues are subject to fragmentation. The fragmentation header is however not included when the frame size is smaller than the user configured fragmentation size. In this mode, the SAP transmits all fragments of a frame before sending the next full or fragmented frame.

The receive direction of the FR SAP supports both modes of operation concurrently, with and without fragment interleaving.

The **no** form of this command restores the default mode of operation.

Default no interleave

scheduling-class

Syntax **scheduling-class** *class-id*
no scheduling-class

Context config>service>vpls>sap>frame-relay

Description This command specifies the scheduling class to use for this SAP. This object is only applicable for a SAP whose bundle type is set to MLFR.

Parameters *class-id* — Specifies the scheduling class

Values 0 to 3

host-connectivity-verify

Syntax **host-connectivity-verify source-ip** *ip-address* [**source-mac** *ieee-address*] [**interval** *interval*] [**action** {**remove** | **alarm**}]

Context config>service>vpls
config>service>vpls>sap

Description This command enables subscriber host connectivity verification on a specified SAP within a VPLS service. This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

Default no host-connectivity-verify

Parameters *ip-address* — Specifies an unused IP address in the same network for generation of subscriber host connectivity verification packets

ieee-address — Specifies the source MAC address to be used for generation of subscriber host connectivity verification packets

interval — Specifies the interval, in minutes, which specifies the time interval in which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

Values 1 to 6000
A zero value can be used by the SNMP agent to disable host-connectivity-verify.

action {**remove** | **alarm**} — Defines the action taken on a subscriber host connectivity verification failure for a specified host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, and so on). DHCP release will be signaled to corresponding DHCP server. Static host will be never removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

3.7.2.7 BGP Auto-Discovery Commands

bgp

Syntax **bgp**

Context config>service>vpls

Description This command enters the context to configure the BGP related parameters for both BGP AD and BGP VPLS.

bgp-vpls

Syntax **bgp-vpls**

Context config>service>vpls

Description This command enters the context to configure the BGP-VPLS parameters and addressing.

max-ve-id

Syntax **max-ve-id** *value*
no max-ve-id

Context config>service>vpls>bgp-vpls

Description This command configures the allowed range for the VE-id value: locally configured and received in a NLRI. Configuration of a VE-id higher than the value specified in this command is not allowed.

Also upon reception of a higher VE-id in an NLRI imported in this VPLS instance (RT is the configured import RT) the following action must be taken:

- a trap must be generated informing the operator of the mismatch.
- NLRI must be dropped
- no service labels are to be installed for this VE-id

- no new NLRI must be generated if a new offset is required for VE-id.

The **no** form of this command sets the max-ve-id to un-configured. The BGP VPLS status should be administratively down for “no max-ve-id” to be used.

The max-ve-id value can be changed without shutting down bgp-vpls if the newly provisioned value does not conflict with the already configured local VE-ID. If the value of the local-VE-ID is higher than the new max-ve-id value the command is rejected. The operator needs to decrease first the VE-ID before running the command.

The actions taken for other max-ve-id values are described below:

- max-ve-id value higher than all VE-IDs (local and received) is allowed and there are no effects.
- max-ve-id higher than the local VE-ID but smaller than the remote VE-IDs:
  - Provisioning is allowed
  - A warning message will be generated stating that “Higher VE-ID values were received in the BGP VPLS context. Related pseudowires will be removed.”
  - The pseudowires associated with the higher VE-IDs will be removed locally.
  - This is a situation that should be corrected by the operator as the pseudowire may be down just at the local PE, consuming unnecessarily core bandwidth. The higher VE-IDs should be removed or lowered.

If the max-ve-id has increased a BGP route refresh is sent to the VPLS community to get the routes which might have been rejected earlier due to max-ve-id check. Default no max-ve-id – max-ve-id is not configured. A max-ve-id value needs to be provisioned for BGP VPLS to be in “no shutdown” state.

Default no max-ve-id

Parameters *value* — Specifies the allowed range of [1-value] for the VE-id. The configured value must be bigger than the existing VE-ids

Values 1 to 65535

ve-name

Syntax **ve-name** *name*
no ve-name

Context config>service>vpls>bgp-vpls

Description This command creates or edits a ve-name. Just one ve-name can be created per BGP VPLS instance.

The **no** form of the command removes the configured ve-name from the bgp vpls node. It can be used only when the BGP VPLS status is shutdown. The **no shutdown** command cannot be used if there is no ve-name configured.

Default no ve-name

Parameters *name* — Specifies the A character string to identify the VPLS Edge instance up to 32 characters in length

ve-id

Syntax **ve-id** *ve-id-value*
no ve-id

Context config>service>vpls>bgp-vpls>ve-name

Description This command configures a ve-id. Just one ve-id can be configured per BGP VPLS instance. The VE-ID can be changed without shutting down the VPLS Instance. When the VE-ID changes, BGP is withdrawing its own previously advertised routes and sending a route-refresh to all the peers which would result in reception of all the remote routes again. The old pseudowires are removed and new ones are instantiated for the new VE-ID value.

The **no** form of the command removes the configured ve-id. It can be used just when the BGP VPLS status is shutdown. The **no shutdown** command cannot be used if there is no ve-id configured.

Default no ve-id

Parameters *value* — Specifies a two-byte identifier that represents the local instance in a VPLS and is advertised through the BGP NLRI. Must be lower or equal with the max-ve-id.

Values 1 to 65535

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>bgp-vpls

Description This command administratively enables/disables the local BGP VPLS instance. On de-activation an MP-UNREACH-NLRI must be sent for the local NLRI.

The **no** form of the command enables the BGP VPLS addressing and the related BGP advertisement. The associated BGP VPLS MP-REACH-NLRI will be advertised in an update message and the corresponding received NLRIs must be considered to instantiate the data plane. RT, RD usage: same like in the BGP AD solution, if the values are not configured here, the value of the VPLS-id from under the bgp-ad node is used. If VPLS-id value is not configured either the MH site cannot be activated – i.e. no shutdown returns an error. Same applies if a pseudowire template is not specified under the bgp node.

Default shutdown

bgp-ad

Syntax [**no**] **bgp-ad**

Context config>service>vpls

Description This command configures BGP auto-discovery.

pw-template-binding

Syntax **pw-template-binding** *policy-id* [**split-horizon-group** *group-name*] [**import-rt** {*ext-community*, ...(up to 5 max)}]
no pw-template-bind *policy-id*

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command binds the advertisements received with the route target (RT) that matches the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present the pw-template is used for all of them.

The pw-template-binding applies to both BGP-AD and BGP-VPLS if these features are enabled in the VPLS.

For BGP VPLS the following additional rules govern the use of pseudowire-template:

- On transmission the settings for the L2-Info extended community in the BGP Update are derived from the pseudowire template attributes. If multiple pseudowire templates (with or without import-rt) are specified for the same VPLS instance the first pw-template entry will be used.
- On reception the values of the parameters in the L2-Info extended community of the BGP Update are compared with the settings from the corresponding pw-template. The following steps are used to determine the local pw-template:
  - The RT values are matched to determine the pw-template.
  - If multiple pw-templates matches are found from the previous steps, the first configured pw-template entry will be considered.
  - If the values used for Layer 2 MTU or C Flag do not match the pseudowire setup fails.

The tools perform commands can be used to control the application of changes in pw-template for both BGP-AD and BGP-VPLS.

The **no** form of the command removes the values from the configuration.

Default none

Parameters *policy-id* — Specifies an existing policy ID

Values 1 to 2147483647

group-name — The specified group-name overrides the split horizon group template settings

import-rt ext-comm — Specifies communities allowed to be accepted from remote PE neighbors. An extended BGP community in the **type**:**x**:**y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values target:{*ip-addr:comm-val*| *2byte-asnumber:ext-comm-val*|*4byte-asnumber:comm-val*}
ip-addr a.b.c.d
comm-val 0 to 65535
2byte-asnumber 0 to 65535
ext-comm-val 0 to 4294967295
4byte-asnumber 0 to 4294967295

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>service>vpls>bgp>pw-template-binding

Description This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a specified protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP/BGP protocol adjacency.

Default no bfd-enable

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>service>vpls>bgp-ad>pw-template-binding
config>service>vpls>bgp>pw-template-binding
config>service>vpls>spoke-sdp

Description This command enables VCCV BFD on the PW associated with the VLL, BGP VPWS, or VPLS service. The parameters for the BFD session are derived from the named BFD template, which must have been first configured using the **bfd-template** command.

Default no bfd-enable

bfd-template

Syntax **bdf-template** *name*
no bfd-template

Context config>service>vpls>bgp-ad>pw-template-binding
config>service>vpls>bgp>pw-template-binding
config>service>vpls>spoke-sdp

Description This command configures a named BFD template to be used by VCCV BFD on PWs belonging to the VLL, BGP VPWS, or VPLS service. The template specifies parameters, such as the minimum transmit and receive control packet timer intervals, to be used by the BFD session. Template parameters are configured under the **config>router>bfd** context.

Default no bfd-template

Parameters *name* — Specifies a text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

oper-group

Syntax **oper-group** *group-name*
no oper-group

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>bgp>pw-template-binding

Description This command associates the context to which it is configured to the operational group specified in the *group-name*. The **oper-group** *group-name* must be already configured under **config>service** context before its name is referenced in this command.

The **no** form of the command removes the association.

Parameters *group-name* — Specifies a character string of maximum 32 ASCII characters identifying the group instance

route-target

Syntax **route-target** {*ext-community*|{[**export** *ext-community*][**import** *ext-community*]}}
no route-target

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command configures the route target (RT) component that will be signaled in the related MP-BGP attribute to be used for BGP auto-discovery, BGP VPLS and BGP multi-homing if these features are configured in this VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The ext-comm can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community.

The following rules apply:

- if BGP AD VPLS-id is configured and no RT is configured under BGP node, the RT is the VPLS-ID
- if BGP AD VPLS-id is not configured then an RT value must be configured under BGP node (this is the case when only BGP VPLS is configured)
- if BGP AD VPLS-id is configured and an RT value is also configured under BGP node, the configured RT value prevails

Parameters **export** *ext-community* — Specifies communities allowed to be sent to remote PE neighbors

import *ext-community* — Specifies communities allowed to be accepted from remote PE neighbors

vpls-id

Syntax **vpls-id** *vpls-id*

Context config>service>vpls>bgp-ad

Description This command configures the VPLS ID component that will be signaled in one of the extended community attributes (*ext-comm*).

Values and format (6 bytes, other 2 bytes of type-subtype will be automatically generated)

Parameters *vpls-id* — Specifies a globally unique VPLS ID for BGP auto-discovery in this VPLS service

Values vpls-id : <ip-addr:comm-val>|<as-number:ext-comm-val>
ip-addra.b.c.d
comm-val0 to 65535
as-number 1..65535
ext-comm-val 0..4294967295

vsi-export

Syntax **vsi-export** *policy-name* [*policy-name*...(up to 5 max)]
no vsi-export

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command specifies the name of the VSI export policies to be used for BGP auto-discovery, BGP VPLS and BGP multi-homing if these features are configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

vsi-id

Syntax **vsi-id**

Context config>service>vpls>bgp-ad

Description This command enters the context to configure the Virtual Switch Instance Identifier (VSI-ID).

prefix

Syntax **prefix** *low-order-vsi-id*
no prefix

Context config>service>vpls>bgp-ad>vsi-id

Description This command specifies the low-order 4 bytes used to compose the Virtual Switch Instance Identifier (VSI-ID) to use for NLRI in BGP auto-discovery in this VPLS service.

If no value is set, the system IP address will be used.

Default no prefix

Parameters *low-order-vsi-id* — Specifies a unique VSI ID

Values 0— 4294967295

route-distinguisher

Syntax **route-distinguisher** *rd*
route-distinguisher auto-rd
no route-distinguisher

Context config>service>vpls>bgp

Description This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for L2VPN and EVPN families. This value will be used for BGP-AD, BGP VPLS and BGP multi-homing NLRI, if these features are configured.

If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:

- if BGP AD VPLS-id is configured and no RD is configured under BGP node - RD is the VPLS-ID
- if BGP AD VPLS-id is not configured then an RD value must be configured under BGP node (this is the case when only BGP VPLS is configured)
- if BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails

Values and format (6 bytes, other 2 bytes of type will be automatically generated)

Alternatively, the **auto-rd** option allows the system to automatically generate an RD based on the **bgp-auto-rd-range** command configured at service level.

Parameters *ip-addr:comm-val* — Specifies the IP address

Values ip-addra.b.c.d
comm-val0 to 65535

as-number:ext-comm-val — Specifies the AS number

Values as-number1 to 65535
ext-comm-val 0 to 4294967295

auto-rd — The system generates an RD for the service according to the IP address and range configured in the bgp-auto-rd-range command

vsi-import

Syntax **vsi-import** *policy-name* [*policy-name*...(up to 5 max)]
no vsi-import

Context config>service>vpls>bgp-ad>vsi-id
config>service>vpls>bgp

Description This command specifies the name of the VSI import policies to be used for BGP auto-discovery, BGP VPLS and BGP multi-homing if these features are configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

bgp-evpn

Syntax [**no**] **bgp-evpn**

Context config>service>vpls

Description This command enters the context to configure the BGP EVPN parameters.

mac-advertisement

Syntax [**no**] **mac-advertisement**

Context config>service>vpls>bgp-evpn

Description The mac-advertisement command enables the advertisement in BGP of the learned MACs on SAPs and SDP bindings. When the mac-advertisement is disabled, the local macs will be withdrawn in BGP.

Default mac-advertisement

mac-duplication

Syntax **mac-duplication**

Context config>service>vpls>bgp-evpn

Description This command enters the context to configure the BGP EVPN mac duplication parameters.

detect

Syntax **detect num-moves** *num-moves* **window** *minutes*

Context config>service>vpls>bgp-evpn>mac-duplication

Description Mac-duplication is always enabled. This command modifies the default behavior. Mac-duplication monitors the number of moves of a MAC address for a period of time (window).

Default num-moves 5 window 3

Parameters **num-moves** — Identifies the number of mac moves in a VPLS service. The counter is incremented when a specified MAC is locally relearned in the FDB or flushed from the FDB due to the reception of a better remote EVPN route for that MAC.

Values 3 to 10 minutes

Default 3 minutes

retry

Syntax **retry** *minutes*
no retry

Context config>service>vpls>bgp-evpn>mac-duplication

Description Specifies the timer after which the MAC in hold-down state is automatically flushed and the mac-duplication process starts again. This value is expected to be equal to two times or more than that of window.

If **no** retry is configured, this implies that, when mac-duplication is detected, mac updates for that mac will be held down till the user intervenes or a network event (that flushes the mac) occurs.

Default 9 minutes

Parameters *minutes* — Specifies the retry timer in minutes

Values 2 to 60 minutes

unknown-mac-route

Syntax [**no**] **unknown-mac-route**

Context config>service>vpls>bgp-evpn

Description This command enables the advertisement of the unknown-mac-route in BGP. This will be coded in an EVPN mac route where the mac address is zero and the mac address length 48. By using this unknown-mac-route advertisement, the user may decide to optionally turn off the advertisement of MAC addresses learned from SAPs and sdp-bindings, hence reducing the control plane overhead and the size of the FDB tables in the data center. All the receiving NVEs supporting this concept will send any unknown-unicast packet to the owner of the unknown-mac-route, as opposed to flooding the unknown-unicast traffic to all other nodes part of the same VPLS. Although the 7450 ESS, 7750 SR, and 7950 XRS can be configured to generate and advertise the unknown-mac-route, the 7450 ESS, 7750 SR, and 7950 XRS will never honor the unknown-mac-route and will flood to the vpls flood list when an unknown-unicast packet arrives to an ingress sap/sdp-binding.

Default no unknown-mac-route

vxlan

Syntax **vxlan**

Context config>service>vpls>bgp-evpn

Description This command enters the context to configure the VXLAN parameters when BGP EVPN is used as the control plane.

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>bgp-evpn>vxlan

Description This command enables/disables the automatic creation of VXLAN auto-bindings by BGP-EVPN.

Default shutdown

3.7.2.8 Redundancy Commands

redundancy

Syntax **redundancy**

Context config

Description This command enters the context to perform redundancy operations.

multi-chassis

Syntax **multi-chassis**

Context config>redundancy

Description This command enters the context to configure multi-chassis parameters.

peer

Syntax [**no**] **peer** *ip-address* **create**

Context config>redundancy>multi-chassis

Description Use this command to configure up to 20 multi-chassis redundancy peers. It is only for mc-lag (20) not for mc-sync (4).

Parameters *ip-address* — Specifies the IP address

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x:-[0 —FFFF]H
	d: [0 — 255]D

create — Specifies to create the peer

sync

Syntax [**no**] **sync**

Context config>redundancy>multi-chassis>peer

Description This command enters the context to configure synchronization parameters.

igmp-snooping

Syntax [**no**] **igmp-snooping**

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether IGMP snooping information should be synchronized with the multi-chassis peer.

Default no igmp-snooping

mld-snooping

Syntax [**no**] **mld-snooping**

Context config>redundancy>multi-chassis>peer>sync

Description This command is not supported. It is not blocked for backwards-compatibility reasons but has no effect on the system if configured.

pim-snooping

Syntax **pim-snooping** [**saps**] [**spoke-sdps**]
no pim-snooping

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether PIM snooping for IPv4 information should be synchronized with a multi-chassis peer. Entering **pim-snooping** without any parameters results in the synchronization being applied only to SAPs.

Specifying the **spoke-sdps** parameter results in the synchronization being applied to manually configured spoke-SDPs. Specifying both the **saps** and **spoke-sdps** parameters results in the synchronization being applied to both SAPs and manually configured spoke-SDPs.

The synchronization of PIM snooping is only supported for manually configured spoke-SDPs but is not supported for spoke-SDPs configured within an endpoint. See PIM Snooping for IPv4 Synchronization for service support.

Default no pim-snooping

Parameters **saps** — Specifies that SAPs are to be synchronized with the multi-chassis peer according to the synchronization tags configured on the port. This is the default when no parameters are specified.

spoke-sdps — Specifies that spoke-SDPs are to be synchronized with the multi-chassis peer according to the synchronization tags configured on spoke-SDPs.

port

Syntax **port** [*port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**create**]
no port [*port-id* | *lag-id*]

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies the port to be synchronized with the multi-chassis peer and a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

Parameters *port-id* — Specifies the port to be synchronized with the multi-chassis peer

lag-id — Specifies the LAG ID to be synchronized with the multi-chassis peer

sync-tag *sync-tag* — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer, up to 32 characters in length

range

Syntax **range** *encap-range* **sync-tag** *sync-tag*
no range *encap-range*

Context config>redundancy>multi-chassis>peer>sync>port

Description This command configures a range of encapsulation values.

Parameters *encap-range* — Specifies a range of encapsulation values on a port to be synchronized with a multi-chassis peer

Values

Dot1q	*start-vlan-end-vlan*
QinQ	Q1.*start-vlan*-Q1.*end-vlan*

sync-tag *sync-tag* — Specifies a synchronization tag, up to 32 characters in length, to be used while synchronizing this encapsulation value range with the multi-chassis peer.

sdp

Syntax **sdp** *sdp-id* [**sync-tag** *sync-tag*] [**create**]
no sdp *sdp-id*

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies the manually configured spoke-SDPs to be synchronized with the multi-chassis peer and a synchronization tag to be used while synchronizing these spoke-SDPs with the multi-chassis peer.

Manually configured spoke-SDPs with the specified *sdp-id* will be synchronized according to the synchronization tag. If synchronization is required only for a subset of the spoke-SDPs using the configured SDP, the **range** sub-command should be used. The **range** command and the **sync-tag** parameters are mutually exclusive.

The synchronization of PIM snooping is only supported for manually configured spoke-SDPs but is not supported for spoke-SDPs configured within an endpoint. See PIM Snooping for IPv4 Synchronization for service support.

The synchronization of PIM snooping is not supported on any of the following when used with the configured *sdp-id*:

- Mesh SDPs
- Spoke SDPs in non-VPLS services
- BGP-AD/BGP-VPLS (FEC 129) spoke-SDPs
- Spoke SDPs configured in endpoints
- Pseudowire SAPs
- ESM-over-MPLS Pseudowires

Non-existent spoke-SDPs may be specified. If these spoke-SDPs are created at a later time, then all states on the spoke-SDPs will be synchronized according to the synchronization tag and the synchronization protocols enabled.

A synchronization tag can be changed by entering the same command with a different synchronization tag. Changing the synchronization tag removes all states relating to the previous synchronization tag for the SDP and a new synchronization tag state is created.

Parameters *sdp-id* — Specifies the SDP of the spoke-SDPs to be synchronized with the multi-chassis peer

Values 1 to 17407

sync-tag — Specifies a synchronization tag, up to 32 characters in length, to be used when synchronizing with the multi-chassis peer.

range

Syntax **range** *vc-id-range* [**sync-tag** *sync-tag*]

no range *vc-id-range*

Context config>redundancy>multi-chassis>peer>sync>sdp

Description This command specifies a range of VC IDs for manually configured spoke-SDPs to be synchronized with the multi-chassis peer and a synchronization tag to be used while synchronizing each range with the multi-chassis peer. The **range** command and the configuration of a synchronization tag on the parent **sdp** command are mutually exclusive.

To synchronize a single spoke-SDP, the *start-vc-id* should be the same as the *end-vc-id*. If the configured *end-vc-id* is lower than the *start-vc-id*, the **range** command fails.

The synchronization tag can be changed by entering the same command with a different synchronization tag. Changing the synchronization tag removes all states relating to the previous synchronization tag for the SDP and a new synchronization tag state is created.

Multiple **range** commands can be configured, however, overlapping ranges for the same SDP (*sdp-id*) are not permitted.

The synchronization of PIM snooping is only supported for manually configured spoke-SDPs but is not supported for spoke-SDPs configured within an endpoint. See PIM Snooping for IPv4 Synchronization for service support.

The synchronization of the PIM snooping state is not supported on any of the following when used with the configured *sdp-id*:

- mesh SDPs
- spoke-SDPs in non-VPLS services
- BGP-AD/BGP-VPLS (FEC 129) spoke-SDPs
- spoke-SDPs configured in endpoints
- pseudowire SAPs
- ESM-over-MPLS pseudowires

Non-existent spoke-SDPs may be specified. If these spoke-SDPs are created at a later time, then all states on the spoke-SDPs will be synchronized according to the synchronization tag and the synchronization protocols enabled.The **sync-tag** can be changed by entering the same command with a different **sync-tag** value. If the synchronization tag is changed, then all states for the previous **sync-tag** are removed for the SDP configured in the command and the state is then built for the new synchronization tag.

Parameters *vc-id-range* — Specifies a non-overlapping range of VC IDs for the spoke-SDPs of the SDP to be synchronized with the multi-chassis peer

Values *start-vc-id-end-vc-id*
start-vc-id: 1 to 4294967295
end-vc-id: 1 to 4294967295

sync-tag — Specifies a synchronization tag, up to 32 characters in length, to be used when synchronizing with the multi-chassis peer

3.8 VPLS Show, Clear, Debug, and Tools Command Reference

3.8.1 Command Hierarchies

3.8.1.1 Show Commands

```
show
    — service
        — active-subscribers summary
        — active-subscribers [subscriber sub-ident-string]  [sap-id sap-id] [sla-profile sla-
            profile-name] [detail | mirror]
        — active-subscribers hierarchy [subscriber sub-ident-string]
        — egress-label egress-label1 [egress-label2]
        — fdb-info
        — fdb-mac ieee-address [expiry]
        — id service-id
            — all
            — arp-host [wholesaler service-id] [sap sap-id | interface interface-name | ip-
                address ip-address[/mask] | mac ieee-address | {[port port-id] [no-inter-
                dest-id | inter-dest-id inter-dest-id]}] [detail]
            — arp-host statistics [sap sap-id | interface interface-name]
            — arp-host summary [interface interface-name]
            — authentication
                — statistics [policy name] [sap sap-id]
            — base [msap]
            — bgp-evpn
            — dhcp
                — lease-state [[sap sap-id] | [sdp sdp-id:vc-id] | [interface interface-
                    name] | [ip-address ip-address]] [detail] |[mac ieee-address]
                    |[wholesaler service-id] [detail]
                — lease-state [sap sap-id]
                — statistics [sdp sdp-id:vc-id]
                — statistics [interface interface-name]
                — summary
            — endpoint [endpoint-name]
            — etree
            — fdb [sap sap-id] [expiry]] | [sdp sdp-id [expiry]] | [mac ieee-address [expiry]] |
                endpoint endpoint | [detail] [expiry] [pbb]
```

- **gsmp**
  - **neighbors group** [*name*] [*ip-address*]
  - **sessions** [**group** *name*] **neighbor** *ip-address*] [**port** *port-number*] [**association**] [**statistics**]
- **host** [**sap** *sap-id*] [**detail**]
- **host summary**
- **i-vpls**
- **igmp-snooping**
  - **all**
  - **base**
  - **mrouters** [**detail**]
  - **mvr**
  - **port-db sap** *sap-id* [**detail**]
  - **port-db sap** *sap-id* **group** *grp-address*
  - **port-db sdp** *sdp-id*:*vc-id* [**detail**]
  - **port-db sdp** *sdp-id*:*vc-id* **group** *grp-address*
  - **port-db vxlan vtep** *ip-address* **vni** [0..4294967295]
  - **proxy-db** [**detail**]
  - **proxy-db** [**group grp-ip-address**]
  - **querier**
  - **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
  - **statistics** [**evpn-mpls** | **sap** *sap-id* | **sdp** *sdp-id:vc-id* | **vxlan vtep** *ip-address* **vni** *vni*]
- **isid-policy**
- **labels**
- **l2pt disabled**
- **l2pt** [**detail**]
- **mac-move**
- **mac-protect**
- **mfib** [**ipv4** | **ipv6** | **mac**]
- **mfib brief**
- **mfib group** *group-address* [**statistics**]
- **mfib statistics** [**ipv4** | **ipv6** | **mac**]
- **mld-snooping**
  - **all**
  - **base**
  - **mrouters** [**detail**]
  - **mvr**
  - **port-db sap** *sap-id*
  - **port-db sap** *sap-id* **detail**
  - **port-db sap** *sap-id* **group** *grp-ipv6-address*
  - **port-db sdp** *sdp-id:vc-id*
  - **port-db sdp** *sdp-id:vc-id* **detail**
  - **port-db sdp** *sdp-id:vc-id* **group** *grp-ipv6-address*
  - **proxy-db** [**detail**]
  - **proxy-db group** *grp-ipv6-address*
  - **querier**
  - **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
  - **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
- **mmrp mac** [*ieee-address*]
- **mrp-policy** [*mrp-policy*]
- **mrp-policy** *mrp-policy* [**association**]
- **mrp-policy** *mrp-policy* [**entry** *entry-id*]
- **mstp-configuration**

- **pim-snooping**
  - **group** [*grp-ip-address*] [**source** *ip-address*] [**type** {**starg** | **sg**}] [**detail**] [*family*]
  - **neighbor** [{**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**address** *ip-address*]] [**detail**] [*family*]
  - **port** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**group** [*grp-ip-address*]] [**detail**] [*family*]
  - **statistics** [**sap** *sap-id*] [**sdp** *sdp-id:vc-id*] [*family*]
  - **status** [*family*]
- **provider-tunnels**
- **proxy-arp** [**ip-address** *ip-address*] [**detail**]
- **proxy-nd** [**ip-address** *ip-address*] [**detail**]
- **retailers**
- **sap** [*sap-id* [*filter*] [**detail**]]
- **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]
- **sdp** *sdp-id:vc-id* {**mrp** | **mmrp**}
- **site** [**detail**]
- **site** *name*
- **split-horizon-group** [*group-name*]
- **stp** **mst-instance** *mst-inst-number*
- **stp** [**detail**]
- **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[*/mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]
- **system**
  - **fdb-usage** [**card** *slot-id*]
- **wholesalers**
- **vxlan**

- **ingress-label** *start-label* [*end-label*]
- **isid-using** [*ISID*]
- **sap-using**[**msap**] [**dyn-script**] [**description**]
- **sap-using** [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**]
- **sap-using** **app-profile** *app-profile-name*
- **sap-using** **authentication-policy** *policy-name* [**msap**]
- **sap-using** **encap-type** *encap-type*
- **sap-using** **eth-cfm** **collect-lmm-stats** [**sap** *sap-id*]
- **sap-using** **eth-ring** [**ring-id** *eth-ring-id*]
- **sap-using** **eth-tunnel** [**tunnel-id** *eth-tunnel-id*]
- **sap-using** **interface** [*ip-address* | *ip-int-name*]
- **sap-using** [**ingress** | **egress**] **atm-td-profile** *td-profile-id*
- **sap-using** [**ingress** | **egress**] **filter** ***f****ilter-id*
- **sap-using** [**ingress** | **egress**] **qos-policy** *qos-policy-id*
- **sap-using** **authentication-policy** *policy-name*
- **sap-using** **mc-ring peer** *ip-address* **ring** *sync-tag*
- **sap-using** **process-cpm-traffic-on-sap-down**
- **sap-using** **etree**
- **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]
- **sdp** [*sdp-id*[*:vc-id*] | **far-end** *ip-address*]
- **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail** | **keep-alive-history**]
- **sdp-using** [*sdp-id*[*:vc-id*] | **far-end** *ip-address*]
- **sdp-using** **e-tree**
- **service-using** [**vpls**] [**b-vpls**] [**i-vpls**] [**m-vpls**]
- **service-using**[**msap**] [**dyn-script**] [**description**] **e-tree**

— **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address*[*/mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]
— **vxlan**

show
— **egress-replication**
— **vpls** *vpls-service-id* **mda** *slot/mda*
— **vpls** *vpls-service-id* **mda** *slot/mda* [**igmp-record** *grp-address* {**source** *source-ip-address* | **starg**}] | [**mRouter**]

show
— **igmp**
— **group** [*grp-ip-address*]
— **ssm-translate**
— **interface** [*ip-int-name* | *ip-address*] [**group** *grp-address*] [**detail**]
— **static** [*ip-int-name* | *ip-addr*]
— **statistics** [*ip-int-name* | *ip-address*]
— **status**

3.8.1.1.1 Show Multi-Chassis Endpoint Commands

show
— **service**
— **id** *service-id*
— **endpoint** [*endpoint-name*]
— **redundancy**
— **multi-chassis**
— **mc-endpoint statistics**
— **mc-endpoint peer** [*ip-address*] **statistics**
— **mc-endpoint endpoint** [*mcep-id*] **statistics**
— **mc-endpoint peer** [*ip-address*]

3.8.1.2 Clear Commands

clear
— **service**
— **id** *service-id*
— **arp-host** { **mac** *ieee-address* | **sap** *sap-id* | **ip-address** *ip-address*[*/mask*]}
— **arp-host** [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id* | **no-inter-dest-id**]
— **arp-host statistics** [**sap** *sap-id* | **interface** *interface-name*]
— **authentication**
— **statistics**
— **dhcp**
— **lease-state** [**no-dhcp-release**]
— **lease-state ip-address** [*ip-address*[*/mask*]] [**no-dhcp-release**]
— **lease-state mac** *ieee-address* [**no-dhcp-release**]

- — **lease-state sap** *sap-id* [**no-dhcp-release**]
- — **lease-state sdp** *sdp-id:vc-id* [**no-dhcp-release**]
- — **statistics** [**sap** *sap-id* | **sdp** *sdp-id*:*vc-id* | **interface** *ip-int-name* | *ip-address*]

— **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id* | **mesh-sdp** *sdp-id*[**:***vc-id*] | **spoke-sdp** *sdp-id*:*vc-id*}

— **igmp-snooping**

- — **port-db sap *sap-id*** [**group** *grp-address* [**source** *ip-address*]]
- — **port-db sdp *sdp-id:vc-id*** [**group** *grp-address* [**source** *ip-address*]]
- — **port-db group** *grp-ip-address* **vxlan vtep** *ip-address* **vni** *vni-id*
- — **port-db group** *grp-ip-address* **source** *src-ip-address* **vxlan-vtep** *ip-address* **vni** *vni-id*]
- — **querier**
- — **statistics** {**evpn-mpls** | **all** | **sap** *sap-id* | **sdp** *sdp-id*:*vc-id* | **vxlan vtep** *ip-address* **vni** *vni-id*}

— **mfib**

- — **statistics** {**all** | **ipv4** | **ipv6** | **mac**}
- — **statistics group** *grp-address*

— **mld-snooping**

- — **port-db sap** *sap-id* [**group** *grp-ipv6-address*]
- — **port-db sap** *sap-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*
- — **port-db sdp** *sdp-id*:*vc-id* [**group** *grp-ipv6-address*]
- — **port-db sdp** *sdp-id*:*vc-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*
- — **querier**
- — **statistics all**
- — **statistics sap** *sap-id*
- — **statistics sdp** *sdp-id*:*vc-id*

— **mesh-sdp** *sdp-id*[**:***vc-id*] **ingress-vc-label**

— **msap** *msap-id*

— **pim-snooping**

- — **database** [[**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**group** *grp-ip-address*] [**source** *src-ip-address*]] [*family*]
- — **neighbor** [*ip-address* | **sap** *sap-id* | **sdp** *sdp-id:vc-id*] [*family*]
- — **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [*family*]

— **spoke-sdp** *sdp-id*:*vc-id* **ingress-vc-label**

— **proxy-arp** {**all** | *ip-address*} [{**dynamic** | **dup**}]

— **proxy-nd** {**all** | *ipv6-address*} [{**dynamic** | **dup**}]

— **stp**

- — **detected-protocols** [**all** | **sap** *sap-id* | **spoke-sdp** [*sdp-id*[:*vc-id*]]

— **statistics**

— **id** *service-id*

- — **capture-sap** *sap-id* [*trigger*]
- — **cem**
- — **counters**
- — **l2pt**
- — **mesh-sdp sdp-id**[*:vc-id*] {**all** | **counters** | **stp** | **mrp**}
- — **mrp**
- — **pip**
- — **spoke-sdp** *sdp-id*[:*vc-id*] {**all** | **counters** | **stp** | **l2pt** | **mrp**}
- — **spoke-sdp**
- — **stp**

— **sap** *sap-id* {**all** | **cem** | **counters** | **l2pt** | **stp** | **mrp**}

— **sdp** *sap-id* {**keep-alive**}

```
clear
    — router
        — dhcp
            — statistics [interface ip-int-name | ip-address]
```

3.8.1.3 Debug Commands

```
debug
    — service
        — id service-id
            — [no] arp-host
            — igmp-snooping
                — detail-level {low | medium | high}
                — no detail-level
                — [no] evpn-mpls
                — [no] mac ieee-address
                — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
                — no mode
                — [no] sap sap-id
                — [no] sdp sdp-id:vc-id
                — [no] vxlan vtep vtep vni vni-id
            — mld-snooping
                — detail-level {low | medium | high}
                — no detail-level
                — [no] mac ieee-address
                — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
                — no mode
                — [no] sap sap-id
                — [no] sdp sdp-id:vc-id
            — [no] mrp
                — all-events
                — [no] applicant-sm
                — [no] leave-all-sm
                — [no] mmrp-mac ieee-address
                — [no] mrpdu
                — [no] periodic-sm
                — [no] registrant-sm
                — [no] sap sap-id
                — [no] sdp sdp-id:vc-id
            — [no] event-type {config-change | svc-oper-status-change | sap-oper-
              status-change | sdpbind-oper-status-change}
            — [no] host-connectivity-verify
                — [no] ip ip-address
                — [no] mac ieee-address
                — [no] sap sap-id
            — [no] pim-snooping
                — [no] adjacency
                — all [group grp-ip-address] [source ip-address] [detail]
                — no all
                — database [group grp-ip-address] [source ip-address] [detail]
                — no database
```

```
                    — jp [group grp-ip-address] [source ip-address] [detail]
                    — no jp
                    — mcs [detail]
                    — no mcs
                    — packet [hello | jp] [sap sap-id | sdp sdp-id:vc-id]
                    — [no] packet
                    — port [sap sap-id | sdp sdp-id:vc-id] [detail]
                    — no port
                    — red [detail]
                    — no red
                — [no] proxy-arp [mac [ieee-address]] [ip [ipaddr] [all]]
                — [no] proxy-nd [mac [ieee-address]] [ip [ipaddr] [all]]
                — [no] sap sap-id
                — stp
                    — all-events
                    — [no] bpdu
                    — [no] core-connectivity
                    — [no] exception
                    — [no] fsm-state-changes
                    — [no] fsm-timers
                    — [no] port-role
                    — [no] port-state
                    — [no] sap sap-id
                    — [no] sdp sdp-id:vc-id

debug
    — router
        — igmp
            — [no] interface [ip-int-name | ip-address]
            — [no] mcs [ip-int-name]
            — [no] mcs
            — [no] packet [query | v1-report| v2-report | v3-report | v2-leave] [ip-int-name |
              ip-address]
```

3.8.1.4 Tools Commands

```
tools
    — dump
        — service
            — id service-id
                — fdb
                    — card-status
                    — mac-status [mac ieee-address] [card slot-id] [pending]
                — provider-tunnels [type {originating | terminating}]
            — proxy-arp usage
            — proxy-nd usage
            — vpls-fdb-stats id
            — vxlan [clear]
                — dup-vtep-egrvni [clear]
                — dup-vtep-egrvni
    — perform
```

```
— service
    — eval-pw-template policy-id [allow-service-impact]
    — id service-id
        — eval-pw-template policy-id [allow-service-impact]
        — eval-vpls-template
        — eval-vpls-sap-template [sap-id]
        — instantiate-data-saps sap-id
        — provider-tunnels
```

Refer to the *7450 ESS, 7750 SR, and 7950 XRS OAM and Diagnostics Guide* for information about CLI commands and syntax for OAM and diagnostics commands.

3.8.2 Command Descriptions

3.8.2.1 VPLS Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

active-subscribers

Syntax **active-subscribers summary**
active-subscribers [**subscriber** *sub-ident-string* [**sap** *sap-id* **sla-profile** *sla-profile-name*]] [**detail**]
active-subscribers hierarchy [**subscriber** *sub-ident-string*]

Context show>service

Description This command displays active subscriber information.

Parameters **sap** *sap-id* — Displays SAP information for the specified SAP ID.

sla-profile *sla-profile-name* — Displays information for the specified SLA profile.

summary — Displays active subscriber information in a brief format.

subscriber *sub-ident-string* — Displays information for the specified subscriber.

hierarchy — Displays the subscriber hierarchy.

detail — Displays detailed output.

Output The following output displays an example of service active subscriber information.

Sample Output

```
A:Dut-A# show service active-subscribers summary
===============================================================================
Active Subscriber table summary
===============================================================================
Total Count : 2
===============================================================================
A:Dut-A#
A:Dut-A# show service active-subscribers hierarchy
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- nokia_110 (sub_default)
   |
   |-- sap:1/1/20:841 - sla:sla-default
   |   |
   |   +-- 192.168.1.56 - mac:00:00:10:10:12:13 - DHCP - svc:1000
   |
   +-- sap:1/1/20:842 - sla:sla-default
       |
       +-- 192.168.1.55 - mac:00:00:10:10:12:11 - DHCP - svc:1000
-- nokia_112 (sub_default)
   |
   +-- sap:1/2/1:112 - sla:sla_default
       |
       +-- 192.168.1.12 - mac:00:00:00:00:00:01 - STATIC - svc:1000
-------------------------------------------------------------------------------
Number of active subscribers : 2
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================
A:Dut-A#
A:Dut-A# show service active-subscribers subscriber nokia_110 hierarchy
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- nokia_110 (sub_default)
   |
   |-- sap:1/1/20:841 - sla:sla-default
   |   |
   |   +-- 192.168.1.56 - mac:00:00:10:10:12:13 - DHCP - svc:1000
   |
   +-- sap:1/1/20:842 - sla:sla-default
       |
       +-- 192.168.1.55 - mac:00:00:10:10:12:11 - DHCP - svc:1000
-------------------------------------------------------------------------------
Number of active subscribers : 2
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================
A:Dut-A#
A:Dut-A# show service active-subscribers subscriber nokia_110
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber nokia_110-2 (sub-default)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/20:841 - sla:sla-default
-------------------------------------------------------------------------------
IP Address
```

```
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
192.168.1.56
                00:00:10:10:12:13    N/A            DHCP         1000       Y
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(2) SLA Profile Instance sap:1/1/20:842 - sla:sla-default
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
192.168.1.55
                00:00:10:10:12:11    N/A            DHCP         1000       Y
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
A:Dut-A#
A:Dut-A# show service active-subscribers subscriber nokia_110 sap 1/2/
1:100 slaprofile
sla_default
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber nokia_110 (sub-default)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/20:841 - sla:sla-default
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
192.168.1.56
                00:00:10:10:12:13    N/A            DHCP         1000       Y
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
A:Dut-A#
A:Dut-A# show service active-subscribers subscriber nokia_110 sap 1/2/
1:100 slaprofile
sla_default detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber nokia_110-2 (sub-default)
-------------------------------------------------------------------------------
I. Sched. Policy : N/A
E. Sched. Policy : N/A                              E. Agg Rate Limit: 10
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy     : N/A                              Collect Stats    : Disabled
ANCP Pol.        : N/A
HostTrk Pol.     : N/A
IGMP Policy      : igmp-policy-01
MLD Policy       : N/A
PIM Policy       : N/A
Sub. MCAC Policy : N/A
NAT Policy       : N/A
UPnP Policy      : N/A
```

```
NAT Prefix List  : N/A
Def. Encap Offset: none                              Encap Offset Mode: none
Avg Frame Size   : N/A
Vol stats type   : full
Preference       : 5
LAG hash class   : 1
LAG hash weight  : 1
Sub. ANCP-String : "nokia_110"
Sub. Int Dest Id : ""
Igmp Rate Adj    : N/A
RADIUS Rate-Limit: N/A
Oper-Rate-Limit  : 10
-------------------------------------------------------------------------------
Radius Accounting
-------------------------------------------------------------------------------
Policy           : acct-01
Session Opti.Stop: False
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:1/1/20:841 (VPLS 1000)
    - sla:sla-default
-------------------------------------------------------------------------------
Description          : (Not Specified)
Host Limits          : No Limit
Egr Sched-Policy     : N/A
Ingress Qos-Policy   : 1                      Egress Qos-Policy : 20
Ingress Queuing Type : Service-queuing (Not Applicable to Policer)
Ingr IP Fltr-Id      : N/A                    Egr IP Fltr-Id    : N/A
Ingr IPv6 Fltr-Id    : N/A                    Egr IPv6 Fltr-Id  : N/A
Ingress Report-Rate  : Maximum
Egress Report-Rate   : Maximum
Egress Remarking     : from Sap Qos
Credit Control Pol.  : N/A
Category Map         : (Not Specified)
Use ing L2TP DSCP    : false
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
192.168.1.56
                00:00:10:10:12:13    N/A            DHCP         1000       Y
-------------------------------------------------------------------------------
-------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------
                       Packets                Octets
Off. HiPrio          : 0                      0
Off. LowPrio         : 0                      0
Off. Uncolor         : 0                      0
Off. Managed         : 0                      0
Queueing Stats (Ingress QoS Policy 1)
Dro. HiPrio          : 0                      0
Dro. LowPrio         : 0                      0
For. InProf          : 0                      0
For. OutProf         : 0                      0
Queueing Stats (Egress QoS Policy 20)
Dro. In/InplusProf          : 0                       0
```

```
Dro. Out/ExcProf          : 0                        0
For. In/InplusProf           : 0                        0
For. Out/ExcProf          : 0                        0
-------------------------------------------------------------------------------
SLA Profile Instance per Queue statistics
-------------------------------------------------------------------------------
                         Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio               : 0                        0
Off. LowPrio              : 0                        0
Dro. HiPrio               : 0                        0
Dro. LowPrio              : 0                        0
For. InProf               : 0                        0
For. OutProf              : 0                        0
Egress Queue 1
Dro. In/InplusProf           : 0                        0
Dro. Out/ExcProf          : 0                        0
For. In/InplusProf           : 0                        0
For. Out/ExcProf          : 0                        0
-------------------------------------------------------------------------------
SLA Profile Instance per Policer statistics
-------------------------------------------------------------------------------
Egress Policer 1 (Stats mode: minimal)
Off. All                  : 0                        0
Dro. All                  : 0                        0
For. All                  : 0                        0
-------------------------------------------------------------------------------
```

egress-label

Syntax **egress-label** *egress-label1* [*egress-label2*]

Context show>service

Description This command displays service information using the range of egress labels.

If only the mandatory *egress-label1* parameter is specified, only services using the specified label are displayed.

If both *egress-label1* and *egress-label2* parameters are specified, the services using the range of labels X where *egress-label1* <= X <= *egress-label2* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *egress-label1* — The starting egress label value for the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0, 2049 to 131071

egress-label2 — The ending egress label value for the label range.

Default The *egress-label1* value.

Values 2049 to 131071

Output The following output displays an example of service egress label information.

Sample Output

```
*A:ALA-12# show service egress-label 0 10000
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          100:1              Mesh 0            0
...
1          107:1              Mesh 0            0
1          108:1              Mesh 0            0
1          300:1              Mesh 0            0
1          301:1              Mesh 0            0
1          302:1              Mesh 0            0
1          400:1              Mesh 0            0
1          500:2              Spok 131070       2001
1          501:1              Mesh 131069       2000
100        300:100            Spok 0            0
200        301:200            Spok 0            0
300        302:300            Spok 0            0
400        400:400            Spok 0            0
-------------------------------------------------------------------------------
Number of Bindings Found : 23
===============================================================================
*A:ALA-12#
```

Table 44 describes show service egress label output fields.

Table 44 **Show Service Egress Command Output Fields**

Label	Description
Svc Id	The ID that identifies a service.
Sdp Id	The ID that identifies an SDP.
Type	Indicates whether the SDP binding is a spoke or a mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.
Number of bindings found	The total number of SDP bindings that exist within the specified egress label range.

fdb-info

Syntax **fdb-info**

Context show>service

Description Displays global FDB usage information.

Output The following output displays an example of service FDB information.

Sample Output

```
*A:PE1# show service fdb-info
===============================================================================
Forwarding Database(FDB) Information
===============================================================================
Service Id        : 1                   Mac Move          : Disabled
Primary Factor    : 3                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                 Allocated Count   : 3
Total In Use      : 3
Learned Count     : 2                   Static Count      : 0
OAM MAC Count     : 0                   DHCP MAC Count    : 0
Host MAC Count    : 0                   Intf MAC Count    : 0
Spb Count         : 0                   Cond MAC Count    : 0
BGP EVPN Count    : 0                   EVPN Static Cnt   : 2
EVPN Dup Det Cnt  : 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
Mac Subnet Len    : 48
Sel Learned FDB   : Disabled
Service Id        : 2                   Mac Move          : Disabled
Primary Factor    : 3                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                 Allocated Count   : 2
Total In Use      : 2
Learned Count     : 4                   Static Count      : 0
OAM MAC Count     : 0                   DHCP MAC Count    : 0
Host MAC Count    : 0                   Intf MAC Count    : 0
Spb Count         : 0                   Cond MAC Count    : 0
BGP EVPN Count    : 0                   EVPN Static Cnt   : 0
EVPN Dup Det Cnt  : 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
Mac Subnet Len    : 48
Sel Learned FDB   : Disabled
-------------------------------------------------------------------------------
Total Service FDBs : 2
Total FDB Configured Size : 500
Total FDB Entries In Use  : 5
PBB MAC Address Indices In Use  : 0
```

```
-------------------------------------------------------------------------------
===============================================================================
*A:PE1#
```

Table 45 describes show FDB-Info command output.

Table 45 **Show FDB Information Fields**

Label	Description
ServID	Displays the service ID.
MAC	Displays the associated MAC address.
Mac Move	Displays the administrative state of the MAC movement feature associated with this service.
Primary Factor	Displays a factor for the primary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Secondary Factor	Displays a factor for the secondary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Mac Move Rate	Displays the maximum rate at which MACs can be re-learned in this service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MAs. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: for example, the default rate of 2 re-learns per second corresponds to 10 re-learns in a 5 second period.
Mac Move Timeout	Displays the time in seconds to wait before a SAP that has been disabled after exceeding the maximum re-learn rate is re-enabled. A value of zero indicates that the SAP will not be automatically re-enabled after being disabled. If after the SAP is re-enabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.
Mac Move Retries	Displays the number of times retries are performed for re-enabling the SAP/SDP.
Table Size	Specifies the maximum number of learned and static entries allowed in the FDB of this service.
Allocated Count	Displays the total number of allocated entries in the FDB of this service.
Total In Use	Displays the total number of entries in use in the FDB of this service.

Table 45 **Show FDB Information Fields (Continued)**

Label	Description (Continued)
Learned Count	Displays the current number of learned entries in the FDB of this service.
Static Count	Displays the current number of static entries in the FDB of this service.
OAM MAC Count	Displays the current number of OAM entries in the FDB of this service.
DHCP MAC Count	Displays the current number of DHCP-learned entries in the FDB of this service.
Host MAC Count	Displays the current number of host-learned entries in the FDB of this service.
Intf MAC Count	Displays the total number of interface MAC entries in the FDB of this service.
SPB Count	Displays the total number of SPB entries in the FDB of this service.
Cond MAC Count	Displays the total number of conditional static MAC entries in the FDB of this service.
BGP EVPN Count	Displays the total number of BGP EVPN entries in the FDB of this service.
EVPN Static Cnt	Displays the total number of BGP EVPN MAC entries with the sticky bit set in the FDB of this service.
EVPN Dup Det Cnt	Displays the total number of times a BGP EVPN duplicate MAC address has been detected in this service.
Remote Age	Displays the number of seconds used to age out FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	Displays the number of seconds used to age out FDB entries learned on local SAPs.
High Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be raised by the agent.
Low Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be cleared by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded.

Table 45 **Show FDB Information Fields (Continued)**

Label	Description (Continued)
Mac Aging	Indicates whether the MAC aging process is enabled.
Relearn Only	When one of the FDB table size limits (service, line card, system) has been reached, the learning of new MAC addresses is temporary disabled and only MAC relearns are allowed. When in this state, the Relearn Only flag is True, otherwise it is False.
Mac Subnet Len	Displays the number of bits to be considered when performing MAC-learning or MAC-switching.
Source-Identifier	The location where the MAC is defined.
Type/Age	Type — Specifies the number of seconds used to age out TLS FDB entries learned on local SAPs. Age — Specifies the number of seconds used to age out TLS FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs. L — Learned - Dynamic entries created by the learning process. OAM — Entries created by the OAM process. H — Host, the entry added by the system for a static configured subscriber host. D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease. P — Indicates the MAC is protected by the MAC protection feature. Static — Statically configured.
Last Change	Indicates the time of the most recent state changes.
Sel Learned FDB	Displays the administrative state of the selective learned FDB feature associated with this service.

fdb-mac

Syntax **fdb-mac** *ieee-address* [*expiry*]

Context show>service

Description This command displays the FDB entry for a specified MAC address.

Parameters *ieee-address* — The 48-bit MAC address for which the FDB entry will be displayed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

expiry — Shows the time until the MAC is aged out.

Output The following output displays an example of FDB MAC information.

Sample Output

```
*A:ian2# show service fdb-mac
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier        Type   Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:00:00:00:00:01 sap:1/1/1                LP/0   01/07/2011 20:25:34
1         00:00:00:00:00:02 sap:1/1/2                L/0    01/07/2011 20:26:25
1         00:00:00:00:00:03 sap:1/1/1                A/0    01/07/2011 20:25:34
-------------------------------------------------------------------------------
No. of Entries: 2
-------------------------------------------------------------------------------
Legend: L=Learned; P=MAC is protected; A=Auto learn protected
===============================================================================
*A:ian2#
```

The following shows the protected MACs in the FDB.

```
A:term17>config>service>vpls>sap>arp-host# show service id  12 fdb detail

===============================================================================
Forwarding Database, Service 12
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
12        00:00:07:00:00:00 sdp:8:1                  LP/0     10/03/11 10:46:00
12        00:00:07:00:00:01 sdp:8:1                  LP/0     10/03/11 10:46:00
12        00:00:07:00:00:62 sdp:8:1                  LP/0     10/03/11 10:46:01
12        00:00:07:00:00:63 sdp:8:1                  LP/0     10/03/11 10:46:01
12        00:11:11:11:11:11 sap:lag-100:12           Static:P 10/03/11 09:42:02
12        00:11:11:11:11:22 sap:lag-1:123            Static   10/03/11 09:42:02
12        00:11:11:11:11:33 sdp:8:1                  Static:P 10/03/11 09:42:02
12        00:11:11:11:11:44 sap:2/1/3:13             Static   10/03/11 09:42:02
12        00:11:11:11:11:55 a(8:80)                  Static   10/03/11 09:42:02
12        00:11:11:11:11:66 sdp:8:10                 Static   10/03/11 09:42:02
12        00:11:11:11:11:77 sap:2/1/3:15             Static   10/03/11 09:42:02
12        00:11:11:11:11:88 sap:2/1/3:14             Static   10/03/11 09:42:02
12        76:1e:ff:00:00:b2 cpm                      Host     10/03/11 09:42:02
-------------------------------------------------------------------------------
No. of MAC Entries: 109
```

The following example shows whether restrict-protected-src is enabled on an SDP.

```
*A:PE# show service id 1 sdp 1:1 detail
===============================================================================
Service Destination Point (Sdp Id : 1:1) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:1  -(1.1.1.2)
```

```
-------------------------------------------------------------------------------
…
RestMacProtSrc Act : SDP-oper-down
```

Table 46 describes the show FDB-MAC command output fields.

Table 46 **Show FDB-MAC Command Output Fields**

Label	Description
Service ID	The service ID number.
MAC	The specified MAC address.
Source-Identifier	The location where the MAC is defined.
Type/Age	Static — FDB entries created by management.
	Learned — Dynamic entries created by the learning process.
	OAM — Entries created by the OAM process.
	H — Host, the entry added by the system for a static configured subscriber host.
	D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease.
	P — Indicates the MAC is protected by the MAC protection feature.

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description Display services using the range of ingress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting ingress label value for the label range. If only *start-label* is specified, services only using *start-label* are displayed.

Values 0, 2048 to 131071

end-label — The ending ingress label value for the label range

Default The *start-label* value.

Values 2049 to 131071

Output The following output displays an example of service ingress label information.

Sample Output

```
*A:ALA-12# show service ingress-label 0
==============================================================================
Martini Service Labels
==============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          50:1               Mesh 0            0
1          100:1              Mesh 0            0
1          101:1              Mesh 0            0
1          102:1              Mesh 0            0
1          103:1              Mesh 0            0
1          104:1              Mesh 0            0
1          105:1              Mesh 0            0
1          106:1              Mesh 0            0
1          107:1              Mesh 0            0
1          108:1              Mesh 0            0
1          300:1              Mesh 0            0
1          301:1              Mesh 0            0
1          302:1              Mesh 0            0
1          400:1              Mesh 0            0
100        300:100            Spok 0            0
200        301:200            Spok 0            0
300        302:300            Spok 0            0
400        400:400            Spok 0            0
------------------------------------------------------------------------------
Number of Bindings Found : 21
------------------------------------------------------------------------------
*A:ALA-12#
```

Table 47 describes show service ingress-label output fields.

Table 47 **Show Service Ingress-Label Fields**

Label	Description
Svc ID	The service identifier.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is spoke or mesh.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.

Table 47 **Show Service Ingress-Label Fields (Continued)**

Label	Description (Continued)
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of Bindings Found	The number of SDP bindings within the label range specified.

sap-using

Syntax **sap-using** [**msap**] [**dyn-script**] [**description**]
sap-using [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**]
sap-using app-profile *app-profile-name*
sap-using authentication-policy *policy-name* [**msap**]
sap-using encap-type *encap-type*
sap-using eth-cfm collect-lmm-stats [**sap** *sap-id*]
sap-using eth-ring [**ring-id** *eth-ring-id*]
sap-using eth-tunnel [**tunnel-id** *eth-tunnel-id*]
sap-using {**ingress** | **egress**} **atm-td-profile** *td-profile-id*
sap-using {**ingress** | **egress**} **filter** *filter-id*
sap-using {**ingress** | **egress**} **qos-policy** *qos-policy-id* [**msap**]
sap-using interface {*ip-address* | *ip-int-name* [**msap**]
sap-using mc-ring peer *ip-address* **ring** *sync-tag*
sap-using process-cpm-traffic-on-sap-down
sap-using etree

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress** — Specifies matching an ingress policy

egress — Specifies matching an egress policy

qos-policy-id — The ingress or egress QoS Policy ID for which the matching SAPs will be displayed. This parameter applies to the 7450 ESS or 7750 SR only.

Values 1 to 65535

td-profile-id — Displays SAPs using this traffic description. This parameters applies to the 7750 SR only.

filter-id — The ingress or egress filter policy ID for which matching SAPs will be displayed

Values 1 to 65535

dyn-script — Displays dynamic service SAPs information

auth-plcy-name — The session authentication policy for which the matching SAPs will be displayed. This parameter applies to the 7450 ESS or 7750 SR only.

sap-id — Specifies the physical port identifier portion of the SAP definition

interface — Specifies matching SAPs with the specified IP interface. This parameter applies to the 7450 ESS or 7750 SR only.

ip-address — The IP address of the interface for which the matching SAPs will be displayed. This parameter applies to the 7450 ESS or 7750 SR only.

Values 1.0.0.0 to 223.255.255.255

ip-int-name — The IP interface name for which the matching SAPs. will be displayed. This command applies to the 7450 ESS or 7750 SR only.

etree — Specifies matching of SAPs configured as E-Tree SAPs and the corresponding role in the E-Tree services: Leaf-AC, Root-AC or Root-leaf-tag SAPs. SAPs listed as Root-leaf-tag "Disabled" and Leaf-Ac "Disabled" function as Root-AC SAPs.

Output The following output displays an example of services associated so particular SAPs.

Sample Output

```
A:ALA-701# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                 SvcId      Ing.   Ing.    Egr.   Egr.    Anti    Adm  Opr
                                  QoS    Fltr    QoS    Fltr    Spoof
-------------------------------------------------------------------------------
1/1/3                  10203041   1      ip4     1      none    none    Up   Up
1/1/4                  10203042   1      none    1      ip4     none    Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 2
-------------------------------------------------------------------------------
A:ALA-701#

show service sap-using process-cpm-traffic-on-sap-down
===============================================================================
SAP Ignore Sap Lag Down Information
===============================================================================
SAP                                    Svc Id              Ignore SapLag Down
-------------------------------------------------------------------------------
lag-1:1100.*                           1100                enabled
-------------------------------------------------------------------------------
Number of lag saps: 1
===============================================================================
```

Sample Output

The following is sample output for VPLS E-Tree configured SAPs.

```
*A:DutA# show service sap-using etree
===============================================================================
Etree SAP Information
===============================================================================
Svc Id      SAP                                Leaf-Tag    Root-    Leaf-Ac
                                                           leaf-tag
-------------------------------------------------------------------------------
2005        1/1/1:2005                         0           Disabled Enabled
2005        1/1/7:2006.200                     2007        Enabled  N/A
2005        1/1/7:0.*                          0           Disabled Disabled
2005        1/1/7:2005.*                       0           Disabled Disabled
2005        1/1/8:1                            0           Disabled Disabled
-------------------------------------------------------------------------------
Number of etree saps: 5
===============================================================================
```

Table 48 describes show service SAP output fields.

Table 48 **Show Service SAP Fields**

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
SapMTU	The SAP MTU value.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
Egr. Fltr	The filter policy ID applied to the egress SAP.
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The administrative state of the SAP.
Opr	The actual state of the SAP.
E-Tree SAP Information	
Svc ID	The service identifier.
SAP	The root SAP including the outer tag used by the root frames.
Leaf-Tag	The outer tag used by the leaf frames on the referred SAP.
Root-Leaf-Tag	The state of the root leaf tag SAPs.
Leaf-AC	The state of the leaf AC SAPs.

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]
sdp [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]
sdp [*sdp-id* | **far-end** *ip-addr*] [**detail** | **keep-alive-history**]

Context show>service>id

Description This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Displays only information for the specified SDP ID. An SDP is a logical mechanism that ties a far-end 7450 ESS or 7750 SR to a particular service without having to specifically define far end SAPs. Each SDP represents a method to reach a router.

Default All SDPs

Values 1 to 17407

far-end *ip-addr* — Displays only SDPs matching with the specified system IP address of the far-end destination 7450 ESS or 7750 SR OS for the Service Distribution Point (SDP) that is the termination point for a service.

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

Output The following output displays an example of service MAC protect information.

Sample Output

```
A:ALA-48# show service id <service-id> mac-protect
===============================================================================
Mac Protection
===============================================================================
ServId    MAC
-------------------------------------------------------------------------------
1         aa:aa:aa:aa:aa:ab
-------------------------------------------------------------------------------
No. of MAC Entries: 1
===============================================================================
```

Table 49 describes show service-id SDP output fields.

Table 49 **Show Service-ID SDP Fields**

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.

Table 49 **Show Service-ID SDP Fields (Continued)**

Label	Description (Continued)
Split Horizon Group	Name of the split horizon group where the SDP belongs.
VC Type	Displays the VC type, ether or vlan.
VC Tag	Displays the explicit dot1q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.

Table 49 **Show Service-ID SDP Fields (Continued)**

Label	Description (Continued)
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
E. Fwd. Octets	Specifies the number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS.

sdp-using

Syntax **sdp-using** [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]
sdp-using etree

Context show>service

Description This command displays services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID

Values 1 to 17407

vc-id — Displays virtual circuit identifier information

Values 1 to 4294967295

far-end *ip-address* — Displays only services matching with the specified far-end IP address

Default Services with any far-end IP address.

etree — Specifies matching of SDP bindings configured as E-Tree SDP bi9ndings and the corresponding role in the E-Tree services: Leaf-AC, Root-AC or Root-leaf-tag SDP binds. SDP binds listed as Root-leaf-tag "Disabled" and Leaf-Ac "Disabled" function as Root-AC SDP binds.

Output The following output displays an example of services using specific SDPs.

Sample Output

```
*A:ALA-1# show service sdp-using 300
===============================================================================
Service Destination Point (Sdp Id : 300)
===============================================================================
SvcId       SdpId               Type Far End        Opr State I.Label  E.Label
-------------------------------------------------------------------------------
1           300:1               Mesh 10.0.0.13      Up        131071   131071
2           300:2               Spok 10.0.0.13      Up        131070   131070
100         300:100             Mesh 10.0.0.13      Up        131069   131069
101         300:101             Mesh 10.0.0.13      Up        131068   131068
102         300:102             Mesh 10.0.0.13      Up        131067   131067
-------------------------------------------------------------------------------
Number of SDPs : 5
-------------------------------------------------------------------------------
*A:ALA-1#
```

Sample Output

The following is sample output for VPLS E-Tree configured SDP bindings.

```
*A:DutA# show service sdp-using etree
===============================================================================
Etree SDP-BIND Information
===============================================================================
Svc Id      SDP-BIND Information             Type        Root-    Leaf-Ac
                                                         leaf-tag
-------------------------------------------------------------------------------
2005        12:2005                          Spoke       Enabled  N/A
2005        12:2006                          Spoke       Disabled Enabled
2005        12:2007                          Spoke       Disabled Enabled
-------------------------------------------------------------------------------
Number of etree sdp-binds: 3
===============================================================================
```

Table 50 describes service sdp-using output fields.

Table 50 **Show Service SDP-Using Fields**

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.
Type	Specifies the type of SDP: Spoke or Mesh.
Far End	The far-end address of the SDP.
Oper State	The operational state of the service.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.

Table 50 **Show Service SDP-Using Fields (Continued)**

Label	Description (Continued)
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Etree SDP Bind Information	
Svc ID	The service identifier.
SDP-Bind	The leaf tag SDP bind identifier.
Type	The type SDP bind.
Root-Leaf-Tag	The state of the root leaf tag SDP bind,
Leaf-AC	The state of the leaf AC SDP bind.

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] [**b-vpls**] [**i-vpls**] [**m-vpls**] [**apipe**] [**fpipe**] [**ipipe**] **sdp** *sdp-id*] [**customer** *customer-id*]
service-using etree

Context show>service

Description This command displays the services matching certain usage properties. If no optional parameters are specified, all services defined on the system are displayed.

Parameters **epipe** — Displays matching Epipe services

ies — Displays matching IES instances. This parameter applies to the 7450 ESS or 7750 SR only.

vpls — Displays matching VPLS instances

vprn — Displays matching VPRN services. This parameter applies to the 7750 SR only.

mirror — Displays matching mirror services

b-vpls — Displays matching B-VPLS services. This parameter applies to the 7450 ESS or 7750 SR only.

i-vpls — Displays matching I-VPLS services. This parameter applies to the 7450 ESS or 7750 SR only.

apipe — Displays matching Apipe services. This parameter applies to the 7750 SR only

fpipe — Displays matching Fpipe services. This parameter applies to the 7750 SR only

ipipe — Displays matching Ipipe services. This parameter applies to the 7450 ESS or 7750 SR only

sdp-id — Displays only services bound to the specified SDP ID. This parameter applies to the 7450 ESS or 7750 SR only.

Default Services bound to any SDP ID

Values 1 to 17407

customer-id — Displays services only associated with the specified customer ID

Default Services associated with a customer

Values 1 to 2147483647

etree — Specifies matching of all VPLS services configured as E-Tree.

Output The following output displays an example of services using certain options.

Sample Output

```
*A:ALA-12# show service service-using customer 10
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            VPLS      Up     Up         10                09/05/2006 13:24:15
100          IES       Up     Up         10                09/05/2006 13:24:15
300          Epipe     Up     Up         10                09/05/2006 13:24:15
-------------------------------------------------------------------------------
Matching Services : 3
===============================================================================
*A:ALA-12#


*A:ALA-12# show service service-using epipe
===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
6            Epipe     Up     Up         6                 09/22/2006 23:05:58
7            Epipe     Up     Up         6                 09/22/2006 23:05:58
8            Epipe     Up     Up         3                 09/22/2006 23:05:58
103          Epipe     Up     Up         6                 09/22/2006 23:05:58
-------------------------------------------------------------------------------
Matching Services : 4
===============================================================================
*A:ALA-12#


*A:ALA-14# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
```

```
101          mVPLS     Up     Up         1                  10/26/2006 15:44:57
102          mVPLS     Up     Up         1                  10/26/2006 15:44:57
-------------------------------------------------------------------------------
Matching Services : 5
-------------------------------------------------------------------------------
*A:ALA-14#


*A:SetupCLI# show service service-using
  - service-using [epipe] [ies] [vpls] [mirror] [ipipe] [b-vpls] [i-vpls]
[m-vpls] [sdp <sdp-id>] [customer <customer-id>]

 <epipe>              : keyword - displays epipe services
 <ies>                : keyword - displays ies services
 <vpls>               : keyword - displays vpls services
 <mirror>             : keyword - displays mirror services
 <ipipe>              : keyword - displays ipipe services
 <sdp-id>             : [1..17407] - display services using this sdp
 <customer-id>        : [1..2147483647] - display services using this customer
 <b-vpls>             : keyword - displays b-vpls services
 <i-vpls>             : keyword - displays i-vpls services
 <m-vpls>             : keyword - displays m-vpls services

*A:SetupCLI# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
23           mVPLS     Up     Down       2                 09/25/2007 21:45:58
100          Epipe     Up     Down       2                 09/25/2007 21:45:58
101          Epipe     Up     Down       2                 09/25/2007 21:45:58
102          Epipe     Up     Down       2                 09/25/2007 21:45:58
105          Epipe     Up     Down       2                 09/25/2007 21:45:58
110          Epipe     Up     Down       1                 09/25/2007 21:45:58
990          IES       Up     Down       1                 09/25/2007 21:45:58
1000         Mirror    Up     Down       1                 09/25/2007 21:45:59
1001         Epipe     Up     Down       1                 09/25/2007 21:45:58
1002         Epipe     Up     Down       1                 09/25/2007 21:45:58
1003         Epipe     Up     Down       1                 09/25/2007 21:45:58
1004         Epipe     Up     Down       1                 09/25/2007 21:45:58
2000         Mirror    Up     Down       1                 09/25/2007 21:45:59
2001         i-VPLS    Up     Down       1                 09/25/2007 21:45:59
2002         b-VPLS    Up     Down       1                 09/25/2007 21:45:59
2003         i-VPLS    Down   Down       1                 09/25/2007 21:45:59
2004         b-mVPLS   Down   Down       1                 09/25/2007 21:45:59
2005         i-mVPLS   Down   Down       1                 09/25/2007 21:45:59
8787         IES       Up     Down       2                 09/25/2007 21:45:58
8888         IES       Up     Down       1                 09/25/2007 21:45:58
10000        IES       Down   Down       1                 09/25/2007 21:45:59
10001        VPLS      Up     Down       1                 09/25/2007 21:45:58
483000       Ipipe     Down   Down       2                 09/25/2007 21:45:59
483001       Ipipe     Up     Down       2                 09/25/2007 21:45:59
483004       Ipipe     Down   Down       2                 09/25/2007 21:45:59
483007       VPLS      Down   Down       2                 09/25/2007 21:45:59
483010       Ipipe     Down   Down       1                 09/25/2007 21:45:59
...
-------------------------------------------------------------------------------
Matching Services : 27
```

```
-------------------------------------------------------------------------------
*A:SetupCLI#


*A:SetupCLI# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
23           mVPLS     Up     Down       2                 09/25/2007 21:45:58
100          Epipe     Up     Down       2                 09/25/2007 21:45:58
101          Epipe     Up     Down       2                 09/25/2007 21:45:58
102          Epipe     Up     Down       2                 09/25/2007 21:45:58
105          Epipe     Up     Down       2                 09/25/2007 21:45:58
110          Epipe     Up     Down       1                 09/25/2007 21:45:58
990          IES       Up     Down       1                 09/25/2007 21:45:58
1000         Mirror    Up     Down       1                 09/25/2007 21:45:59
1001         Epipe     Up     Down       1                 09/25/2007 21:45:58
1002         Epipe     Up     Down       1                 09/25/2007 21:45:58
1003         Epipe     Up     Down       1                 09/25/2007 21:45:58
1004         Epipe     Up     Down       1                 09/25/2007 21:45:58
2000         Mirror    Up     Down       1                 09/25/2007 21:45:59
2001         i-VPLS    Up     Down       1                 09/25/2007 21:45:59
2002         b-VPLS    Up     Down       1                 09/25/2007 21:45:59
2003         i-VPLS    Down   Down       1                 09/25/2007 21:45:59
2004         b-mVPLS   Down   Down       1                 09/25/2007 21:45:59
2005         i-mVPLS   Down   Down       1                 09/25/2007 21:45:59
8787         IES       Up     Down       2                 09/25/2007 21:45:58
8888         IES       Up     Down       1                 09/25/2007 21:45:58
10000        IES       Down   Down       1                 09/25/2007 21:45:59
10001        VPLS      Up     Down       1                 09/25/2007 21:45:58
483000       Ipipe     Down   Down       2                 09/25/2007 21:45:59
483001       Ipipe     Up     Down       2                 09/25/2007 21:45:59
483004       Ipipe     Down   Down       2                 09/25/2007 21:45:59
483007       VPLS      Down   Down       2                 09/25/2007 21:45:59
483010       Ipipe     Down   Down       1                 09/25/2007 21:45:59
...
-------------------------------------------------------------------------------
Matching Services : 27
-------------------------------------------------------------------------------
*A:SetupCLI#

*A:term17>config>service>epipe# show service id 2000 epipe
===============================================================================
Related Epipe services for bVpls service 2000
===============================================================================
Epipe SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
1                   1                   Down                Down
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
*A:term17>config>service>epipe#
```

The following sample outputs show VPLS Services configured as E-Tree.

```
*A:DutA# show service service-using
```

```
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            VPLS      Up   Up   1          evpn-vxlan-1
2            VPRN      Up   Up   1
2005         VPLS-Etr* Up   Up   1
2006         VPRN      Up   Up   1
2147483648   IES       Up   Down 1          _tmnx_InternalIesService
2147483649   intVpls   Up   Down 1          _tmnx_InternalVplsService
-------------------------------------------------------------------------------
Matching Services : 6
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.


*A:DutA# show service service-using etree
===============================================================================
Services [etree]
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
2005         VPLS-Etr* Up   Up   1
-------------------------------------------------------------------------------
Matching Services : 1
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Table 51 describes show service service-using output fields.

Table 51 **Show Service Service-Using Fields**

Label	Description
Service Id	The service identifier.
Type	Specifies the service type configured for the service ID including VPLS, VPRN, VPLS-ETR, VPRN, IES and INTVPLS
Adm	The administrative state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.

subscriber-using

Syntax **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address*[*/mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]

Context show>service>subscriber-using

Description This command displays subscribers using specified options.

Parameters *service-id* — Displays subscriber information about the specified service ID

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters in length

sap-id — Specifies the physical port identifier portion of the SAP definition

ip-int-name — Displays subscriber information about the specified interface

ip-address[/mask — Displays subscriber information about the specified IP address

ieee-address — Displays subscriber information about the specified MAC address

sub-profile-name — Displays subscriber information about the specified subscriber profile name

sla-profile-name — Displays subscriber information about the specified SLA profile name

id

Syntax **id** *service-id*

Context show>service

Description This command displays information for a particular service ID.

Parameters *service-id* — Displays information for the unique service identification number that identifies the service in the service domain.

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters in length

all — Displays detailed information about the service

arp — Displays ARP entries for the service. This parameter applies to the 7450 ESS or 7750 SR only.

authentication — Displays subscriber authentication information. This parameter applies to the 7450 ESS or 7750 SR only.

base — Displays basic service information

dhcp — Displays DHCP information. This parameter applies to the 7450 ESS or 7750 SR only.

endpoint — Displays service endpoint information

epipe — Displays Epipe services associated with the B-VPLS service. This parameter applies to the 7450 ESS or 7750 SR only.

fdb — Displays FDB entries

gsmp — Displays GSMP information. This parameter applies to the 7450 ESS or 7750 SR only.

host — Displays static hosts configured on the service. This parameter applies to the 7450 ESS or 7750 SR only.

i-vpls — Displays I-VPLS services associated with the B-VPLS. This parameter applies to the 7450 ESS or 7750 SR only.

igmp-snooping — Displays IGMP snooping information. This parameter applies to the 7450 ESS or 7750 SR only.

interface — Displays service interfaces. This parameter applies to the 7450 ESS or 7750 SR only.

l2-route-table — Displays Layer 2 route information associated with the service. This parameter applies to the 7450 ESS or 7750 SR only.

l2pt — Displays L2PT information of SAPs and Spokes. This parameter applies to the 7450 ESS or 7750 SR only.

labels — Displays labels being used by this service

mac-move — Displays Mac move-related information about the service. This parameter applies to the 7450 ESS or 7750 SR only.

mac-protect — Displays MAC protect information. This parameter applies to the 7450 ESS or 7750 SR only.

mfib — Displays MFIB related information. This parameter applies to the 7450 ESS or 7750 SR only.

mld-snooping — Displays MLD snooping information. This parameter applies to the 7450 ESS or 7750 SR only.

mmrp — Displays MMRP information. This parameter applies to the 7450 ESS or 7750 SR only.

mrp — Displays MRP information. This parameter applies to the 7450 ESS or 7750 SR only.

msap — Displays MSAPs associated to the service. This parameter applies to the 7450 ESS or 7750 SR only.

pim-snooping — Displays PIM snooping information. This parameter applies to the 7450 ESS or 7750 SR only.

pppoe — Displays PPPoE information. This parameter applies to the 7450 ESS or 7750 SR only.

retailers — Displays service retailer information. This parameter applies to the 7450 ESS or 7750 SR only.

sap — Displays SAPs associated to the service

sdp — Displays SDPs associated with the service

source-address — Displays source-address configured for applications. This parameter applies to the 7450 ESS or 7750 SR only.

split-horizon-group — Displays split horizon group information. This parameter applies to the 7450 ESS or 7750 SR only.

stp — Displays STP information

subscriber-host — Displays subscriber host information. This parameter applies to the 7450 ESS or 7750 SR only.

wholesalers — Displays service wholesaler information. This parameter applies to the 7450 ESS or 7750 SR only.

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output The following output displays an example a service displaying all information.

Sample Output

```
*A:ALA-48# show service id 700 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 700                 Vpn Id            : 0
Service Type      : VPLS
Description       : IMA VPLS
Customer Id       : 7
Last Status Change: 02/02/2009 09:27:55
Last Mgmt Change  : 02/02/2009 09:27:57
Admin State       : Up                  Oper State        : Down
MTU               : 1514                Def. Mesh VC Id   : 700
SAP Count         : 1                   SDP Bind Count    : 2
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
-------------------------------------------------------------------------------
BGP Auto-discovery Information
-------------------------------------------------------------------------------
Admin State       : Down                Vpls Id           : None
Route Dist        : None                Prefix            : 10.10.10.103
Rte-Target Import : None                Rte-Target Export : None
Vsi-Import        : None
Vsi-Export        : None
PW-Template Id    : None
-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
Split Horizon Group : DSL-group1
-------------------------------------------------------------------------------
Description       : (Not Specified)
```

```
Instance Id       : 1                   Last Change       : 02/02/2009 09:27:57
-------------------------------------------------------------------------------
Split Horizon Group : SHG_test
-------------------------------------------------------------------------------
Description       : test
Instance Id       : 2                   Last Change       : 02/02/2009 09:27:57
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
 Sdp Id 2:222  -(10.10.10.104)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.104
SDP Id            : 2:222                 Type              : Spoke
Split Horiz Grp   : (Not Specified)
VC Type           : Ether                 VC Tag            : n/a
Admin Path MTU    : 0                     Oper Path MTU     : 0
Far End           : 10.10.10.104          Delivery          : GRE

Admin State       : Up                    Oper State        : Down
Acct. Pol         : None                  Collect Stats     : Disabled
Ingress Label     : 0                     Egress Label      : 0
Ing mac Fltr      : n/a                   Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                   Egr ip Fltr       : n/a
Ing ipv6 Fltr     : n/a                   Egr ipv6 Fltr     : n/a
Admin ControlWord : Not Preferred         Oper ControlWord  : False
Last Status Change: 02/02/2009 09:27:55   Signaling         : TLDP
Last Mgmt Change  : 02/02/2009 09:27:57   Force Vlan-Vc     : Disabled
Endpoint          : N/A                   Precedence        : 4
Class Fwding State: Down
Flags             : SdpOperDown
                    NoIngVCLabel NoEgrVCLabel
                    PathMTUTooSmall
Time to RetryReset: never                 Retries Left      : 3
Mac Move          : Blockable             Blockable Level   : Tertiary
Peer Pw Bits      : None
Peer Fault Ip     : None
Max Nbr of MAC Addr: No Limit             Total MAC Addr    : 0
Learned MAC Addr  : 0                     Static MAC Addr   : 0

MAC Learning      : Enabled               Discard Unkwn Srce: Disabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
Ignore Standby Sig: False                 Block On Mesh Fail: False

KeepAlive Information :
Admin State       : Disabled              Oper State        : Disabled
Hello Time        : 10                    Hello Msg Len     : 0
Max Drop Count    : 3                     Hold Down Time    : 10

Statistics          :
I. Fwd. Pkts.     : 0                     I. Dro. Pkts.     : 0
I. Fwd. Octs.     : 0                     I. Dro. Octs.     : 0
E. Fwd. Pkts.     : 0                     E. Fwd. Octets    : 0
MCAC Policy Name  :
MCAC Max Unconst BW: no limit             MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                    MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                    MCAC Avail Opnl BW: unlimited
```

```
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : Disabled                  Port State        : Discarding
Port Number        : 2049                      Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : False
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : Rstp
Last BPDU from     : N/A
Designated Bridge  : N/A                       Designated Port Id: 0

Fwd Transitions    : 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
-------------------------------------------------------------------------------
 Sdp Id 2:700  -(10.10.10.104)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.104
SDP Id             : 2:700                     Type              : Mesh
Split Horiz Grp    : (Not Specified)
VC Type            : Ether                     VC Tag            : n/a
Admin Path MTU     : 0                         Oper Path MTU     : 0
Far End            : 10.10.10.104              Delivery          : GRE

Admin State        : Up                        Oper State        : Down
Acct. Pol          : None                      Collect Stats     : Disabled
Ingress Label      : 0                         Egress Label      : 0
Ing mac Fltr       : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                       Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                       Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord  : False
Last Status Change : 02/02/2009 09:27:55       Signaling         : TLDP
Last Mgmt Change   : 02/02/2009 09:27:57       Force Vlan-Vc     : Disabled
Endpoint           : N/A                       Precedence        : 4
Class Fwding State : Down
Flags              : SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Peer Pw Bits       : None
Peer Fault Ip      : None
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Disabled                  Oper State        : Disabled
Hello Time         : 10                        Hello Msg Len     : 0
Max Drop Count     : 3                         Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                         I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                         I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                         E. Fwd. Octets    : 0
MCAC Policy Name   :
MCAC Max Unconst BW: no limit                  MCAC Max Mand BW  : no limit
```

```
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
SAP 1/1/9:0
-------------------------------------------------------------------------------
Service Id         : 700
SAP                : 1/1/9:0                  Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Multi Svc Site     : None
Last Status Change : 02/02/2009 09:27:55
Last Mgmt Change   : 02/02/2009 09:27:57
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : 10
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
Q Frame-Based Acct : Disabled
ARP Reply Agent    : Enabled                  Host Conn Verify  : Enabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : Ip                       Avl Static Hosts  : 1
                                              Tot Static Hosts  : 1
Calling-Station-Id : n/a
Application Profile: None

MCAC Policy Name   :                          MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Restr MacProt Src  : Enabled                  Restr MacUnpr Dst : Disabled
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy        : none
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
```

```
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : Disabled                  Port State        : Discarding
Port Number        : 2048                      Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : False
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : Rstp
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                       Designated Port   : N/A

Forward transitions: 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                         MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                         Min Auth Interval : 15 minutes
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 100                       Egress qos-policy : 1
Shared Q plcy      : default                   Multipoint shared : Enabled
I. Sched Pol       : SLA1
E. Sched Pol       : SLA1
-------------------------------------------------------------------------------
Ingress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
Egress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down                      Lease Populate    : 0
DHCP Snooping      : Down                      Action            : Keep

Proxy Admin State  : Down
Proxy Lease Time   : N/A
Emul. Server Addr  : Not Configured
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Down                      MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : None
Def SLA-Profile    : None
Def App-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
 Prof Traffic Only : False
```

```
 Non-Sub-Traffic   : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                          Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0

Queueing Stats(Ingress QoS Policy 100)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                          Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 10 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 12 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 13 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
```

```
...
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Down                     Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Down
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 10:02.90:30:ff:00:00:00  Bridge Instance Id: 2
Bridge Priority    : 4096                     Tx Hold Count     : 5
Topology Change    : Inactive                 Bridge Hello Time : 5
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 25
Top. Change Count  : 0                        Bridge Fwd Delay  : 20
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 20
Rcvd Hello Time    : 5                        Root Max Age      : 25
Root Priority      : 4098                     Root Port         : N/A
-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 700                Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 1
Learned Count     : 0                  Static Count      : 0
OAM-learned Count : 0                  DHCP-learned Count: 0
Host-learned Count: 1
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False
Mac Subnet Len    : 48
-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp               Oper    MRtr Pim  Send    Max Num Max Num MVR       Num
Id                    State   Port Port Queries Grps    Srcs    From-VPLS Grps
-------------------------------------------------------------------------------
sap:1/1/9:0           Down    No   No   No      None    None    Local     0
sdp:2:222             Down    No   No   No      None    None    N/A       0
sdp:2:700             Down    No   No   No      None    None    N/A       0
-------------------------------------------------------------------------------
MLD Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                 Oper    MRtr  Send      Max Num   MVR        Num
Id                      State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
```

```
sap:1/1/9:0              Down    No    Disabled  No Limit  Local      0
sdp:2:222                Down    No    Disabled  No Limit  N/A        0
sdp:2:700                Down    No    Disabled  No Limit  N/A        0
-------------------------------------------------------------------------------
DHCP Summary, service 700
-------------------------------------------------------------------------------
Sap/Sdp                Snoop  Used/       Arp Reply   Info     Admin
                              Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/9:0            No     0/0         Yes         Keep     Down
sdp:2:222              No     N/A         N/A         N/A      N/A
sdp:2:700              No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 3
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
ARP host Summary, service 700
-------------------------------------------------------------------------------
Sap                    Used        Provided    Admin State
-------------------------------------------------------------------------------
sap:1/1/9:0            0           1           outOfService
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:ALA-48#


*A:SetupCLI# show service id 2001 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 2001                Vpn Id            : 0
Service Type      : i-VPLS
Customer Id       : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change  : 09/25/2007 21:45:59
Admin State       : Up                  Oper State        : Down
MTU               : 1514                Def. Mesh VC Id   : 2001
SAP Count         : 1                   SDP Bind Count    : 0
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
b-vpls Id         : 2002                Oper ISID         : 122
Snd Flush in bVpls: Disabled
Snd Flush in bVpls: All-from-me         b-vpls-status     : Up
                  : All-but-mine
-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
SAP 1/1/12:2001.2001
```

```
-------------------------------------------------------------------------------
Service Id         : 2001
SAP                : 1/1/12:2001.2001         Encap             : qinq
Sub Type           : regular
QinQ Dot1p         : Default
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Last Status Change : 09/25/2007 21:12:01
Last Mgmt Change   : 09/25/2007 21:45:59
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1522                     Oper MTU          : 1522
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
tod-suite          : None                     qinq-pbit-marking : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Multi Svc Site     : None
Acct. Pol          : None                     Collect Stats     : Disabled
Restr MacProt Src  : Disabled                 Restr MacUnpr Dst : Disabled
Mac Move           : Non Blockable            Mac Move Block Lvl: Tertiary
Egr MCast Grp      :
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Unknown
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A

Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
SAP MRP Information
-------------------------------------------------------------------------------
Rx Pdus            : 0                        Tx Pdus           : 0
Dropped Pdus       : 0                        Tx Pdus           : 0
Rx New Event       : 0                        Rx Join-In Event  : 0
Rx In Event        : 0                        Rx Join Empty Evt : 0
Rx Empty Event     : 0                        Rx Leave Event    : 0
```

```
Tx New Event       : 0                        Tx Join-In Event  : 0
Tx In Event        : 0                        Tx Join Empty Evt : 0
Tx Empty Event     : 0                        Tx Leave Event    : 0
-------------------------------------------------------------------------------
SAP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
Number of MACs=0 Registered=0 Declared=0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
Off. Uncolor          : 0                        0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                  Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 0                        0
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 0                        0
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
Egress Queue 1
For. InProf           : 0                        0
For. OutProf          : 0                        0
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Down                     Core Connectivity : Down
Stp Admin State    : Down                     Stp Oper State    : Down
```

```
Mode               : Rstp                       Vcp Active Prot.  : N/A
Bridge Id          : 80:00.70:ec:ff:00:00:00    Bridge Instance Id: 0
Bridge Priority    : 32768                      Tx Hold Count     : 6
Topology Change    : Inactive                   Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00                Bridge Max Age    : 20
Top. Change Count  : 0                          Bridge Fwd Delay  : 15
MST region revision: 0                          Bridge max hops   : 20
MST region name    :
Root Bridge        : N/A
Primary Bridge     : N/A
Root Path Cost     : 0                          Root Forward Delay: 15
Rcvd Hello Time    : 2                          Root Max Age      : 20
Root Priority      : 32768                      Root Port         : N/A
-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 2001                Mac Move          : Disabled
Primary Factor    : 3                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
OAM-learned Count : 0                   DHCP-learned Count: 0
Host-learned Count: 0
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Dsabl               Relearn Only      : False
-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                Oper    MRtr Pim   Send      Max Num   MVR        Num
Id                     State   Port Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/1/12:2001.2001   Down    No   No    Disabled  No Limit  Local      0
-------------------------------------------------------------------------------
DHCP Summary, service 2001
-------------------------------------------------------------------------------
Sap/Sdp                Snoop  Used/       Arp Reply   Info     Admin
                              Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/12:2001.2001   No     0/0         No          Keep     Down
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
MRP Information
-------------------------------------------------------------------------------
Admin State        : Down              Failed Register Cnt: 0
Max Attributes     : 2048              Attribute Count    : 0
-------------------------------------------------------------------------------
*A:SetupCLI#


*A:SetupCLI# show service id 2002 all
===============================================================================
Service Detailed Information
===============================================================================
```

```
Service Id        : 2002               Vpn Id            : 0
Service Type      : b-VPLS
Customer Id       : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change  : 09/25/2007 21:45:59
Admin State       : Up                 Oper State        : Down
MTU               : 1530               Def. Mesh VC Id   : 2002
SAP Count         : 2                  SDP Bind Count    : 2
Snd Flush on Fail : Disabled           Host Conn Verify  : Disabled
Oper Backbone Src : 00:f7:f7:f7:f7:f7
-------------------------------------------------------------------------------
Related iVpls services for bVpls service 2002
-------------------------------------------------------------------------------
iVpls SvcId         Oper ISID          Admin               Oper
-------------------------------------------------------------------------------
2001                122                Up                  Down
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
 Sdp Id 2000:2001  -(101.101.101.101)
-------------------------------------------------------------------------------
SDP Id            : 2000:2001              Type              : Spoke
VC Type           : Ether                  VC Tag            : n/a
Admin Path MTU    : 1500                   Oper Path MTU     : 1500
Far End           : 101.101.101.101        Delivery          : MPLS

Admin State       : Down                   Oper State        : Down
Acct. Pol         : None                   Collect Stats     : Disabled
Ingress Label     : 0                      Egress Label      : 0
Ing mac Fltr      : n/a                    Egr mac Fltr      : n/a
Admin ControlWord : Not Preferred          Oper ControlWord  : False
Last Status Change : 09/25/2007 21:12:01   Signaling         : TLDP
Last Mgmt Change  : 09/25/2007 21:45:59    Force Vlan-Vc     : Disabled
Endpoint          : N/A                    Precedence        : 4
Class Fwding State : Down
Flags             : SdpOperDown SdpBindAdminDown
                    NoIngVCLabel NoEgrVCLabel
                    PathMTUTooSmall
Peer Pw Bits      : None
Peer Fault Ip     : None
Max Nbr of MAC Addr: No Limit              Total MAC Addr    : 0
Learned MAC Addr  : 0                      Static MAC Addr   : 0

MAC Learning      : Enabled                Discard Unkwn Srce: Disabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled
Ignore Standby Sig : False

KeepAlive Information :
Admin State       : Enabled                Oper State        : No response
Hello Time        : 600                    Hello Msg Len     : 1500
Max Drop Count    : 3                      Hold Down Time    : 10

Statistics          :
I. Fwd. Pkts.     : 0                      I. Dro. Pkts.     : 0
```

```
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0

Associated LSP LIST :
No LSPs Associate
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : disabled
Default LSP        : Uknwn                    Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable                Blockable Level   : Tertiary
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 2050                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Sdp Id 2000:2001 MRP Information
-------------------------------------------------------------------------------
Rx Pdus            : 0                        Tx Pdus           : 0
Dropped Pdus       : 0
Rx New Event       : 0                        Rx Join-In Event  : 0
Rx In Event        : 0                        Rx Join Empty Evt : 0
Rx Empty Event     : 0                        Rx Leave Event    : 0
Tx New Event       : 0                        Tx Join-In Event  : 0
Tx In Event        : 0                        Tx Join Empty Evt : 0
Tx Empty Event     : 0                        Tx Leave Event    : 0
-------------------------------------------------------------------------------
SDP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
Number of MACs=0 Registered=0 Declared=0
-------------------------------------------------------------------------------
 Sdp Id 2000:2002  -(101.101.101.101)
-------------------------------------------------------------------------------
SDP Id             : 2000:2002                Type              : Mesh
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 1500                     Oper Path MTU     : 1500
Far End            : 101.101.101.101          Delivery          : MPLS

Admin State        : Down                     Oper State        : Down
```

```
Acct. Pol          : None                      Collect Stats     : Disabled
Ingress Label      : 2050                      Egress Label      : 2050
Ing mac Fltr       : n/a                       Egr mac Fltr      : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord  : False
Last Status Change : 09/25/2007 21:12:01       Signaling         : TLDP
Last Mgmt Change   : 09/25/2007 21:45:58       Force Vlan-Vc     : Disabled
Endpoint           : N/A                       Precedence        : 4
Class Fwding State : Down
Flags              : SdpOperDown SdpBindAdminDown
                     PathMTUTooSmall
Peer Pw Bits       : None
Peer Fault Ip      : None
Ignore Standby Sig : False

KeepAlive Information :
Admin State        : Enabled                   Oper State        : No response
Hello Time         : 600                       Hello Msg Len     : 1500
Max Drop Count     : 3                         Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                         I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                         E. Fwd. Octets    : 0

Associated LSP LIST :
No LSPs Associated
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : disabled
Default LSP        : Uknwn                     Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Sdp Id 2000:2002 MRP Information
-------------------------------------------------------------------------------
Rx Pdus            : 0                         Tx Pdus           : 0
Dropped Pdus       : 0
Rx New Event       : 0                         Rx Join-In Event  : 0
Rx In Event        : 0                         Rx Join Empty Evt : 0
Rx Empty Event     : 0                         Rx Leave Event    : 0
Tx New Event       : 0                         Tx Join-In Event  : 0
Tx In Event        : 0                         Tx Join Empty Evt : 0
Tx Empty Event     : 0                         Tx Leave Event    : 0
-------------------------------------------------------------------------------
SDP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
Number of MACs=0 Registered=0 Declared=0
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
SAP 1/1/12:2002.2002
-------------------------------------------------------------------------------
```

```
Service Id         : 2002
SAP                : 1/1/12:2002.2002         Encap             : qinq
Sub Type           : regular
QinQ Dot1p         : Default
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
PBB Ethertype      : 0x88e7

Admin State        : Down                     Oper State        : Down
Flags              : SapAdminDown
                     PortOperDown PortMTUTooSmall
Last Status Change : 09/25/2007 21:12:01
Last Mgmt Change   : 09/25/2007 21:45:58
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1522                     Oper MTU          : 1522
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
tod-suite          : None                     qinq-pbit-marking : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Multi Svc Site     : None
Acct. Pol          : None                     Collect Stats     : Disabled
Restr MacProt Src  : Disabled                 Restr MacUnpr Dst : Disabled
Mac Move           : Blockable                Mac Move Block Lvl: Tertiary
Egr MCast Grp      :
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Unknown
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A

Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
SAP MRP Information
-------------------------------------------------------------------------------
Rx Pdus            : 0                        Tx Pdus           : 0
Dropped Pdus       : 0                        Tx Pdus           : 0
Rx New Event       : 0                        Rx Join-In Event  : 0
Rx In Event        : 0                        Rx Join Empty Evt : 0
Rx Empty Event     : 0                        Rx Leave Event    : 0
```

```
Tx New Event       : 0                        Tx Join-In Event  : 0
Tx In Event        : 0                        Tx Join Empty Evt : 0
Tx Empty Event     : 0                        Tx Leave Event    : 0
-------------------------------------------------------------------------------
SAP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
Number of MACs=0 Registered=0 Declared=0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
Off. Uncolor          : 0                        0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                  Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 0                        0
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 0                        0
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Egress Queue 1
For. InProf           : 0                        0
For. OutProf          : 0                        0
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
-------------------------------------------------------------------------------
```

```
SAP 1/1/30:2002
-------------------------------------------------------------------------------
Service Id         : 2002
SAP                : 1/1/30:2002             Encap             : q-tag
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
PBB Ethertype      : 0x88e7

Admin State        : Down                    Oper State        : Down
Flags              : SapAdminDown
                     PortOperDown PortMTUTooSmall
Last Status Change : 09/25/2007 21:12:01
Last Mgmt Change   : 09/25/2007 21:45:58
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1518                    Oper MTU          : 1518
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None                    qinq-pbit-marking : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Multi Svc Site     : None
Acct. Pol          : None                    Collect Stats     : Disabled
Restr MacProt Src  : Disabled                Restr MacUnpr Dst : Disabled
Mac Move           : Blockable               Mac Move Block Lvl: Tertiary
Egr MCast Grp      :

-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Unknown
Port Number        : 2048                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
SAP MRP Information
-------------------------------------------------------------------------------
Rx Pdus            : 0                       Tx Pdus           : 0
Dropped Pdus       : 0                       Tx Pdus           : 0
Rx New Event       : 0                       Rx Join-In Event  : 0
```

```
Rx In Event        : 0                         Rx Join Empty Evt : 0
Rx Empty Event     : 0                         Rx Leave Event    : 0
Tx New Event       : 0                         Tx Join-In Event  : 0
Tx In Event        : 0                         Tx Join Empty Evt : 0
Tx Empty Event     : 0                         Tx Leave Event    : 0
-------------------------------------------------------------------------------
SAP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
Number of MACs=0 Registered=0 Declared=0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
```

```
VPLS Spanning Tree Information
------------------------------------------------------------------------------
VPLS oper state    : Down                     Core Connectivity : Down
Stp Admin State    : Down                     Stp Oper State    : Down
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.70:ec:ff:00:00:00  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 20
Top. Change Count  : 0                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 32768                    Root Port         : N/A
------------------------------------------------------------------------------
Forwarding Database specifics
------------------------------------------------------------------------------
Service Id        : 2002               Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
OAM-learned Count : 0                  DHCP-learned Count: 0
Host-learned Count: 0
Remote Age        : 900                Local Age         : 300
High WaterMark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabl              Discard Unknown   : Dsabl
Mac Aging         : Dsabl              Relearn Only      : False
------------------------------------------------------------------------------
IGMP Snooping Base info
------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
------------------------------------------------------------------------------
Sap/Sdp                Oper    MRtr Pim   Send      Max Num   MVR        Num
Id                     State   Port Port  Queries   Groups    From-VPLS  Groups
------------------------------------------------------------------------------
sap:1/1/12:2002.2002   Down    No   No    Disabled  No Limit  Local      0
sap:1/1/30:2002        Down    No   No    Disabled  No Limit  Local      0
sdp:2000:2001          Down    No   No    Disabled  No Limit  N/A        0
sdp:2000:2002          Down    No   No    Disabled  No Limit  N/A        0
------------------------------------------------------------------------------
DHCP Summary, service 2002
------------------------------------------------------------------------------
Sap/Sdp                Snoop  Used/       Arp Reply   Info     Admin
                              Provided    Agent       Option   State
------------------------------------------------------------------------------
sap:1/1/12:2002.2002   No     0/0         No          Keep     Down
sap:1/1/30:2002        No     0/0         No          Keep     Down
sdp:2000:2001          No     N/A         N/A         N/A      N/A
sdp:2000:2002          No     N/A         N/A         N/A      N/A
------------------------------------------------------------------------------
Number of Entries : 4
```

```
-------------------------------------------------------------------------------
MRP Information
-------------------------------------------------------------------------------
Admin State        : Up                Failed Register Cnt: 0
Max Attributes     : 2048              Attribute Count    : 2
-------------------------------------------------------------------------------
*A:SetupCLI#


*A:alcag1-R6# show service id 5000 all | match post-lines 15 eth-cfm
eth-cfm Configuration Information
-------------------------------------------------------------------------------
Md-index           : 1                         Direction         : Up
Ma-index           : 1                         Admin             : Enabled
MepId              : 51                        CCM-Enable        : Enabled
LowestDefectPri    : allDef                    HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:ae:ae:ae:ae:ae
CcmTx              : 11548                     CcmSequenceErr    : 2
LbRxReply          : 1                         LbRxBadOrder      : 0
LbRxBadMsdu        : 0                         LbTxReply         : 2
LbNextSequence     : 3                         LtNextSequence    : 3
LtRxUnexplained    : 0
-------------------------------------------------------------------------------


*A:alcmtu1-R6#
*A:alcmtu1-R6# show service id 5000 all | match post-lines 15 eth-cfm
eth-cfm Configuration Information
-------------------------------------------------------------------------------
Md-index           : 1                         Direction         : Up
Ma-index           : 1                         Admin             : Enabled
MepId              : 56                        CCM-Enable        : Enabled
LowestDefectPri    : allDef                    HighestDefect     : defMACstatus
Defect Flags       : bDefMACstatus
Mac Address        : 00:af:af:af:af:af
CcmTx              : 815                       CcmSequenceErr    : 0
LbRxReply          : 0                         LbRxBadOrder      : 0
LbRxBadMsdu        : 0                         LbTxReply         : 0
LbNextSequence     : 3                         LtNextSequence    : 1
LtRxUnexplained    : 0
-------------------------------------------------------------------------------
*A:alcmtu1-R6#

*A:Dut-B# show service id 1 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id         : 1                     Vpn Id            : 0
Service Type       : VPLS
Name               : vpls_1
Description        : (Not Specified)
Customer Id        : 1                     Creation Origin   : manual
Last Status Change: 01/28/2015 21:59:54
Last Mgmt Change   : 01/28/2015 21:59:54
Etree Mode         : Disabled
Admin State        : Up                    Oper State        : Up
MTU                : 1514                  Def. Mesh VC Id   : 1
SAP Count          : 1                     SDP Bind Count    : 1
```

```
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
VSD Domain        : <none>
SPI load-balance  : Disabled
-------------------------------------------------------------------------------
BGP Information
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore              V-Mep Extensions  : Enabled
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 230:1  -(10.20.1.3)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 230:1                   Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp   : (Not Specified)
Etree Root Leaf Tag: Disabled               Etree Leaf AC     : Disabled
VC Type           : Ether                   VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 1582
Delivery          : MPLS
Far End           : 10.20.1.3
Tunnel Far End    : n/a                     LSP Types         : SR-ISIS
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled

Admin State       : Up                      Oper State        : Up
Acct. Pol         : None                    Collect Stats     : Disabled
Ingress Label     : 262135                  Egress Label      : 262135
Ingr Mac Fltr-Id  : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id   : n/a                     Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id : n/a                     Egr IPv6 Fltr-Id  : n/a
Admin ControlWord : Not Preferred           Oper ControlWord  : False
BFD Template      : None
BFD-Enabled       : no                      BFD-Encap         : ipv4
Last Status Change : 01/28/2015 22:00:07    Signaling         : TLDP
Last Mgmt Change  : 01/28/2015 21:59:53
Endpoint          : N/A                     Precedence        : 4
PW Status Sig     : Enabled
Force Vlan-Vc     : Disabled                Force Qinq-Vc     : Disabled
Class Fwding State : Down
Flags             : None
Time to RetryReset : never                  Retries Left      : 3
Mac Move          : Blockable               Blockable Level   : Tertiary
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None
```

```
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel

Application Profile: None
Transit Policy     : None
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
OAM MAC Addr       : 0                       DHCP MAC Addr     : 0
Host MAC Addr      : 0                       Intf MAC Addr     : 0
SPB MAC Addr       : 0                       Cond MAC Addr     : 0
BGP EVPN Addr      : 0                       EVPN Static Addr  : 0

MAC Learning       : Enabled                 Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                   Block On Mesh Fail: False
Oper Group         : (none)                  Monitor Oper Grp  : (none)
Rest Prot Src Mac  : Disabled
Auto Learn Mac Prot: Disabled                RestProtSrcMacAct : Disable
Ing. Vlan Trans.   : 0

Ingress Qos Policy : (none)                  Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                       I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0
-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false

MCAC Policy Name   :
MCAC Max Unconst BW: no limit                MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited
-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
V-MEP Filtering    : Disabled
Squelch Levels     : None
-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
```

```
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                   Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 0                       Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                     Designated Port Id: 0

Fwd Transitions    : 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                       TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAP 1/1/8:1.1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/8:1.1               Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 01/28/2015 21:59:54
Last Mgmt Change   : 01/28/2015 21:59:53
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Etree Root Leaf Tag: Disabled                Etree Leaf Tag    : 0
Etree Leaf AC      : Disabled
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
OAM MAC Addr       : 0                       DHCP MAC Addr     : 0
Host MAC Addr      : 0                       Intf MAC Addr     : 0
SPB MAC Addr       : 0                       Cond MAC Addr     : 0
BGP EVPN Addr      : 0                       EVPN Static Addr  : 0
```

```
Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
                                             Egr Agg Rate Limit: max
Q Frame-Based Acct : Disabled                Limit Unused BW   : Disabled
ARP Reply Agent    : Disabled                Host Conn Verify  : Disabled
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None                    Collect Stats     : Disabled

Anti Spoofing      : None                    Dynamic Hosts     : Enabled
Avl Static Hosts   : 0                       Tot Static Hosts  : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy     : None

Oper Group         : (none)                  Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Cflowd             : Disabled
MCAC Policy Name   :                         MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited
Use LAG port weight: no
Restr MacProt Src  : Disabled                Restr MacUnpr Dst : Disabled
Auto Learn Mac Prot: Disabled                RestProtSrcMacAct : Disable
Time to RetryReset : never                   Retries Left      : 3
Mac Move           : Blockable               Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy        : None
-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : accept                  AIS               : Disabled
MC Prop-Hold-Timer : n/a                     V-MEP Filtering   : Disabled
Squelch Levels     : None
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : N/A                     Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
```

```
Cfg BPDUs rcvd     : 0                          Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                          TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                          TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                          RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                          MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                          Min Auth Interval : 15 minutes
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 2                          Egress qos-policy : 2
Ingress FP QGrp    : (none)                     Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                     Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                        Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down                       Lease Populate    : 0
DHCP Snooping      : Down                       Action            : Keep

Proxy Admin State  : Down
Proxy Lease Time   : N/A
Emul. Server Addr  : Not Configured
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Down                       MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : None
Def SLA-Profile    : None
Def Inter-Dest-Id  : None
Def App-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
 Prof Traffic Only : False
 Non-Sub-Traffic   : N/A

Static host management
MAC learn options  : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                  Octets
CPM Ingress           : 0                        0

Forwarding Engine Stats
Dropped               : 0                        0
```

```
Received Valid        : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats(Ingress QoS Policy 2)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 2)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Managed          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Down                     Stp Oper State    : Down
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.00:03:fa:32:16:62  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 20
Top. Change Count  : 0                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    :
```

```
Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 0
Rcvd Hello Time    : 0                        Root Max Age      : 0
Root Priority      : 0                        Root Port         : N/A

-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 1                  Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
OAM MAC Count     : 0                  DHCP MAC Count    : 0
Host MAC Count    : 0                  Intf MAC Count    : 0
Spb Count         : 0                  Cond MAC Count    : 0
BGP EVPN Count    : 0                  EVPN Static Cnt   : 0
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False
Mac Subnet Len    : 48
-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                   Oper MRtr Pim  Send Max   Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps  Srcs Grp   From-VPLS Grps
                                                      Srcs
-------------------------------------------------------------------------------
sap:1/1/8:1.1             Up   No   No   No   None  None None  Local     0
sdp:230:1                 Up   No   No   No   None  None None  N/A       0

-------------------------------------------------------------------------------
MLD Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR        Num
Id                       State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/1/8:1.1            Up      No    Disabled  No Limit  Local      0
sdp:230:1                Up      No    Disabled  No Limit  N/A        0

-------------------------------------------------------------------------------
DHCP Summary, service 1
-------------------------------------------------------------------------------
Sap/Sdp                Snoop  Used/       Arp Reply   Info     Admin
                              Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/8:1.1          No     0/0         No          Keep     Down
sdp:230:1              No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
```

```
Number of Entries : 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
ARP host Summary, service 1
-------------------------------------------------------------------------------
Sap                    Used        Provided    Admin State
-------------------------------------------------------------------------------
sap:1/1/8:1.1          0           1           outOfService
-------------------------------------------------------------------------------
Number of SAPs : 1     0
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
-------------------------------------------------------------------------------
WLAN Gateway specifics
-------------------------------------------------------------------------------
Admin State            : disabled
Description            : (Not Specified)
SAP-template           : (Not Specified)
Last management change : (Not Specified)
No associated WLAN Gateway interface VLAN tag ranges found.
===============================================================================

===============================================================================
Service VPLS Group Information
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 0
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address           Egress VNI     Num. MACs    In Mcast List?  Oper State
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
-------------------------------------------------------------------------------
===============================================================================
VPLS Sites
===============================================================================
Site                Site-Id   Dest                Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
*A:Dut-B#

*A:Dut-B>config>service>sdp# show service id 1 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
Name              : vpls_1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 05/27/2015 06:55:40
```

```
Last Mgmt Change  : 05/27/2015 06:55:40
Etree Mode        : Disabled
Admin State       : Up                  Oper State        : Up
MTU               : 1514                Def. Mesh VC Id   : 1
SAP Count         : 1                   SDP Bind Count    : 1
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled            Fwd-IPv4-Mcast-To*: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
VSD Domain        : <none>
SPI load-balance  : Disabled
TEID load-balance : Disabled

-------------------------------------------------------------------------------
BGP Information
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore              V-Mep Extensions  : Enabled
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 230:1  -(10.20.1.3)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 230:1                 Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp   : (Not Specified)
Etree Root Leaf Tag: Disabled             Etree Leaf AC     : Disabled
VC Type           : Ether                 VC Tag            : n/a
Admin Path MTU    : 0                     Oper Path MTU     : 1582
Delivery          : MPLS
Far End           : 10.20.1.3
Tunnel Far End    : n/a                   LSP Types         : SR-OSPF
Hash Label        : Disabled              Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled

Admin State       : Up                    Oper State        : Up
Acct. Pol         : None                  Collect Stats     : Disabled
Ingress Label     : 262142                Egress Label      : 262141
Ingr Mac Fltr-Id  : n/a                   Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id   : n/a                   Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id : n/a                   Egr IPv6 Fltr-Id  : n/a
Admin ControlWord : Not Preferred         Oper ControlWord  : False
BFD Template      : None
BFD-Enabled       : no                    BFD-Encap         : ipv4
Last Status Change : 05/27/2015 06:59:46  Signaling         : TLDP
Last Mgmt Change  : 05/27/2015 06:55:40
Endpoint          : N/A                   Precedence        : 4
```

```
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                Force Qinq-Vc     : Disabled
Class Fwding State : Down
Flags              : None
Time to RetryReset : never                   Retries Left      : 3
Mac Move           : Blockable               Blockable Level   : Tertiary
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel

Application Profile: None
Transit Policy     : None
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
OAM MAC Addr       : 0                       DHCP MAC Addr     : 0
Host MAC Addr      : 0                       Intf MAC Addr     : 0
SPB MAC Addr       : 0                       Cond MAC Addr     : 0
BGP EVPN Addr      : 0                       EVPN Static Addr  : 0

MAC Learning       : Enabled                 Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                   Block On Mesh Fail: False
Oper Group         : (none)                  Monitor Oper Grp  : (none)
Rest Prot Src Mac  : Disabled
Auto Learn Mac Prot: Disabled                RestProtSrcMacAct : Disable
Ing. Vlan Trans.   : 0

Ingress Qos Policy : (none)                  Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                       I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0
-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false

MCAC Policy Name   :
MCAC Max Unconst BW: no limit                MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited
-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
```

```
-------------------------------------------------------------------------------
V-MEP Filtering    : Disabled
Squelch Levels     : None
-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                   Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
OSPF               : enabled                 LSP Id            : 524291
Oper Instance Id   : 0
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 0                       Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                     Designated Port Id: 0

Fwd Transitions    : 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                       TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAP 1/1/8:1.1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/8:1.1               Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
```

```
Last Status Change : 05/27/2015 06:55:40
Last Mgmt Change   : 05/27/2015 06:55:40
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Etree Root Leaf Tag: Disabled                Etree Leaf Tag    : 0
Etree Leaf AC      : Disabled
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
OAM MAC Addr       : 0                       DHCP MAC Addr     : 0
Host MAC Addr      : 0                       Intf MAC Addr     : 0
SPB MAC Addr       : 0                       Cond MAC Addr     : 0
BGP EVPN Addr      : 0                       EVPN Static Addr  : 0
Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : both
                                             Egr Agg Rate Limit: max
Q Frame-Based Acct : Disabled                Limit Unused BW   : Disabled
ARP Reply Agent    : Disabled                Host Conn Verify  : Disabled
SHCV pol IPv4      : None
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None                    Collect Stats     : Disabled

Anti Spoofing      : None                    Dynamic Hosts     : Enabled
Avl Static Hosts   : 0                       Tot Static Hosts  : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy     : None

Oper Group         : (none)                  Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Cflowd             : Disabled
MCAC Policy Name   :                         MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited
Use LAG port weight: no
Restr MacProt Src  : Disabled                Restr MacUnpr Dst : Disabled
Auto Learn Mac Prot: Disabled                RestProtSrcMacAct : Disable
Time to RetryReset : never                   Retries Left      : 3
Mac Move           : Blockable               Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy        : None
-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : accept                  AIS               : Disabled
MC Prop-Hold-Timer : n/a                     V-MEP Filtering   : Disabled
Squelch Levels     : None
```

```
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Forwarding
Port Number        : N/A                       Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                       Designated Port   : N/A

Forward transitions: 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                         TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                         MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                         Min Auth Interval : 15 minutes
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 2                         Egress qos-policy : 2
Ingress FP QGrp    : (none)                    Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                    Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                       Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down                      Lease Populate    : 0
DHCP Snooping      : Down                      Action            : Keep

Proxy Admin State  : Down
Proxy Lease Time   : N/A
Emul. Server Addr  : Not Configured
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Down                      MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : None
Def SLA-Profile    : None
Def Inter-Dest-Id  : None
Def App-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
```

```
 Prof Traffic Only : False
 Non-Sub-Traffic   : N/A

Static host management
MAC learn options  : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                          Packets                 Octets
CPM Ingress           : 0                         0

Forwarding Engine Stats
Dropped               : 0                         0
Received Valid        : 0                         0
Off. HiPrio           : 0                         0
Off. LowPrio          : 0                         0
Off. Uncolor          : 0                         0
Off. Managed          : 0                         0

Queueing Stats(Ingress QoS Policy 2)
Dro. HiPrio           : 0                         0
Dro. LowPrio          : 0                         0
For. InProf           : 0                         0
For. OutProf          : 0                         0

Queueing Stats(Egress QoS Policy 2)
Dro. InProf           : 0                         0
Dro. OutProf          : 0                         0
For. InProf           : 0                         0
For. OutProf          : 0                         0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                          Packets                 Octets

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                         0
Off. LowPrio          : 0                         0
Off. Managed          : 0                         0
Dro. HiPrio           : 0                         0
Dro. LowPrio          : 0                         0
For. InProf           : 0                         0
For. OutProf          : 0                         0
-------------------------------------------------------------------------------
Sap per Policer stats
-------------------------------------------------------------------------------
                          Packets                 Octets

Ingress Policer 1 (Stats mode: minimal)
Off. All              : 0                         0
Dro. All              : 0                         0
For. All              : 0                         0

Egress Policer 1 (Stats mode: minimal)
Off. All              : 0                         0
Dro. All              : 0                         0
For. All              : 0                         0
```

```
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                        Core Connectivity : Down
Stp Admin State    : Down                      Stp Oper State    : Down
Mode               : Rstp                      Vcp Active Prot.  : N/A

Bridge Id          : 80:00.66:30:ff:00:00:00   Bridge Instance Id: 0
Bridge Priority    : 32768                     Tx Hold Count     : 6
Topology Change    : Inactive                  Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00               Bridge Max Age    : 20
Top. Change Count  : 0                         Bridge Fwd Delay  : 15
MST region revision: 0                         Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                         Root Forward Delay: 0
Rcvd Hello Time    : 0                         Root Max Age      : 0
Root Priority      : 0                         Root Port         : N/A

-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 1                  Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
OAM MAC Count     : 0                  DHCP MAC Count    : 0
Host MAC Count    : 0                  Intf MAC Count    : 0
Spb Count         : 0                  Cond MAC Count    : 0
BGP EVPN Count    : 0                  EVPN Static Cnt   : 0
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False
Mac Subnet Len    : 48
-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                  Oper MRtr Pim  Send Max   Max  Max   MVR       Num
Id                       Stat Port Port Qrys Grps  Srcs Grp   From-VPLS Grps
                                                      Srcs
-------------------------------------------------------------------------------
sap:1/1/8:1.1            Up   No   No   No   None  None None  Local     0
sdp:230:1                Up   No   No   No   None  None None  N/A       0

-------------------------------------------------------------------------------
MLD Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
```

```
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR         Num
Id                       State   Port  Queries   Groups    From-VPLS   Groups
-------------------------------------------------------------------------------
sap:1/1/8:1.1            Up      No    Disabled  No Limit  Local       0
sdp:230:1                Up      No    Disabled  No Limit  N/A         0
-------------------------------------------------------------------------------
DHCP Summary, service 1
-------------------------------------------------------------------------------
Sap/Sdp                Snoop  Used/       Arp Reply   Info     Admin
                              Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/8:1.1          No     0/0         No          Keep     Down
sdp:230:1              No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
ARP host Summary, service 1
-------------------------------------------------------------------------------
Sap                    Used        Provided    Admin State
-------------------------------------------------------------------------------
sap:1/1/8:1.1          0           1           outOfService
-------------------------------------------------------------------------------
Number of SAPs : 1     0
===============================================================================
-------------------------------------------------------------------------------
WLAN Gateway specifics
-------------------------------------------------------------------------------
Admin State            : disabled
Description            : (Not Specified)
SAP-template           : (Not Specified)
Last management change : (Not Specified)
No associated WLAN Gateway interface VLAN tag ranges found.
===============================================================================
===============================================================================
Service VPLS Group Information
===============================================================================
===============================================================================
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 0

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address           Egress VNI     Num. MACs    In Mcast List?  Oper State
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
-------------------------------------------------------------------------------
===============================================================================
VPLS Sites
===============================================================================
Site                Site-Id   Dest                Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
No Matching Entries
```

```
===============================================================================
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Table 52 describes the command output fields.

Table 52 **Show Service ID All Output Fields**

Label	Description
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	Specifies the type of service.
SDP Id	The SDP identifier.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group	Name of the split horizon group for this service.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
SDP Id	The SDP identifier.
Type	Indicates whether this service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.

Table 52 **Show Service ID All Output Fields (Continued)**

Label	Description (Continued)
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS
Number of SDPs	The total number SDPs applied to this service ID.
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.

Table 52 **Show Service ID All Output Fields (Continued)**

Label	Description (Continued)
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Dropped In Profile	The number of in-profile packets or octets discarded.
Dropped Out Profile	The number of out-of-profile packets or octets discarded.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.

Table 52 **Show Service ID All Output Fields (Continued)**

Label	Description (Continued)
Ingress Queue 1	The index of the ingress QoS queue of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.
High priority dropped	The number of high priority traffic packets/octets dropped.
Low priority offered	The packets or octets count of the low priority traffic.
Low priority dropped	The number of low priority traffic packets/octets dropped.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded.
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
State	Specifies whether DHCP Relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop.
Circuit ID	Specifies whether the If Index is inserted in Circuit ID sub-option of Option 82.
Remote ID	Specifies whether the far-end MAC address is inserted in Remote ID sub-option of Option 82.
Managed by Service	Specifies the service-id of the management VPLS managing this SAP.
Managed by MSTI	Specifies the MST instance inside the management VPLS managing this SAP.
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by SAP	Specifies the sap-id inside the management VPLS managing this SAP.

Table 52 **Show Service ID All Output Fields (Continued)**

Label	Description (Continued)
Prune state	Specifies the STP state inherited from the management VPLS.
Managed by Service	Specifies the service-id of the management VPLS managing this spoke-SDP.
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke-SDP.
Prune state	Specifies the STP state inherited from the management VPLS.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the above failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signaling method to indicate faults. pwNotForwarding — Pseudowire not forwarding. lacIngressFault Local — Attachment circuit RX fault. lacEgressFault Local — Attachment circuit TX fault. psnIngressFault Local — PSN-facing PW RX fault. psnEgressFault Local — PSN-facing PW TX fault. pwFwdingStandby — Pseudowire in standby mode.

arp

Syntax **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

Context show>service>id

Description This command displays the ARP table for the VPLS instance. The ARP entries for a subscriber interface are displayed uniquely. Each MAC associated with the subscriber interface child group-interfaces is displayed with each subscriber interface ARP entry for easy lookup.

Parameters *ip-address* — Displays all IP addresses

ieee-address — Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address is in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee and ff are hexadecimal numbers.

Default All MAC addresses.

sap-id — Displays SAP information for the specified SAP ID

interface — Specifies matching service ARP entries associated with the IP interface

ip-address — Displays the IP address of the interface for which the matching ARP entries will be displayed

Values 1.0.0.0 to 223.255.255.255

ip-int-name — Displays the IP interface name for which the matching ARPs will be displayed

Output The following output displays an example of service ARP information.

Table 53 describes show service-id ARP output fields.

Table 53 **Show Service-ID ARP Fields**

Label	Description
IP Address	The IP address.
MAC Address	The specified MAC address. Type Static — FDB entries created by management. Learned — Dynamic entries created by the learning process. Other — Local entries for the IP interfaces created.
Expiry	The age of the ARP entry.
Interface	The interface applied to the service.
SAP	The SAP ID.

authentication

Syntax **authentication**

Context show>service>id

Description This command enables the context to show session authentication information.

statistics

Syntax **statistics** [**policy** *name*] [**sap** *sap-id*]

Context show>service>id>auth

Description This command displays subscriber authentication statistics.

arp-host

Syntax **arp-host** [**wholesaler** *service-id*] [**sap** *sap-id* | **interface** *interface-name* | **ip-address** *ip-address*[/*mask*] | **mac** *ieee-address* | {[**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *inter-dest-id*]}] [**detail**]

arp-host statistics [**sap** *sap-id* | **interface** *interface-name*]

arp-host summary [**interface** *interface-name*]

Context show>service>id

Description This command displays ARP host related information.

Output The following output is an example of service ARP host information.

Sample Output

```
*A:Dut-C# show service id 2 arp-host
===============================================================================
ARP host table, service 2
===============================================================================
IP Address      Mac Address       Sap Id              Remaining           MC
                                                      Time                Stdby
-------------------------------------------------------------------------------
128.128.1.2     00:80:00:00:00:01 2/1/5:2             00h04m41s
128.128.1.3     00:80:00:00:00:02 2/1/5:2             00h04m42s
128.128.1.4     00:80:00:00:00:03 2/1/5:2             00h04m43s
128.128.1.5     00:80:00:00:00:04 2/1/5:2             00h04m44s
128.128.1.6     00:80:00:00:00:05 2/1/5:2             00h04m45s
128.128.1.7     00:80:00:00:00:06 2/1/5:2             00h04m46s
128.128.1.8     00:80:00:00:00:07 2/1/5:2             00h04m47s
128.128.1.9     00:80:00:00:00:08 2/1/5:2             00h04m48s
128.128.1.10    00:80:00:00:00:09 2/1/5:2             00h04m49s
128.128.1.11    00:80:00:00:00:0a 2/1/5:2             00h04m50s
-------------------------------------------------------------------------------
Number of ARP hosts : 10
===============================================================================
*A:Dut-C#


*A:Dut-C# show service id 2 arp-host ip-address 128.128.1.2 detail
===============================================================================
ARP hosts for service 2
===============================================================================
Service ID           : 2
IP Address           : 128.128.1.2
MAC Address          : 00:80:00:00:00:01
SAP                  : 2/1/5:2
Remaining Time       : 00h04m58s

Sub-Ident            : "alu_1_2"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
RADIUS-User-Name     : "128.128.1.2"
```

```
Session Timeout (s)  : 301
Start Time           : 02/09/2009 16:35:07
Last Auth            : 02/09/2009 16:36:34
Last Refresh         : 02/09/2009 16:36:38
Persistence Key      : N/A
-------------------------------------------------------------------------------
Number of ARP hosts : 1
===============================================================================
*A:Dut-C#


*A:Dut-C# show service id 2 arp-host statistics
===============================================================================
ARP host statistics
===============================================================================
Num Active Hosts             : 20
Received Triggers            : 70
Ignored Triggers             : 10
Ignored Triggers (overload)  : 0
SHCV Checks Forced           : 0
Hosts Created                : 20
Hosts Updated                : 40
Hosts Deleted                : 0
Authentication Requests Sent : 40
===============================================================================
*A:Dut-C#


*A:Dut-C# show service id 2 arp-host summary
============================================================
ARP host Summary, service 2
============================================================
Sap                    Used        Provided    Admin State
------------------------------------------------------------
sap:2/1/5:2            20          8000        inService
------------------------------------------------------------
Number of SAPs : 1
------------------------------------------------------------
============================================================
*A:Dut-C#
```

base

Syntax **base** [**msap**]

Context show>service>id
show>service>id>igmp-snooping

Description This command displays basic information about the service ID including service type, description, SAPs and SDPs.

Parameters **msap** — Displays management SAPs. This parameter applies to the 7450 ESS or 7750 SR only.

Output The following output is an example of service base information.

Sample Output

```
A:Sr-4# show service id 300 sap 1/1/1:300.* detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id :300
SAP :1/1/1:300.* Encap :qinq
QinQ Dot1p :Default
Admin State :Up Oper State:Up
Flags :None
Multi Svc Site :None
Last Status Change :11/19/2007 20:42:34
Last Mgmt Change :11/19/2007 20:42:25
Sub Type :regular
Dot1Q Ethertype    :0x8100 QinQ Ethertype:0x8100

Admin MTU :1522 Oper MTU:1522
Ingr IP Fltr-Id    :n/a Egr IP Fltr-Id :n/a
Ingr Mac Fltr-Id   :n/a Egr Mac Fltr-Id :n/a
Ingr IPv6 Fltr-Id  :n/a Egr IPv6 Fltr-Id :n/a
tod-suite :None qinq-pbit-marking :both
Egr Agg Rate Limit :max Endpoint:N/A
Q Frame-Based Acct :Disabled
Vlan-translation   :None

Acct. Pol : None Collect Stats:Disabled
Ingress qos-policy : 1 Egress qos-policy :1
Shared Q plcy : n/a Multipoint shared :Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time: 11/19/2007 21:23:45

 Packets Octets
Forwarding Engine Stats
Dropped       : 00
Off. HiPrio      : 00
Off. LowPrio      : 00
Off. Uncolor      : 00

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio      : 00
Dro. LowPrio      : 00
For. InProf      : 00
For. OutProf      : 00

Queueing Stats(Egress QoS Policy 1)
Dro. InProf      : 00
Dro. OutProf      : 00
For. InProf      : 00
For. OutProf      : 00
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
Packets Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 00
Off. LoPrio      : 00
```

```
Dro. HiPrio      : 00
Dro. LoPrio      : 00
For. InProf      : 00
For. OutProf      : 00

Egress Queue 1
For. InProf      : 00
For. OutProf      : 00
Dro. InProf      : 00
Dro. OutProf      : 00
===============================================================================
*A:Sr-4#
*A:SetupCLI# show service id 2001 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2001               Vpn Id            : 0
Service Type      : i-VPLS
Customer Id       : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change  : 09/25/2007 21:45:59
Admin State       : Up                 Oper State        : Down
MTU               : 1514               Def. Mesh VC Id   : 2001
SAP Count         : 1                  SDP Bind Count    : 0
Snd Flush on Fail : Disabled           Host Conn Verify  : Disabled
b-vpls Id         : 2002               Oper ISID         : 122
Snd Flush in bVpls: Disabled
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                       Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/12:2001.2001             qinq         1522    1522    Up      Down
-------------------------------------------------------------------------------
[<sap-id>] indicates a Managed SAP
===============================================================================
*A:SetupCLI#


*A:SetupCLI# show service id 2002 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2002               Vpn Id            : 0
Service Type      : b-VPLS
Customer Id       : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change  : 09/25/2007 21:45:59
Admin State       : Up                 Oper State        : Down
MTU               : 1530               Def. Mesh VC Id   : 2002
SAP Count         : 2                  SDP Bind Count    : 2
Snd Flush on Fail : Disabled           Host Conn Verify  : Disabled
Oper Backbone Src : 00:f7:f7:f7:f7:f7
-------------------------------------------------------------------------------
Related iVpls services for bVpls service 2002
-------------------------------------------------------------------------------
iVpls SvcId         Oper ISID          Admin               Oper
-------------------------------------------------------------------------------
2001                122                Up                  Down
```

```
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/12:2002.2002            qinq         1522    1522    Down    Down
sap:1/1/30:2002                 q-tag        1518    1518    Down    Down
sdp:2000:2001 S(101.101.101.101) n/a         1500    1500    Down    Down
sdp:2000:2002 M(101.101.101.101) n/a         1500    1500    Down    Down
-------------------------------------------------------------------------------
[<sap-id>] indicates a Managed SAP
===============================================================================


A:ALA-48>config>service>vpls# show service id 700 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 700                 Vpn Id            : 0
Service Type      : VPLS
Description       : IMA VPLS
Customer Id       : 7
Last Status Change: 11/21/2008 17:33:20
Last Mgmt Change  : 11/21/2008 17:33:34
Admin State       : Up                  Oper State        : Down
MTU               : 1514                Def. Mesh VC Id   : 700
SAP Count         : 1                   SDP Bind Count    : 2
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None

-------------------------------------------------------------------------------
BGP Auto-discovery Information
-------------------------------------------------------------------------------
Admin State       : Down                Vpls Id           : None
Route Dist        : None                Prefix            : 10.10.10.103
Rte-Target Import : None                Rte-Target Export : None
Vsi-Import        : None
Vsi-Export        : None
PW-Template Id    : None
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm   Opr
-------------------------------------------------------------------------------
sap:1/1/9:0                              q-tag        1518    1518    Up    Down
sdp:2:222 S(10.10.10.104)                n/a          0       0       Up    Down
sdp:2:700 M(10.10.10.104)                n/a          0       0       Up    Down
===============================================================================
A:ALA-48>config>service>vpls#
```

Table 54 describes show service-id base output fields.

Table 54 **Show Service-ID Base Fields**

Label	Description
Service Id	The service identifier.
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	Displays the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The administrative state of the service.
Oper	The operational state of the service.
Mtu	The largest frame size (in octets) that the service can handle.
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service.
SDP Bind Count	The number of SDPs bound to the service.
Identifier	Specifies the service access (SAP) and destination (SDP) points.
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
AdmMTU	Specifies the largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented.
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this service to the far-end ESR, without requiring the packet to be fragmented.
Opr	The operating state of the SAP.

epipe

Syntax **epipe**

Context show>service>id

Description This command displays Epipe services associated with the B-VPLS service. The command only applies when the service is a B-VPLS.

Output The following output is an example of Epipe service information.

Sample Output

```
*A:term17>show>service>id# epipe
===============================================================================
Related Epipe services for bVpls service 2000
===============================================================================
Epipe SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
100                 100                 Down                Down
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
*A:term17>show>service>id#
```

fdb

Syntax **fdb** [**sap** *sap-id* [**expiry**]] | [**sdp** *sdp-id* [**expiry**]] | [**mac** *ieee-address* [**expiry**]] | **endpoint** *endpoint* | [**detail**] [**expiry**] [**pbb**]

Context show>service>id
show>service>fdb-mac

Description This command displays FDB entries for a specified MAC address.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP

detail — Displays detailed information

expiry — Displays time until MAC is aged out

pbb — Displays PBB related information. This keyword is only applicable to B-VPLS or I-VPLS services. This parameter applies to the 7450 ESS or 7750 SR only.

endpoint-name — Specifies an endpoint name up to 32 characters in length. This parameter applies to the 7450 ESS or 7750 SR only.

Output The following output is an example of service FDB information.

Sample Output

```
A:ALA-48>show>service>id# fdb mac detail
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
6         00:aa:00:00:00:00 sap:lag-2             L/0       06/27/2006 15:04:31
6         00:aa:00:00:00:01 sap:lag-2             L/0       06/27/2006 15:04:31
6         00:aa:00:00:00:02 sap:lag-2             L/0       06/27/2006 15:04:31
```

```
6         00:aa:00:00:00:03 sap:lag-2             L/0       06/27/2006 15:04:31
6         00:aa:00:00:00:04 sap:lag-2             L/0       06/27/2006 15:04:31
10        12:12:12:12:12:12 sap:1/1/1:100         S         06/26/2006 10:03:29
===============================================================================
A:ALA-48>show>service>id#


A:PE-1# show service id 1 fdb detail
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:00:00:00:00:01 sap:1/1/1                LP/0     02/24/12 11:40:07
-------------------------------------------------------------------------------
No. of MAC Entries: 1
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC
===============================================================================
A:PE-1#

A:PE1# show service id 1 fdb
===============================================================================
Forwarding Database, Service 1
===============================================================================
Service Id        : 1                   Mac Move          : Disabled
Primary Factor    : 3                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                 Allocated Count   : 4
Total In Use      : 4
Learned Count     : 2                   Static Count      : 0
OAM MAC Count     : 0                   DHCP MAC Count    : 0
Host MAC Count    : 0                   Intf MAC Count    : 0
Spb Count         : 0                   Cond MAC Count    : 0
BGP EVPN Count    : 0                   EVPN Static Cnt   : 2
EVPN Dup Det Cnt  : 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
Mac Subnet Len    : 48
Sel Learned FDB   : Disabled
===============================================================================
A:PE1#

*A:cses-B0102>show>service>id# fdb detail
===============================================================================
Forwarding Database, Service 510
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
510       00:00:00:aa:aa:aa sap:1/1/22:510           CStatic  06/14/13 20:16:19
510       00:00:00:bb:bb:bb sap:1/1/22:510           CStatic  06/14/13 20:14:49
510       00:00:00:dd:dd:dd sdp:7:2                  Spb      06/14/13 20:03:23
510       d8:da:ff:00:00:00 sap:1/1/22:510           CStatic  06/14/13 21:06:38
510       d8:e0:ff:00:00:00 sdp:7:2                  Spb      06/14/13 21:09:29
```

```
-------------------------------------------------------------------------------
No. of MAC Entries: 5
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC
===============================================================================

A:term17>config>service# show service id  2000 fdb pbb
(BVPLS = 2000, IVPLS = 2100)
===============================================================================
Forwarding Database, bVpls Service 2000
===============================================================================
MAC               Source-Identifier     iVplsMACs  Type/Age Last Change
-------------------------------------------------------------------------------
00:f4:f4:f4:f4:f4 sdp:100:2000          10         L/0      09/25/2007 15:34:19
===============================================================================
A:term17>config>service#

*A:SetupCLI# show service id 2100 fdb pbb
===============================================================================
Forwarding Database, iVpls Service 2100
===============================================================================
MAC               Source-Identifier     B-Svc     bVpls MAC         Type/Age
-------------------------------------------------------------------------------
76:55:ff:00:01:a4 b-sdp:100:2000        2000      00:f4:f4:f4:f4:ff L/0
76:55:ff:00:01:bb sap:1/1/1:2100        2000      N/A               Static
===============================================================================
*A:SetupCLI#


A:term17>config>service# show service id  2100 fdb pbb
===============================================================================
Forwarding Database, iVpls Service 2100
===============================================================================
MAC               Source-Identifier     B-Svc     bVpls MAC         Type/Age
-------------------------------------------------------------------------------
00:f4:f4:f4:00:00 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:01 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:02 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:03 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:04 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:05 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:06 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:07 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:08 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:09 b-sdp:100:2000        2000      00:f4:f4:f4:f4:f4 L/0
00:f7:f7:f7:00:00 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:01 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:02 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:03 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:04 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:06 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:07 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:08 sap:lag-1:2100        2000      N/A               L/0
00:f7:f7:f7:00:09 sap:lag-1:2100        2000      N/A               L/0
===============================================================================
A:term17>config>service#


*A:SetupCLI# show service id 2100 fdb pbb
```

```
===============================================================================
Forwarding Database, iVpls Service 2100
===============================================================================
MAC               Source-Identifier     B-Svc     bVpls MAC         Type/Age
-------------------------------------------------------------------------------
76:55:ff:00:01:a4 b-sdp:100:2000        2000      00:f4:f4:f4:f4:ff L/0
76:55:ff:00:01:bb sap:1/1/1:2100        2000      N/A               Static
===============================================================================
*A:SetupCLI#


*A:term17>config>service>epipe# show service id 2000 fdb detail pbb
===============================================================================
Forwarding Database, bVpls Service 2000
===============================================================================
MAC               Source-Identifier     iVplsMACs  Epipes     Type/Age
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:term17>config>service>epipe#


*A:term17>config>service>epipe# show service id 2100 fdb detail
===============================================================================
Forwarding Database, Service 2100
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:term17>config>service>epipe# show service id 2100 fdb detail pbb
===============================================================================
Forwarding Database, iVpls Service 2100
===============================================================================
MAC               Source-Identifier     B-Svc     bVpls MAC         Type/Age
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:term17>config>service>epipe#

# show service id 1 fdb detail
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier         Type     Last Change
                                                      Age
-------------------------------------------------------------------------------
1         00:ca:ca:ba:ca:01 eES:                      Evpn     06/29/15 23:21:34
                            01:00:00:00:00:71:00:00:00:01
1         00:ca:ca:ba:ca:06 eES:                      Evpn     06/29/15 23:21:34
                            01:74:13:00:74:13:00:00:74:13
1         00:ca:00:00:00:00 sap:1/1/1:2               CStatic  06/29/15 23:20:58
1         00:ca:fe:ca:fe:00 black-hole                EvpnD:P  06/29/15 23:20:00
1         00:ca:fe:ca:fe:69 eMpls:                    EvpnS    06/29/15 20:40:13
                            192.0.2.69:262133
1         00:ca:fe:ca:fe:70 eMpls:                    EvpnS    06/29/15 20:43:29
                            192.0.2.70:262132
1         00:ca:fe:ca:fe:72 eMpls:                    EvpnS    06/29/15 20:47:39
                            192.0.2.72:262132
```

```
-------------------------------------------------------------------------------
No. of MAC Entries: 7
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================

A:PE-5# show service id 30 fdb detail
===============================================================================
Forwarding Database, Service 30
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
30        00:ca:ca:00:00:01 eMpls:                   Evpn     12/15/16 18:34:17
                            192.0.2.3:262129
30        00:ca:ca:00:00:05 eMpls:                   Evpn     12/15/16 18:34:17
                            192.0.2.3:262129
30        00:ca:fe:00:00:01 eES:                     Evpn     12/15/16 18:32:44
                                                     Lf
                            01:23:23:23:23:23:23:23:23:23
30        00:ca:fe:00:00:02 eES:                     Evpn     12/15/16 18:32:44
                                                     Lf
                            01:23:23:23:23:23:23:23:23:23
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Table 55 describes service FDB output fields.

Table 55 **Show FDB Information Fields**

Label	Description
ServID	Displays the service ID.
MAC	Displays the associated MAC address.
Source Identifier	Displays the id of the source MAC.

Table 55 **Show FDB Information Fields (Continued)**

Label	Description (Continued)
Type/Age	Type — Specifies the number of seconds used to age out TLS FDB entries learned on local SAPs. Age — Specifies the number of seconds used to age out TLS FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs. L — Learned - Dynamic entries created by the learning process. OAM — Entries created by the OAM process. H — Host, the entry added by the system for a static configured subscriber host. D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease. P — Indicates the MAC is protected by the MAC protection feature. Static — Statically configured.
Last Change	Indicates the time of the most recent state changes.

gsmp

Syntax **gsmp**

Context show>service>id

Description This command displays GSMP information.

neighbors

Syntax **neighbors group** [*name*] [*ip-address*]

Context show>service>id>gsmp

Description This command displays GSMP neighbor information.

Parameters **group** — A GSMP group defines a set of GSMP neighbors which have the same properties.

name — Specifies a GSMP group name is unique only within the scope of the service in which it is defined.

ip-address — Specifies the ip-address of the neighbor.

Output The following output is an example of service GSMP neighbor information.

Sample Output

These commands show the configured neighbors per service, regardless of the fact there exists an open TCP connection with this neighbor. The admin state is shown because for a neighbor to be admin enabled, the service, gsmp node, group node and the neighbor node in this service must all be in 'no shutdown' state. Session gives the number of session (open TCP connections) for each configured neighbor.

```
A:active>show>service>id>gsmp# neighbors
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
dslam1                         192.168.1.3        Enabled     0
-------------------------------------------------------------------------------
Number of neighbors shown: 2
===============================================================================
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslam1
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
dslam1                         192.168.1.3        Enabled     0
-------------------------------------------------------------------------------
Number of neighbors shown: 2
===============================================================================
A:active>show>service>id>gsmp#


A:active>show>service>id>gsmp# neighbors group dslam1 192.168.1.2
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
===============================================================================
A:active>show>service>id>gsmp#
```

sessions

Syntax **sessions** [**group** *name*] **neighbor** *ip-address*] [**port** *port-number*] [**association**] [**statistics**]

Context show>service>id>gsmp

Description This command displays GSMP sessions information.

Parameters **group** — A GSMP group defines a set of GSMP neighbors which have the same properties

name — Specifies a GSMP group name within the scope of the service in which it is defined

ip-address — Specifies the ip-address of the neighbor

port — Specifies the neighbor TCP port number use for this ANCP session

Values 0 to 65535

association — Displays to what object the ANCP-string is associated.

statistics — Displays statistics information about an ANCP session known to the system

Output The following output is an example of service GSMP sessions information.

Sample Output

This show command gives information about the open TCP connections with DSLAMs.

```
A:active>show>service>id>gsmp# sessions
=======================================================
GSMP sessions for service 999 (VPRN)
=======================================================
Port   Ngbr-IpAddr     Gsmp-Group
-------------------------------------------------------
40590  192.168.1.2     dslam1
-------------------------------------------------------
Number of GSMP sessions : 1
=======================================================
A:active>show>service>id>gsmp#


A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590
===============================================================================
GSMP sessions for service 999 (VPRN), neighbor 192.168.1.2, Port 40590
===============================================================================
State           : Established
Peer Instance   : 1                  Sender Instance : a3cf58
Peer Port       : 0                  Sender Port     : 0
Peer Name       : 12:12:12:12:12:12  Sender Name     : 00:00:00:00:00:00
timeouts        : 0                  Max. Timeouts   : 3
Peer Timer      : 100                Sender Timer    : 100
Capabilities      : DTD OAM
Conf Capabilities : DTD OAM
Priority Marking  : dscp nc2
Local Addr.       : 192.168.1.4
Conf Local Addr.  : N/A
===============================================================================
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
===============================================================================
ANCP-Strings
===============================================================================
ANCP-String                                                  Assoc. State
-------------------------------------------------------------------------------
No ANCP-Strings found
===============================================================================
A:active>show>service>id>gsmp#
```

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 statistics
===============================================================
GSMP session stats, service 999 neighbor 192.168.1.2, Port 40590
===============================================================
Event                                  Received  Transmitted
---------------------------------------------------------------
Dropped                                0         0
Syn                                    1         1
Syn Ack                                1         1
Ack                                    14        14
Rst Ack                                0         0
Port Up                                0         0
Port Down                              0         0
OAM Loopback                           0         0
===============================================================
A:active>show>service>id>gsmp#
```

Note: The association command gives an overview of each ANCP string received from this session.

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
===============================================================================
ANCP-Strings
===============================================================================
ANCP-String                                                    Assoc.
State
-------------------------------------------------------------------------------
7330-ISAM-E47 atm 1/1/01/01:19425.64048                        ANCP    Up
-------------------------------------------------------------------------------
Number of ANCP-Strings : 1
===============================================================================
A:active>show>service>id>gsmp#
```

Table 56 describes show sessions neighbor output fields.

Table 56 **Show Sessions Neighbor Output Fields**

Label	Description
State	The current state of the ANCP session.
Peer Instance	The instance number of the ANCP session at the neighbor's side.
Sender Instance	The instance number of the ANCP session at our side.
Peer Port	The port number of the ANCP session at the neighbor's side.
Sender Port	The port number of the ANCP session at the local side.
Peer Name	The MAC address of the ANCP session at the neighbor's side.

Table 56 **Show Sessions Neighbor Output Fields (Continued)**

Label	Description (Continued)
Sender name	The MAC address of the ANCP session at the local side.
timeouts	The number of adjacency protocol message timeouts.
Max. Timeouts	The maximum allowed of the above timeouts before closing.
Peer Timer	The timer value for the neighbor periodic adjacency protocol messages.
Sender Timer	The timer value for the local periodic adjacency protocol messages.
Capabilities	The negotiated capabilities for the Established ANCP session (DTD: dynamic topology discovery - OAM: operation and maintenance).
Conf Cap	The configured local capabilities.
Priority Marking	The DSCP bits for the IP messages used in the ANCP session.
Local Addr.	The destination IP address for this ANCP session.
Conf Local Addr.	The destination IP address accepted for ANCP connections.

host

Syntax **host** [**sap** *sap-id*] [**detail**]
host summary

Context show>service>id

Description This command displays static host information configured on this service.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

summary — Displays summary host information

host-connectivity-verify

Syntax **host-connectivity-verify statistics** [**sap** *sap-id*]

Context show>service>id

Description This command displays host connectivity check statistics.

Parameters **statistics** — Displays host connectivity verification data

sap-id — Specifies the physical port identifier portion of the SAP definition

Output The following output is an example of service host connectivity information.

Sample Output

```
A:ALA-48>show>service>id# host-connectivity-verify statistics sap 1/1/9:0
===============================================================================
Host connectivity check statistics
===============================================================================
Svc    SapId/          DestIp          Last                  Time        Oper
Id     SdpId           Address         Response              Expired     State
-------------------------------------------------------------------------------
10001/2/3:0143.144.145.1                                   Up
===============================================================================
A:ALA-48>show>service>id#
```

Table 57 describes show service-id host connectivity verification output fields.

Table 57 **Show Service Id Host Connectivity Verify Fields**

Label	Description
Svc Id	The service identifier.
SapId/SdpId	The SAP and SDP identifiers.
DestIp Address	The destination IP address.
Last Response	The time when the last response was received.
Time Expired	Displays whether the interval value has expired.
Oper State	Displays the current operational state of the service.

i-vpls

Syntax **i-vpls**

Context show>service>id

Description Displays I-VPLS services associated with the B-VPLS service. This command only applies when the service is a B-VPLS.

Output The following output is an example of service I-VPLS information.

Sample Output

```
*A:SetupCLI# show service id 2002 i-vpls
===============================================================
Related iVpls services for bVpls service 2002
===============================================================
iVpls SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
```

```
2001                122                 Up                  Down
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------

*A:term17>show>service>id# i-vpls
===============================================================================
Related iVpls services for bVpls service 2000
===============================================================================
iVpls SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
2100                2100                Up                  Up
2110                123                 Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

isid-using

Syntax **isid-using** [*ISID*]

Context show>service

Description This command displays services using an ISID.

Parameters *ISID* — Specifies a 24 bit (0 to 16777215) service instance identifier for this service. As part of the Provider Backbone Bridging frames, it is used at the destination PE as a demultiplexor field.

Values 0 to 16777215

Output The following output is an example of services using ISID information.

Sample Output

```
*A:SetupCLI# show service isid-using
===============================================================
Services
===============================================================
SvcId      ISID     Type   b-Vpls     Adm  Opr  SvcMtu CustId
---------------------------------------------------------------
2001       122      i-VPLS 2002       Up   Down 1514   1
2005       2005     i-mVP* 2004       Down Down 1500   1
---------------------------------------------------------------
Matching Services : 2
-===============================================================
*A:SetupCLI#

A:term17# show service isid-using
===============================================================
Services
===============================================================
SvcId      ISID     Type   b-Vpls     Adm  Opr  SvcMtu CustId
---------------------------------------------------------------
```

```
2000       0        b-VPLS 0          Up   Up   1530   1
2110       123      i-VPLS 2000       Up   Up   1514   1
2299       0        b-VPLS 0          Down Down 1514   1
------------------------------------------------------------------
Matching Services : 3
------------------------------------------------------------------
A:term17#
```

labels

Syntax **labels**

Context show>service>id

Description This command displays the labels being used by the service.

Output The following output is an example of service label information.

Sample Output

```
*A:ALA-12# show service id 1 labels
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          40:1               Mesh 130081       131061
1          60:1               Mesh 131019       131016
1          100:1              Mesh 0            0
-------------------------------------------------------------------------------
Number of Bound SDPs : 6
-------------------------------------------------------------------------------
*A:ALA-12#
```

Table 58 describes show service-id labels output fields.

Table 58 **Show Service-ID Labels Fields**

Label	Description
Svc Id	The service identifier.
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is spoke or mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.

Table 58 **Show Service-ID Labels Fields (Continued)**

Label	Description
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.

l2pt

Syntax **l2pt disabled**
l2pt [**detail**]

Context show>service>id

Description This command displays Layer 2 Protocol Tunnel (L2-PT) route information associated with this service.

Parameters **disabled** — Displays only entries with termination disabled. This helps identify configuration errors.

detail — Displays detailed information.

Output The following output is an example of service L2PT information.

Sample Output

```
A:ALA-48>show>service>id# l2pt
===============================================================================
L2pt summary, Service id 700
===============================================================================
       L2pt-term   L2pt-term   Bpdu-trans  Bpdu-trans  Bpdu-trans  Bpdu-trans
       enabled     disabled    auto        disabled    pvst        stp
-------------------------------------------------------------------------------
SAP's  0           1           0           1           0           0
SDP's  0           1           0           1           0           0
-------------------------------------------------------------------------------
Total  0           2           0           2           0           0
===============================================================================
A:ALA-48>show>service>id#


A:ALA-48>show>service>id# l2pt disabled
===============================================================================
L2pt details, Service id 700
===============================================================================
Service Access Points
-------------------------------------------------------------------------------
SapId          L2pt-                        Admin Bpdu-    Oper Bpdu-
               termination                  translation    translation
-------------------------------------------------------------------------------
1/1/9:0        disabled                     disabled       disabled
-------------------------------------------------------------------------------
Number of SAPs : 1
```

```
Service Destination Points
-------------------------------------------------------------------------------
SdpId          L2pt-                          Admin Bpdu-    Oper Bpdu-
               termination                    translation    translation
-------------------------------------------------------------------------------
2:222          disabled                       disabled       disabled
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
L2pt summary, Service id 700
===============================================================================
       L2pt-term   L2pt-term   Bpdu-trans  Bpdu-trans  Bpdu-trans  Bpdu-trans
       enabled     disabled    auto        disabled    pvst        stp
-------------------------------------------------------------------------------
SAP's  0           1           0           1           0           0
SDP's  0           1           0           1           0           0
-------------------------------------------------------------------------------
Total  0           2           0           2           0           0
===============================================================================
A:ALA-48>show>service>id#


A:ALA-48>show>service>id# l2pt detail
===============================================================================
L2pt details, Service id 700
===============================================================================
Service Access Points
-------------------------------------------------------------------------------
SapId          L2pt-                          Admin Bpdu-    Oper Bpdu-
               termination                    translation    translation
-------------------------------------------------------------------------------
1/1/9:0        disabled                       disabled       disabled
-------------------------------------------------------------------------------
Number of SAPs : 1

Service Destination Points
-------------------------------------------------------------------------------
SdpId          L2pt-                          Admin Bpdu-    Oper Bpdu-
               termination                    translation    translation
-------------------------------------------------------------------------------
2:222          disabled                       disabled       disabled
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
L2pt summary, Service id 700
===============================================================================
       L2pt-term   L2pt-term   Bpdu-trans  Bpdu-trans  Bpdu-trans  Bpdu-trans
       enabled     disabled    auto        disabled    pvst        stp
-------------------------------------------------------------------------------
SAP's  0           1           0           1           0           0
SDP's  0           1           0           1           0           0
-------------------------------------------------------------------------------
Total  0           2           0           2           0           0
===============================================================================
A:ALA-48>show>service>id#
```

Table 59 describes show L2PT output fields.

Table 59 **Show L2PT Fields**

Label	Description
Service id	Displays the 24 bit (0 to 16777215) service instance identifier for the service.
L2pt-term enabled	Indicates if L2-PT-termination and/or Bpdu-translation is in use in this service by at least one SAP or spoke-SDP binding. If in use, at least one of L2PT-termination or Bpdu-translation is enabled. When enabled it is not possible to enable STP on this service.
L2pt-term disabled	Indicates that L2-PT-termination is disabled.
Bpdu-trans auto	Specifies the number of L2-PT PDUs are translated before being sent out on a port or sap.
Bpdu-trans disabled	Indicates that Bpdu-translation is disabled.
SAPs	Displays the number of SAPs with L2PT or BPDU translation enabled or disabled.
SDPs	Displays the number of SDPs with L2PT or BPDU translation enabled or disabled.
Total	Displays the column totals of L2PT entities.
SapId	The ID of the access point where this SAP is defined.
L2pt-termination	Indicates whether L2pt termination is enabled or disabled.
Admin Bpdu-translation	Specifies whether Bpdu translation is administratively enabled or disabled.
Oper Bpdu-translation	Specifies whether Bpdu translation is operationally enabled or disabled.
SdpId	Specifies the SAP ID.

mac-move

Syntax **mac-move**

Context show>service>id

Description This command displays MAC move related information about the service.

Output The following output is an example of service MAC move information.

Sample Output

```
*A:ALA-2009>config>service>vpls>mac-move# show service id 500 mac-move
===============================================================================
Service Mac Move Information
===============================================================================
Service Id        : 500                 Mac Move          : Enabled
Primary Factor    : 4                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
-------------------------------------------------------------------------------
SAP Mac Move Information: 2/1/3:501
-------------------------------------------------------------------------------
Admin State       : Up                  Oper State        : Down
Flags              : RelearnLimitExceeded
Time to come up   : 1 seconds           Retries Left      : 1
Mac Move          : Blockable           Blockable Level   : Tertiary
-------------------------------------------------------------------------------
SAP Mac Move Information: 2/1/3:502
-------------------------------------------------------------------------------
Admin State       : Up                  Oper State        : Up
Flags              : None
Time to RetryReset: 267 seconds         Retries Left      : none
Mac Move          : Blockable           Blockable Level   : Tertiary
-------------------------------------------------------------------------------
SDP Mac Move Information: 21:501
-------------------------------------------------------------------------------
Admin State       : Up                  Oper State        : Up
Flags              : None
Time to RetryReset: never               Retries Left      : 3
Mac Move          : Blockable           Blockable Level   : Secondary
-------------------------------------------------------------------------------
SDP Mac Move Information: 21:502
-------------------------------------------------------------------------------
Admin State       : Up                  Oper State        : Down
Flags              : RelearnLimitExceeded
Time to come up   : never               Retries Left      : none
Mac Move          : Blockable           Blockable Level   : Tertiary
===============================================================================
*A:ALA-2009>config>service>vpls>mac-move#
```

mac-protect

Syntax **mac-protect**

Context show>service>id

Description This command displays MAC protect-related information about the service.

Output The following output is an example of service MAC protect information.

Sample Output

```
*A:ALA-48>show>service>id# mac-protect
===============================================================================
Protected MACs, Service 700
===============================================================================
```

```
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
700       ff:ff:ff:ff:ff:ff not learned           n/a       n/a
-------------------------------------------------------------------------------
No. of MAC Entries: 1
===============================================================================
*A:ALA-48>show>service>id# mac-protect
```

mld-snooping

Syntax **mld-snooping**

Context show>service>id

Description This command displays MLD snooping information.

all

Syntax **all**

Context show>service>id>mld-snooping

Description This command displays detailed information about MLD snooping.

base

Syntax **base**

Context show>service>id>mld-snooping

Description This command displays basic MLD snooping information.

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>mld-snooping

Description This command displays all multicast routers.

mvr

Syntax **mvr**

Context show>service>id>mld-snooping

Description This command displays multicast VPLS registration information.

port-db

Syntax **port-db sap** *sap-id*
port-db sap *sap-id* **detail**
port-db sap *sap-id* **group** *grp-ipv6-address*
port-db sdp *sdp-id:vc-id* [**detail**]
port-db sdp *sdp-id:vc-id* **group** *grp-ipv6-address*

Context show>service>id>mld-snooping

Description This command displays MLD snooping information related to a specific SAP.

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-ip-address*

Context show>service>id>mld-snooping

Description This command displays proxy-reporting database entries.

Parameters *grp-ip-address* — Displays the IGMP snooping proxy reporting database for a specific multicast group address.

detail — Displays detailed information about the proxy-reporting database,

querier

Syntax **querier**

Context show>service>id>mld-snooping

Description This command displays information about the current querier.

static

Syntax **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>mld-snooping

Description This command displays MLD snooping static group membership data.

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>mld-snooping

Description This command displays MLD snooping statistics.

mrp-policy

Syntax **mrp-policy** [*mrp-policy*]
mrp-policy *mrp-policy* [**association**]
mrp-policy *mrp-policy* [**entry** *entry-id*]

Context show>service>id

Description This command displays information on an MRP policy.

Parameters *mrp-policy* — Specifies the MRP policy name

Values 32 chars max

entry-id — Specifies the entry ID number

Values 1 to 65535

Output The following output is an example of service MRP policy information.

Sample Output

```
*A:PE-B# show service mrp-policy
===============================================================================
Mrp Policies
===============================================================================
Mrp-Policy                      Scope     Applied Description
-------------------------------------------------------------------------------
1                               template  Yes
2                               template  Yes
-------------------------------------------------------------------------------
Total: 2
===============================================================================


*A:PE-B# show service mrp-policy "1"
===============================================================================
Mrp Policy
===============================================================================
Policy Name : 1                                 Applied         : Yes
Scope       : template                          Def. Action     : block
Entries     : 1
Description : (Not Specified)
-------------------------------------------------------------------------------
Mrp Policy Entries
-------------------------------------------------------------------------------
```

```
Entry       : 1                                 Match action    : end-station
Description : (Not Specified)
isid        : 10..11
===============================================================================
*A:PE-B#
```

mmrp

Syntax **mmrp mac** [*ieee-address*]

Context show>service>id

Description This command displays information on MACs. If a MAC address is specified, information will be displayed relevant to the specific group. No parameter will display information on all group MACs on a server.

Parameters *ieee-address* — Displays hex string information

Output The following output is an example of service MMRP MAC information.

Sample Output

```
*A:PE-A# show service id 10 mmrp mac 01:1E:83:00:00:65
-------------------------------------------------------------------------------
SAP/SDP                                 MAC Address       Registered  Declared
-------------------------------------------------------------------------------
sap:1/1/4:10                            01:1e:83:00:00:65 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:65 No          Yes
sap:2/2/5:10                            01:1e:83:00:00:65 Yes         Yes
-------------------------------------------------------------------------------
*A:PE-A#

*A:PE-A# show service id 10 mmrp mac
-------------------------------------------------------------------------------
SAP/SDP                                 MAC Address       Registered  Declared
-------------------------------------------------------------------------------
sap:1/1/4:10                            01:1e:83:00:00:65 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:66 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:67 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:68 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:69 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6a No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6b No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6c No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6d No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6e No          Yes
sap:1/2/2:10                            01:1e:83:00:00:65 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:66 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:67 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:68 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:69 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6a No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6b No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6c No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6d No          Yes
```

```
sap:1/2/2:10                            01:1e:83:00:00:6e No          Yes
sap:2/2/5:10                            01:1e:83:00:00:65 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:66 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:67 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:68 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:69 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6a Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6b Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6c Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6d Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6e Yes         Yes
-------------------------------------------------------------------------------
*A:PE-A#
```

mstp-configuration

Syntax **mstp-configuration**

Context show>service>id

Description This command displays the MSTP specific configuration data. This command is only valid on a management VPLS.

pim-snooping

Syntax **pim-snooping**

Context show>service>id

Description This command enables the context to display PIM snooping information.

group

Syntax **group** [*grp-ip-address*] [**source** *ip-address*] [**type** {**starg** | **sg**}] [**detail**] [*family*]

Context show>service>id>pim-snooping

Description This command displays the multicast group information.

Parameters *grp-ip-address* — Specifies the IP multicast group address for which this entry contains information

ip-address — Specifies the source address for which this entry contains information.

starg — Specifies that only (*,G) entries be displayed

sg — Specifies that only (S,G) entries be displayed

detail — Displays detailed group information.

family — Displays either IPv4 or IPv6 information.

Values ipv4 or ipv6

Output The following output is an example of service PIM snooping information.

Sample Output

```
*A:PE# show service id 1 pim-snooping group
===============================================================================
PIM Snooping Groups ipv4
===============================================================================
Group Address           Source Address          Type     Incoming          Num
                                                         Intf              Oifs
-------------------------------------------------------------------------------
233.252.0.1             10.0.0.2                (S,G)    SAP:1/1/2         2
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
*A:PE#
```

neighbor

Syntax **neighbor** [{**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**address** *ip-address*]] [**detail**] [*family*]

Context show>service>id>pim-snooping

Description This command displays PIM neighbor information.

Parameters *ip-int-name* — Only displays the interface information associated with the specified IP interface name

sap-id — Displays the neighbor information associated with the specified SAP

sdp-id:vc-id — Displays the neighbor information associated with the specified SDP

ip-address — Displays information for the neighbor with the specified IP address

detail — Displays detailed neighbor information

family — Displays either IPv4 or IPv6 information for the specified neighbor

Values ipv4 or ipv6

Output The following output is an example of service PIM snooping neighbor information.

Sample Output

```
*A:PE#  show service id 1 pim-snooping neighbor
===============================================================================
PIM Snooping Neighbors ipv4
===============================================================================
Port Id                 Nbr DR Prty     Up Time       Expiry Time     Hold Time
  Nbr Address
-------------------------------------------------------------------------------
```

```
SAP:1/1/9:1             1               0d 00:05:23   0d 00:01:23     105
  10.0.0.1
EVPN-MPLS               1               0d 00:05:02   0d 00:01:43     105
  10.0.0.2
EVPN-MPLS               1               0d 00:05:02   0d 00:01:43     105
  10.0.0.3
-------------------------------------------------------------------------------
Neighbors : 3
===============================================================================
```

port

Syntax **port** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**group** [*grp-ip-address*]] [**detail**] [*family*]

Context show>service>id>pim-snooping

Description This command displays PIM port information.

Parameters *sap-id* — Displays the port information associated with the specified SAP

sdp-id:vc-id — Displays the port information associated with the specified SDP

grp-ip-address — Specifies the IP multicast group address for which this entry contains information

detail — Displays detailed port information

family — Displays either IPv4 or IPv6 information for the specified port

Values ipv4 or ipv6

Output The following output is an example of service PIM snooping information.

Sample Output

```
*A:PE# show service id 1 pim-snooping port
===============================================================================
PIM Snooping Ports ipv4
===============================================================================
Sap/Sdp Id                          Opr
-------------------------------------------------------------------------------
SAP:1/1/1                           Up
SAP:1/1/2                           Up
===============================================================================
*A:PE#
```

statistics

Syntax **statistics** [**sap** *sap-id*] [**sdp** *sdp-id:vc-id*] [*family*]

Context show>service>id>pim-snooping

Description This command displays PIM statistics information.

Parameters *sap-id* — Displays the statistics associated with the specified SAP

sdp-id:vc-id — Displays the statistics associated with the specified SDP

family — Displays either IPv4 or IPv6 statistics

Values ipv4 or ipv6

Output The following output is an example of service PIM snooping statistics information.

Sample Output

```
*A:PE# show service id 1 pim-snooping statistics
===============================================================
PIM Snooping Statistics ipv4
===============================================================
Message Type        Received       Transmitted    Rx Errors
---------------------------------------------------------------
Hello               36             -              0
Join Prune          8              8              0
Total Packets       44             8
-------------------------------------------------------------------------------
General Statistics
-------------------------------------------------------------------------------
Rx Neighbor Unknown               : 0
Rx Bad Checksum Discard           : 0
Rx Bad Encoding                   : 0
Rx Bad Version Discard            : 0
Join Policy Drops                 : 0
-------------------------------------------------------------------------------
Source Group Statistics
-------------------------------------------------------------------------------
(S,G)                             : 1
(*,G)                             : 0
===============================================================
*A:PE#
```

status

Syntax **status** [*family*]

Context show>service>id>pim-snooping

Description This command displays PIM status information.

Parameters *family* — Displays either IPv4 or IPv6 status information

Values ipv4 or ipv6

Output The following output is an example of service PIM snooping status information.

Sample Output

```
*A:PE# show service id 1 pim-snooping status
===============================================================================
PIM Snooping Status ipv4
===============================================================================
Admin State                      : Up
Oper State                       : Up
Mode Admin                       : Proxy
Mode Oper                        : Proxy
Hold Time                        : 90
Designated Router                : 10.0.1.2
J/P Tracking                     : Inactive
Up Time                          : 0d 00:08:43
Group Policy                     : None
===============================================================================
*A:PE#
```

provider-tunnels

Syntax **provider-tunnels**

Context show>service>id

Description This command displays the service provider tunnel information.

Output The following output is an example of service provider tunnel information.

Sample Output

```
*A:Dut-B# show service id 1 provider-tunnel

===============================================================================
Service Provider Tunnel Information
===============================================================================
Type               : inclusive          Root and Leaf      : enabled
Admin State        : inService          Data Delay Intvl   : 3 secs
PMSI Type          : ldp                LSP Template       :
Remain Delay Intvl : 0 secs             LSP Name used      : 8193
===============================================================================
*A:Dut-B# /tools dump service id 1 provider-tunnels type originating

===============================================================================
VPLS 1 Inclusive Provider Tunnels Originating
===============================================================================
ipmsi (LDP)                                      P2MP-ID  Root-Addr
-------------------------------------------------------------------------------
8193                                             8193    10.20.1.2

-------------------------------------------------------------------------------
*A:Dut-B# /tools dump service id 1 provider-tunnels type terminating

===============================================================================
VPLS 1 Inclusive Provider Tunnels Terminating
===============================================================================
ipmsi (LDP)                                      P2MP-ID  Root-Addr
-------------------------------------------------------------------------------
```

```
                                                  8193    10.20.1.3
                                                  8193    10.20.1.4
                                                  8193    10.20.1.6
                                                  8193    10.20.1.7
-------------------------------------------------------------------------
*A:Dut-B# /tools dump service id 1 provider-tunnels
=========================================================================
VPLS 1 Inclusive Provider Tunnels Originating
=========================================================================
ipmsi (LDP)                                       P2MP-ID  Root-Addr
-------------------------------------------------------------------------
8193                                              8193    10.20.1.2
-------------------------------------------------------------------------
=========================================================================
VPLS 1 Inclusive Provider Tunnels Terminating
=========================================================================
ipmsi (LDP)                                       P2MP-ID  Root-Addr
-------------------------------------------------------------------------
                                                  8193    10.20.1.3
                                                  8193    10.20.1.4
                                                  8193    10.20.1.6
                                                  8193    10.20.1.7
-------------------------------------------------------------------------

*A:Dut-B# show service id 1 provider-tunnel
=========================================================================
Service Provider Tunnel Information
=========================================================================
Type               : inclusive          Root and Leaf      : enabled
Admin State        : inService          Data Delay Intvl   : 3 secs
PMSI Type          : ldp                LSP Template       :
Remain Delay Intvl : 0 secs             LSP Name used      : 8193
=========================================================================

*A:Dut-C# show service id 1001 provider-tunnel
=========================================================================
Service Provider Tunnel Information
=========================================================================
Type               : inclusive          Root and Leaf      : enabled

Admin State        : inService          Data Delay Intvl   : 3 secs

PMSI Type          : rsvp               LSP Template       : ipmsi

Remain Delay Intvl : 0 secs             LSP Name used      : ipmsi-1001-73728
=========================================================================
```

proxy-arp

Syntax **proxy-arp** [**ip-address** *ip-address*] [**detail**]

Context show>service

Description This command displays the proxy-ARP entries existing for a particular service. A 7750 SR, 7450 ESS or 7950 XRS router receiving an ARP request from a SAP or SDP-binding will perform a lookup in the proxy-arp table for the service. If the router finds a match, it will reply to the ARP and will not let the ARP be flooded in the VPLS service. If the router does not find a match, the ARP will be flooded within the service. The command allows for specific IP addresses to be shown.

The **detail** parameter allows the user to display all the entries. An individual IP address entry can also be shown.

Output The following output is an example of service proxy ARP information.

Sample Output

```
:PE71(1)# show service id 600 proxy-arp
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : 200 secs            Send Refresh      : 120 secs

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 3
Hold down         : max
Anti Spoof MAC    : 00:ca:ca:ca:ca:ca

EVPN
-------------------------------------------------------------------------------
Garp Flood        : disabled            Req Flood         : disabled
-------------------------------------------------------------------------------
A:PE71(1)# show service id 600 proxy-arp detail
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : 200 secs            Send Refresh      : 120 secs

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 3
Hold down         : max
Anti Spoof MAC    : 00:ca:ca:ca:ca:ca

EVPN
-------------------------------------------------------------------------------
Garp Flood        : disabled            Req Flood         : disabled
-------------------------------------------------------------------------------

===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
```

```
172.16.0.1          00:ca:fe:ca:fe:02   evpn      active    12/01/2014 12:02:27
172.16.0.61         00:ca:de:ba:ca:00   dyn       active    12/01/2014 15:40:10
172.16.0.100        00:00:00:00:00:01   stat      inActv    12/01/2014 12:01:57
172.16.0.102        00:00:00:00:00:02   stat      inActv    12/01/2014 12:01:57
-------------------------------------------------------------------------------
Number of entries : 4
===============================================================================
A:PE71(1)#
```

proxy-nd

Syntax **proxy-nd** [**ip-address** *ip-address*] [**detail**]

Context show>service

Description This command displays the information about the proxy ND settings configured in a specified service. The **detail** parameter allows the user to display all the entries. An individual IP address entry can also be shown.

Output The following output is an example of service proxy ND information.

Sample Output

```
:PE71(1)# show service id 600 proxy-nd
-------------------------------------------------------------------------------
Proxy nd
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : 200 secs            Send Refresh      : 120 secs

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 3
Hold down         : max
Anti Spoof MAC    : 00:ca:ca:ca:ca:ca

EVPN
-------------------------------------------------------------------------------
Garp Flood        : disabled            Req Flood         : disabled
-------------------------------------------------------------------------------
A:PE71(1)# show service id 600 proxy-nd detail
-------------------------------------------------------------------------------
Proxy nd
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : 200 secs            Send Refresh      : 120 secs

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 3
Hold down         : max
Anti Spoof MAC    : 00:ca:ca:ca:ca:ca
```

```
EVPN
------------------------------------------------------------------------------
Garp Flood        : disabled            Req Flood         : disabled
------------------------------------------------------------------------------

==============================================================================
VPLS Proxy ND Entries
==============================================================================
IP Address          Mac Address         Type      Status    Last Update
------------------------------------------------------------------------------
172.16.0.1          00:ca:fe:ca:fe:02   evpn      active    12/01/2014 12:02:27
172.16.0.61         00:ca:de:ba:ca:00   dyn       active    12/01/2014 15:40:10
172.16.0.100        00:00:00:00:00:01   stat      inActv    12/01/2014 12:01:57
172.16.0.102        00:00:00:00:00:02   stat      inActv    12/01/2014 12:01:57
------------------------------------------------------------------------------
Number of entries : 4
==============================================================================
A:PE71(1)#
```

retailers

Syntax **retailers**

Context show>service>id

Description This command displays the service ID of the retailer subscriber service to which this DHCP lease belongs.

wholesalers

Syntax **wholesalers**

Context show>service>id

Description This command displays service wholesaler information.

vxlan

Syntax **vxlan**

Context show>service>id
show>service

Description This command displays the VXLAN bindings auto-created in a specified service. A VXLAN binding is composed of the remote VTEP (VXLAN Termination Endpoint) and the corresponding egress VNI (VXLAN Network Identifier) to identify the service at the egress node. The command shows the number of MACs associated to each binding as well as the operational status if the binding is part of the multicast list. The binding will be operationally down when the VTEP address is not found in the base routing table (the VTEP address cannot be reached). A binding will be part of the multicast list if a valid BGP EVPN inclusive multicast route exists for it.

Output The following output is an example of service VXLAN information.

Sample Output

```
*A:DutA# show service id 1 vxlan
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 1

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address           Egress VNI     Num. MACs    In Mcast List?  Oper State
-------------------------------------------------------------------------------
192.0.0.71             1              1            Yes             Up
192.0.0.72             1              0            Yes             Up
192.0.0.74             1              0            Yes             Up
192.0.0.76             1              1            Yes             Down
192.168.45.2           1              0            Yes             Down
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 5
-------------------------------------------------------------------------------
===============================================================================
A:DutB# show service vxlan
<vtep>
 192.0.2.65   192.0.2.66

A:PE63# show service vxlan 192.0.2.65
===============================================================================
VXLAN Tunnel Endpoint: 192.0.2.65
===============================================================================
Egress VNI                    Service Id               Oper State
-------------------------------------------------------------------------------
60                            60                       Up
-------------------------------------------------------------------------------
```

sap

Syntax **sap** [*sap-id* [*filter*] [**detail**]]

Context show>service>id

Description This command displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters *sap-id* — The ID that displays SAPs for the service in the *slot*/*mda*/*port*[.*channel*] form

detail — Displays detailed information for the SAP

filter — Specifies a search term to narrow down the results. This parameter applies to the 7450 ESS or 7750 SR only.

Values atm, base, detail, dhcp, mc-ring, mcac, mrp, qos, sap-stats, stats, stp, sub-mgmt

Output The following output is an example of service SAP information.

Sample Output

```
*A:PE# show service id 1 sap 1/1/1:1 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:1                  Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 01/29/2015 10:51:49
Last Mgmt Change   : 01/28/2015 11:48:21
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Etree Root Leaf Tag: Disabled                 Etree Leaf Tag    : 0
Etree Leaf AC      : Disabled
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
OAM MAC Addr       : 0                        DHCP MAC Addr     : 0
Host MAC Addr      : 0                        Intf MAC Addr     : 0
SPB MAC Addr       : 0                        Cond MAC Addr     : 0
BGP EVPN Addr      : 0                        EVPN Static Addr  : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
                                              Egr Agg Rate Limit: max
Q Frame-Based Acct : Disabled                 Limit Unused BW   : Disabled
ARP Reply Agent    : Disabled                 Host Conn Verify  : Disabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Dynamic Hosts     : Enabled
Avl Static Hosts   : 0                        Tot Static Hosts  : 0
Calling-Station-Id : n/a
```

```
Application Profile: None
Transit Policy     : None

Oper Group         : (none)                    Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Cflowd             : Disabled
MCAC Policy Name   :                           MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                  MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                         MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                         MCAC Avail Opnl BW: unlimited
Use LAG port weight: no
Restr MacProt Src  : Disabled                  Restr MacUnpr Dst : Disabled
Auto Learn Mac Prot: Disabled                  RestProtSrcMacAct : Disable
Time to RetryReset : never                     Retries Left      : 3
Mac Move           : Blockable                 Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy        : None
-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : n/a                       AIS               : Disabled
MC Prop-Hold-Timer : n/a                       V-MEP Filtering   : Disabled
Squelch Levels     : None
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Forwarding
Port Number        : N/A                       Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                       Designated Port   : N/A

Forward transitions: 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                         TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                         MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                         Min Auth Interval : 15 minutes
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                         Egress qos-policy : 30
Ingress FP QGrp    : gq1                       Egress Port QGrp  : gq1
Ing FP QGrp Inst   : 1                         Egr Port QGrp Inst: 1
Shared Q plcy      : n/a                       Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : test2
I. Policer Ctl Pol : (Not Specified)
```

```
E. Policer Ctl Pol : (Not Specified)
I. QGrp Redir. List: list1
E. QGrp Redir. List: list1
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down                    Lease Populate    : 0
DHCP Snooping      : Down                    Action            : Keep

Proxy Admin State  : Down
Proxy Lease Time   : N/A
Emul. Server Addr  : Not Configured
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Down                    MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : None
Def SLA-Profile    : None
Def Inter-Dest-Id  : None
Def App-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
Prof Traffic Only : False
Non-Sub-Traffic   : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
CPM Ingress           : 0                       0

Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 30)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets
```

```
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
Off. Managed          : 0                        0
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0

Egress Queue 1
For. InProf           : 0                        0
For. OutProf          : 0                        0
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
===============================================================================
*A:PE#

*A:PE-A# show service id 10 sap 2/2/5:10 mrp
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 10
SAP                : 2/2/5:10                 Encap             : q-tag
Description        : Default sap description for service id 10
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 01/16/2008 09:37:57
Last Mgmt Change   : 01/16/2008 09:37:41
-------------------------------------------------------------------------------
SAP MRP Information
-------------------------------------------------------------------------------
Join Time          : 0.2 secs                 Leave Time        : 1.0 secs
Leave All Time     : 10.0 secs                Periodic Time     : 1.0 secs
Periodic Enabled   : false
Rx Pdus            : 11                       Tx Pdus           : 12
Dropped Pdus       : 0                        Tx Pdus           : 12
Rx New Event       : 0                        Rx Join-In Event  : 150
Rx In Event        : 10                       Rx Join Empty Evt : 10
Rx Empty Event     : 10                       Rx Leave Event    : 0
Tx New Event       : 0                        Tx Join-In Event  : 140
Tx In Event        : 0                        Tx Join Empty Evt : 20
Tx Empty Event     : 10                       Tx Leave Event    : 0
-------------------------------------------------------------------------------
SAP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
01:1e:83:00:00:65 Yes               Yes
01:1e:83:00:00:66 Yes               Yes
01:1e:83:00:00:67 Yes               Yes
```

```
01:1e:83:00:00:68 Yes               Yes
01:1e:83:00:00:69 Yes               Yes
01:1e:83:00:00:6a Yes               Yes
01:1e:83:00:00:6b Yes               Yes
01:1e:83:00:00:6c Yes               Yes
01:1e:83:00:00:6d Yes               Yes
01:1e:83:00:00:6e Yes               Yes
-------------------------------------------------------------------------------
Number of MACs=10 Registered=10 Declared=10
-------------------------------------------------------------------------------
*A:PE-A#
```

Table 60 describes show service SAP fields.

Table 60 **Show Service-ID SAP Fields**

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operational state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch,RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown,
	NoSapIpipeCeIpAddr, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	Specifies the time of the most recent operating status change to this SAP.
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Admin MTU	The largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.

Table 60 **Show Service-ID SAP Fields (Continued)**

Label	Description (Continued)
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Forwarding Engine Stats	
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, etc.
Received Valid	The number of valid packets and octets received on the SAP.
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy.
Queuing Stats (Ingress QoS Policy)	
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped due to: MBS exceeded, buffer pool limit exceeded, etc.
Dro. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, dropped due to: MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded.
For. OutProf	The number of out-of-profile packets and octets discarded due to MBS exceeded, buffer pool limit exceeded, etc.
Queuing Stats (Egress QoS Policy)	

Table 60 **Show Service-ID SAP Fields (Continued)**

Label	Description (Continued)
Dro. InProf	The number of in-profile packets and octets discarded due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. OutProf	The number of out-of-profile packets and octets due to MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded.
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded.
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.
Alarm Cell Handling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.

sdp

Syntax **sdp** *sdp-id:vc-id* {**mrp**}
sdp [*sdp-id* | **far-end** *ip-addr*] [**detail**]

Context show>service>id

Description This command displays information for the SDPs associated with the service. If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Displays only information for the specified SDP ID

Default All SDPs

Values 1 to 17407

ip-addr — Displays only SDPs matching with the specified far-end IP address

Default SDPs with any far-end IP address

detail — Displays detailed SDP information.

Output The following output is an example of service SDP information.

Sample Output

```
*A:Dut-C# show service id 1001 sdp 17407:4294967295 detail
===============================================================================
Service Destination Point (Sdp Id : 17407:4294967295) Details
```

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

-------------------------------------------------------------------------------

Sdp Id 17407:4294967295  -(0.0.0.0)

-------------------------------------------------------------------------------

Description     : (Not Specified)

SDP Id             : 17407:4294967295         Type            : VplsPmsi

Split Horiz Grp    : (Not Specified)

VC Type            : Ether                    VC Tag            : n/a

Admin Path MTU     : 9194                     Oper Path MTU     : 9194

Far End            : not applicable           Delivery          : MPLS

Tunnel Far End     : n/a                    LSP Types         : None

Hash Label         : Disabled               Hash Lbl Sig Cap  : Disabled

Oper Hash Label    : Disabled


Admin State        : Up                     Oper State        : Up

Acct. Pol          : None                   Collect Stats     : Disabled

Ingress Label      : 0                        Egress Label      : 3

Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a

Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a

Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a

Admin ControlWord  : Not Preferred            Oper ControlWord  : False

Last Status Change : 01/31/2012 00:51:46      Signaling         : None

Last Mgmt Change   : 01/31/2012 00:49:58      Force Vlan-Vc     : Disabled

Endpoint           : N/A                      Precedence        : 4

PW Status Sig      : Enabled

Class Fwding State : Down

Flags              : None

Time to RetryReset : never                    Retries Left      : 3

Mac Move           : Blockable                Blockable Level   : Tertiary

Local Pw Bits      : None
```

```
Peer Pw Bits       : None

Peer Fault Ip      : None

Application Profile: None

Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0

Learned MAC Addr   : 0                       Static MAC Addr   : 0


MAC Learning       : Enabled                 Discard Unkwn Srce: Disabled

MAC Aging          : Enabled

BPDU Translation   : Disabled

L2PT Termination   : Disabled

MAC Pinning        : Disabled

Ignore Standby Sig : False                   Block On Mesh Fail: False

Oper Group         : (none)                  Monitor Oper Grp  : (none)

Rest Prot Src Mac  : Disabled

Auto Learn Mac Prot: Disabled                RestProtSrcMacAct : Disable


Ingress Qos Policy : (none)                  Egress Qos Policy : (none)

Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)

Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)


-------------------------------------------------------------------------------

ETH-CFM SDP-Bind specifics

-------------------------------------------------------------------------------

V-MEP Filtering    : Disabled

KeepAlive Information :

Admin State        : Disabled                Oper State        : Disabled

Hello Time         : 10                      Hello Msg Len     : 0

Max Drop Count     : 3                       Hold Down Time    : 10


Statistics           :
```

```
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0

I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0

E. Fwd. Pkts.      : 5937639                  E. Fwd. Octets    : 356258340

MCAC Policy Name   :

MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit

MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited

MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited


-------------------------------------------------------------------------------

RSVP/Static LSPs

-------------------------------------------------------------------------------
Associated LSP List :

No LSPs Associated


-------------------------------------------------------------------------------
Class-based forwarding :

-------------------------------------------------------------------------------

Class forwarding   : Disabled                 EnforceDSTELspFc  : Disabled

Default LSP        : Uknwn                    Multicast LSP     : None

===============================================================================
FC Mapping Table

===============================================================================

FC Name             LSP Name

-------------------------------------------------------------------------------

No FC Mappings

-------------------------------------------------------------------------------
Stp Service Destination Point specifics

-------------------------------------------------------------------------------

Stp Admin State    : Down                     Stp Oper State    : Down

Core Connectivity  : Down

Port Role          : N/A                      Port State        : Forwarding

Port Number        : 0                        Port Priority     : 128
```

```
Port Path Cost     : 10                       Auto Edge         : Enabled

Admin Edge         : Disabled                 Oper Edge         : N/A

Link Type          : Pt-pt                    BPDU Encap        : Dot1d

Root Guard         : Disabled                 Active Protocol   : N/A

Last BPDU from     : N/A

Designated Bridge  : N/A                      Designated Port Id: N/A


Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0

Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0

TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0

TC bit BPDUs rcvd  : 0                        TC bit BPDUs tx   : 0

RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0

-------------------------------------------------------------------------------

Number of SDPs : 1

-------------------------------------------------------------------------------

===============================================================================



A:Dut-A# show service id  1 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
 Sdp Id 1:1  -(10.20.1.2)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1:1                      Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 9186
Far End            : 10.20.1.2                Delivery          : MPLS

Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 2048                     Egress Label      : 2048
Ing mac Fltr       : n/a                      Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                      Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                      Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 05/31/2007 00:45:43      Signaling         : None
Last Mgmt Change   : 05/31/2007 00:45:43
Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
```

```
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0

MAC Learning       : Enabled                 Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                BPDU Translation  : Disabled
MAC Pinning        : Disabled

KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                       I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0
MCAC Policy Name   :
MCAC Max Unconst BW: no limit                MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited

Associated LSP LIST :
Lsp Name           : A_B_1
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_2
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_3
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_4
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_5
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_6
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_7
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_8
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_9
```

```
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_10
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : enabled
Default LSP        : A_B_10                   Multicast LSP     : A_B_9
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
af                  A_B_3
be                  A_B_1
ef                  A_B_6
h1                  A_B_7
h2                  A_B_5
l1                  A_B_4
l2                  A_B_2
nc                  A_B_8
===============================================================================
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
A:Dut-A#
show service id x all
-------------------------------------------------------------------------------
SAP 1/1/4:500
-------------------------------------------------------------------------------
Service Id         : 500
SAP                : 1/1/4:500                Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Multi Svc Site     : None
```

```
Last Status Change : 09/19/2013 11:43:04
Last Mgmt Change   : 09/19/2013 11:43:05
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1518                      Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                       Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                       Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                       Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                      qinq-pbit-marking : both
                                               Egr Agg Rate Limit: max
Endpoint           : N/A
Q Frame-Based Acct : Disabled
Vlan-translation   : None

Acct. Pol          : None                      Collect Stats     : Disabled

Application Profile: None
Transit Policy     : None

Oper Group         : (none)                    Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Ignore Oper Down   : Disabled
Lag Link Map Prof  : (none)
Cflowd             : Disabled

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : n/a                       AIS               : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : 0 1 2 3 4 5 6 7

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                         Egress qos-policy : 1
.
.
.
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 1:2  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 1:2                       Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp    : (Not Specified)
VC Type            : Ether                     VC Tag            : n/a
Admin Path MTU     : 0                         Oper Path MTU     : 0
Delivery           : GRE
Far End            : 1.1.1.1
Tunnel Far End     : n/a                       LSP Types         : n/a
Hash Label         : Disabled                  Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
```

```
Admin State        : Up                       Oper State        : Down
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 0                        Egress Label      : 0
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 09/11/2013 20:02:40      Signaling         : TLDP
Last Mgmt Change   : 09/15/2013 13:56:56      Force Vlan-Vc     : Disabled
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None
Transit Policy     : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
OAM MAC Addr       : 0                        DHCP MAC Addr     : 0
Host MAC Addr      : 0                        Intf MAC Addr     : 0
SPB MAC Addr       : 0                        Cond MAC Addr     : 0

MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                    Block On Mesh Fail: False
Oper Group         : (none)                   Monitor Oper Grp  : (none)
Rest Prot Src Mac  : Disabled
Auto Learn Mac Prot: Disabled                 RestProtSrcMacAct : Disable

Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)

-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
V-MEP Filtering    : Disabled

KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
Max Drop Count     : 3                        Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
```

```
Squelch Levels     : 0 1 2 3 4 5 6 7
```

Table 61 describes show service-id SDP output fields.

Table 61 **Show Service-ID SDP Fields**

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is spoke or mesh.
Split Horizon Group	Indicates the name of the split horizon group that the SDP belongs to.
VC Type	Displays the VC type: ether, vlan, or vpls.
VC Tag	Displays the explicit dot1q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current status of the SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.

Table 61 **Show Service-ID SDP Fields (Continued)**

Label	Description (Continued)
Oper State	The operational state of the Keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the Keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
E. Fwd. Octets	Specifies the number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the above failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signaling method to indicate faults. pwNotForwarding — Pseudowire not forwarding. lacIngressFault Local — Attachment circuit RX fault. lacEgressFault Local — Attachment circuit TX fault. psnIngressFault Local — PSN-facing PW RX fault. psnEgressFault Local — PSN-facing PW TX fault. pwFwdingStandby — Pseudowire in standby mode.

site

Syntax **site** [**detail**]
site *name*

Context show>service>id

Description This command displays sites configures for the service.

Parameters *name* — Specifies the site name up to 32 characters in length

split-horizon-group

Syntax **split-horizon-group** [*group-name*]

Context show>service>id

Description This command displays service split horizon groups.

```
*A:ALA-1# show service id 700 split-horizon-group
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
R DSL-group1                    Split horizon group for DSL
DSL-group1                      Split horizon group for DSL
-------------------------------------------------------------------------------
R = Residential Split Horizon Group
No. of Split Horizon Groups: 1
===============================================================================
*A:ALA-1#

*A:ALA-1# show service id 700 split-horizon-group DSL-group1
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
R DSL-group1                    Split horizon group for DSL
DSL-group1                      Split horizon group for DSL
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
SAP                             1/1/3:1
SDP                             108:1
SDP                             109:1
-------------------------------------------------------------------------------
R = Residential Split Horizon Group
SAPs Associated : 1             SDPs Associated : 2
===============================================================================
*A:ALA-1#
```

stp

Syntax **stp** [**detail**]
stp mst-instance *mst-inst-number*

Context show>service>id

Description This command displays information for the spanning tree protocol instance for the service.

Parameters **detail** — Displays detailed information.

mst-inst-number — Displays information about the specified MST.

Values 1 to 4094

Output

Sample Output

```
*A:ALA-12# show service id 11 stp
===============================================================================
Stp info, Service 11
===============================================================================
Bridge Id          : 80:00.22:68:ff:00:00:00  Top. Change Count : 1
Root Bridge        : 00:00.22:69:ff:00:00:00  Stp Oper State    : Syncing Vcp
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Mstp                     Last Top. Change  : 0d 19:12:58
Vcp Active Prot.   : N/A
Root Port          : 2048                     External RPC      : 10
===============================================================================
MSTP specific info for CIST
===============================================================================
Regional Root      : This Bridge              Root Port         : 2048
Internal RPC       : 0                        Remaining Hopcount: 20

===============================================================================
Stp port info for CIST
===============================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Oper-  Link-  Active
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
1/1/1:0            Up        Root       Forward     2048   False  Pt-pt  Mstp
1/1/3:0            Up        N/A        Forward     2049   N/A    Pt-pt  N/A
1/1/4:*            Up        Designated Forward     2050   False  Pt-pt  Mstp
===============================================================================
MSTP specific info for MSTI 111
===============================================================================
Regional Root      : 80:6f.1c:65:ff:00:00:00  Root Port         : 2050
Internal RPC       : 10                       Remaining Hopcount: 19

=================================================================
MSTP port info for MSTI 111
=================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Same
                   State     Role       State       Num    Region
-----------------------------------------------------------------
1/1/1:0            Up        Master     Forward     2048   False
```

```
1/1/3:0            Up        N/A        Forward     2049   N/A
1/1/4:*            Up        Root       Forward     2050   True
===============================================================
*A:ALA-12#


*A:ALA-12# show service id stp detail
===============================================================================
Spanning Tree Information
===============================================================================
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Mstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.22:68:ff:00:00:00  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 19:14:34              Bridge Max Age    : 20
Top. Change Count  : 1                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    : abc

Root Bridge        : 00:00.22:69:ff:00:00:00
Primary Bridge     : N/A

Root Path Cost     : 10                       Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 0                        Root Port         : 2048

MSTP info for CIST :
Regional Root      : This Bridge              Root Port         : 2048
Internal RPC       : 0                        Remaining Hopcount: 20
MSTP info for MSTI 111 :
Regional Root      : 80:6f.1c:65:ff:00:00:00  Root Port         : 2050
Internal RPC       : 10                       Remaining Hopcount: 19
-------------------------------------------------------------------------------
Spanning Tree Virtual Core Port (VCP) Specifics
-------------------------------------------------------------------------------
Mesh Sdp Id        Sdp           Sdp Bind     Mesh Sdp     HoldDown   Awaiting
                   Oper-state    Oper-state   Port-state   Timer      Agreement
-------------------------------------------------------------------------------
3:11               Down          Down         Discard      Inactive   N/A
4:11               Down          Down         Discard      Inactive   N/A
-------------------------------------------------------------------------------
Spanning Tree Sap/Spoke SDP Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/1/1:0                  Stp Admin State   : Up
Port Role          : Root                     Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : 00:00.22:69:ff:00:00:00  Inside Mst Region : False
CIST Desig Bridge  : 00:00.22:69:ff:00:00:00  Designated Port   : 34816
MSTI 111 Port Prio : 128                      Port Path Cost    : 10
MSTI 111 Desig Brid: This Bridge              Designated Port   : 34816
```

```
Forward transitions: 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 34638                    MST BPDUs tx      : 3

SAP Identifier     : 1/1/3:0                  Stp Admin State   : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : 0
MSTI 111 Port Prio : 128                      Port Path Cost    : 10
MSTI 111 Desig Brid: N/A                      Designated Port   : 0
Forward transitions: 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

SAP Identifier     : 1/1/4:*                  Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2050                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : 50:00.1c:65:ff:00:00:00  Inside Mst Region : True
CIST Desig Bridge  : This Bridge              Designated Port   : 34818
MSTI 111 Port Prio : 128                      Port Path Cost    : 10
MSTI 111 Desig Brid: 80:6f.1c:65:ff:00:00:00  Designated Port   : 34819
Forward transitions: 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 34636                    MST BPDUs tx      : 34640
===============================================================================
*A:ALA-12#
*A:SetupCLI# show service id 2001 stp
===============================================================================
Stp info, Service 2001
===============================================================================
Bridge Id          : 80:00.70:ec:ff:00:00:00  Top. Change Count : 0
Root Bridge        : N/A                      Stp Oper State    : Down
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Rstp                     Last Top. Change  : 0d 00:00:00
Vcp Active Prot.   : N/A
Root Port          : N/A                      External RPC      : 0
===============================================================================
Stp port info
===============================================================================
Sap/Sdp/PIP Id     Oper-     Port-      Port-       Port-  Oper-  Link-  Active
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
Backbone VPLS      Down      N/A        Discard     2048   N/A    N/A    N/A
1/1/12:2001.2001   Down      N/A        Disabled    2049   N/A    Pt-pt  N/A
```

```
===============================================================================
*A:SetupCLI#


*A:SetupCLI# show service id 2001 stp detail
===============================================================================
Spanning Tree Information
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Down                     Core Connectivity : Down
Stp Admin State    : Down                     Stp Oper State    : Down
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.70:ec:ff:00:00:00  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 20
Top. Change Count  : 0                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 32768                    Root Port         : N/A
-------------------------------------------------------------------------------
Spanning Tree Sap/Spoke SDP Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/1/12:2001.2001         Stp Admin State   : Up
Port Role          : N/A                      Port State        : Unknown
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A
Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
Spanning Tree PIP (Provider Internal Port) Specifics
-------------------------------------------------------------------------------
Oper Status        : Down                     mVPLS Prune State : N/A
Port Num           : 2048                     Oper Protocol     : N/A
Port Role          : N/A                      Port State        : Discarding
CIST Desig Bridge  : N/A                      Designated Port   : N/A
b-Vpls STP state   : Disabled
Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
===============================================================================
```

```
*A:SetupCLI#
```

Table 62 describes show service-id STP output fields.

Table 62 **Show Service-ID STP Output Fields**

Label	Description
RSTP Admin State	Indicates the administrative state of the Rapid Spanning Tree Protocol instance associated with this service.
Core Connectivity	Indicates the connectivity status to the core.
RSTP Oper State	Indicates the operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Bridge-id	Specifies the MAC address used to identify this bridge in the network.
Hold Time	Specifies the interval length during which no more than two Configuration BPDUs shall be transmitted by this bridge.
Bridge fwd delay	Specifies how fast a bridge changes its state when moving toward the forwarding state.
Bridge Hello time	Specifies the amount of time between the transmission of Configuration BPDUs.
Bridge max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. This is the actual value that this bridge is currently using.
Bridge priority	Defines the priority of the Spanning Tree Protocol instance associated with this service.
Topology change	Specifies whether a topology change is currently in progress.
Last Top. change	Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service.
Top. change count	Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management entity was last reset or initialized.
Root bridge-id	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.

Table 62 **Show Service-ID STP Output Fields (Continued)**

Label	Description (Continued)
Root path cost	Specifies the cost of the path to the root bridge as seen from this bridge.
Root forward delay	Specifies how fast the root changes its state when moving toward the forwarding state.
Root hello time	Specifies the amount of time between the transmission of configuration BPDUs.
Root max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded.
Root priority	This object specifies the priority of the bridge that is currently selected as root-bridge for the network.
Root port	Specifies the port number of the port which offers the lowest cost path from this bridge to the root bridge.
SAP Identifier	The ID of the access port where this SAP is defined.
RSTP State	The operational state of RSTP.
STP Port State	Specifies the port identifier of the port on the designated bridge for this port's segment.
BPDU encap	Specifies the type of encapsulation used on BPDUs sent out and received on this SAP.
Port Number	Specifies the value of the port number field which is contained in the least significant 12 bits of the 16-bit port ID associated with this SAP.
Priority	Specifies the value of the port priority field which is contained in the most significant 4 bits of the 16-bit port ID associated with this SAP.
Cost	Specifies the contribution of this port to the path cost of paths toward the spanning tree root which include this port.
Fast Start	Specifies whether Fast Start is enabled on this SAP.
Designated Port	Specifies the port identifier of the port on the designated bridge for this port's segment.
Designated Bridge	Specifies the bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment.
Service Access Points	

Table 62 **Show Service-ID STP Output Fields (Continued)**

Label	Description (Continued)
Managed by Service	Specifies the service ID of the management VPLS managing this SAP or spoke-SDP.
Managed by SAP or spoke	Specifies the SAP ID or SDP ID inside the management VPLS managing this SAP or spoke-SDP.
Prune state	Specifies the STP state inherited from the management VPLS.

subscriber-hosts

Syntax **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]

Context show>service>id

Description This command displays subscriber host information.

Parameters *sap sap-id* — Displays the specified subscriber host SAP information

ip-address/mask — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets).
mask: 1 to 32

mac ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sub-profile-name — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile-name — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

detail — Displays detailed information

Output The following output is an example of service subscriber host information.

Sample Output

```
A:ALA#-SR12# show service id 20 subscriber-hosts
===============================================================================
```

```
Subscriber Host table
===============================================================================
Sap Id                 IP Address      MAC Address       Origin(*) Subscriber
-------------------------------------------------------------------------------
1/2/6:0                101.1.1.10      00:bb:bb:00:00:00 S/-/-
   Eval-20-static
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:ALA#

A:ALA# show service id 10 subscriber-hosts
===============================================================================
Subscriber Host table
===============================================================================
Sap Id                 IP Address      MAC Address       Origin(*) Subscriber
-------------------------------------------------------------------------------
1/2/5:0                100.1.1.10      00:aa:aa:00:00:01 -/D/-
   SUB-10-00aaaa000001
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:ALA-SR12#
```

fdb-usage

Syntax **fdb-usage** [**card** *slot-id*]

Context show>service>system

Description This command displays the FDB usage, excluding the pending updates (which can be seen using the **tools dump service id** *id* **fdb {card-status | mac-status}** command) for the system and all line cards.

Parameters *slot-id* — Displays the information for the line card in the specified slot IDs, expressed as an integer.

Values 1 to 20

Output The following output is an example of FDB usage information.

Sample Output

```
*A:PE1# show service system fdb-usage
===============================================================================
FDB Usage
===============================================================================
System
-------------------------------------------------------------------------------
Limit:     511999
Allocated: 8
Free:      511991
Global:    2
```

```
-------------------------------------------------------------------------------
Line Cards
-------------------------------------------------------------------------------
Card        Selective         Allocated         Limit             Free
-------------------------------------------------------------------------------
1           0                 2                 511999            511997
2           4                 6                 511999            511993
5           2                 4                 511999            511995
-------------------------------------------------------------------------------
===============================================================================
*A:PE1#
*A:PE1# show service system fdb-usage card 1
===============================================================================
FDB Usage
===============================================================================
Card        Selective         Allocated         Limit             Free
-------------------------------------------------------------------------------
1           0                 2                 511999            511997
===============================================================================
===============================================================================
*A:PE1#
```

statistics

Syntax **statistics** [**policy** *name*] [**sap** *sap-id*]

Context show>service>id>authentication

Description This command displays session authentication statistics for this service.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

Output

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
===============================================================
Authentication statistics
===============================================================
Interface / SAP                Authentication  Authentication
                               Successful      Failed
---------------------------------------------------------------
vpls-11-90.1.0.254              1582            3
---------------------------------------------------------------
Number of entries: 1
===============================================================
*A:ALA-1#
```

3.8.2.2 IGMP Snooping Show Commands

igmp-snooping

Syntax **igmp-snooping**

Context show>service>id

Description This command enables the context to display IGMP snooping information.

all

Syntax **all**

Context show>service>id>igmp-snooping

Description This command displays detailed information for all aspects of IGMP snooping on the VPLS service.

Output

Sample Output

```
A:ALA-48>show>service>id>igmp-snooping>snooping# all
===============================================================================
IGMP Snooping info for service 750
===============================================================================
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                Oper      MRtr  Send      Max Num   Num
Id                     State     Port  Queries   Groups    Groups
-------------------------------------------------------------------------------
sap:1/1/7:0            Down      No    Disabled  No Limit  0
sdp:1:22               Down      No    Disabled  No Limit  0
sdp:8:750              Down      No    Disabled  No Limit  0
-------------------------------------------------------------------------------
IGMP Snooping Querier info
-------------------------------------------------------------------------------
No querier found for this service.
-------------------------------------------------------------------------------
IGMP Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
Number of mrouters: 0
-------------------------------------------------------------------------------
IGMP Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
```

```
Group Address    Mode     Type     Up Time         Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/7:0 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time         Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SDP 1:22 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time         Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SDP 8:750 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time         Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups
-------------------------------------------------------------------------------
IGMP Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted      Forwarded
-------------------------------------------------------------------------------
General Queries         0             0                0
Group Queries           0             0                0
Group-Source Queries    0             0                0
V1 Reports              0             0                0
V2 Reports              0             0                0
V3 Reports              0             0                0
V2 Leaves               0             0                0
Unknown Type            0             N/A              0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
===============================================================================
A:ALA-48>show>service>id>snooping#
```

Table 63 describes the show all service-id command output fields.

Table 63 **IGMP Snooping Fields**

Label	Description
Admin State	The administrative state of the IGMP instance.
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached.
Sap Id	Displays the SAP IDs of the service ID.
Oper State	Displays the operational state of the SAP IDs of the service ID.
Mrtr Port	Specifies if the port is a multicast router port.
Send Queries	Specifies whether the send-queries command is enabled or disabled.
Max Num Groups	Specifies the maximum number of multicast groups that can be joined on this SAP.
MVR From VPLS	Specifies MVR from VPLS.
Num Groups	Specifies the actual number of multicast groups that can be joined on this SAP.

mfib

Syntax **mfib** [**ipv4** | **ipv6** | **mac**]
mfib brief
mfib group *group-address* [**statistics**]
mfib statistics [**ipv4** | **ipv6** | **mac**]

Context show>service>id

Description This command displays the multicast FIB on the VPLS service.

Parameters **brief** — Displays a brief output

statistics — Displays statistics on the multicast FIB

ipv4 — Displays IPv4 address information

ipv6 — Displays IPv6 address information

mac — Displays MAC address information

group-address — Displays the multicast FIB for a specific multicast group address

Output

Sample Output

```
*A:PE# show service id 1 mfib
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                       Svc Id   Fwd
                                                                            Blk
-------------------------------------------------------------------------------
10.0.0.2        233.252.0.1           sap:1/1/1                     Local    Fwd
                                      sap:1/1/2                     Local    Fwd
2001:db8:1000:* ff0e:db8:1000::1      sap:1/1/1                     Local    Fwd
                                      sap:1/1/2                     Local    Fwd
2001:db8:1001:* ff0e:db8:1001::1      sap:1/1/1                     Local    Fwd
                                      sap:1/1/2                     Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 3
===============================================================================
*A:PE# show service id 1 mfib ipv4
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                       Svc Id   Fwd
                                                                            Blk
-------------------------------------------------------------------------------
10.0.0.2        233.252.0.1           sap:1/1/1                     Local    Fwd
                                      sap:1/1/2                     Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
*A:PE# show service id 1 mfib ipv6
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address
                Group Address
                                      Port Id                       Svc Id   Fwd
                                                                            Blk
-------------------------------------------------------------------------------
2001:db8:1000::1
                ff0e:db8:1000::1
                                      sap:1/1/1                     Local    Fwd
                                      sap:1/1/2                     Local    Fwd
2001:db8:1001::1
                ff0e:db8:1001::1
                                      sap:1/1/1                     Local    Fwd
                                      sap:1/1/2                     Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
*A:PE# show service id 1 mfib statistics
===============================================================================
Multicast FIB Statistics, Service 1
===============================================================================
Source Address  Group Address       Matched Pkts          Matched Octets
-------------------------------------------------------------------------------
10.0.0.2        233.252.0.1         0                     0
2001:db8:1000:* ff0e:db8:1000::1    0                     0
2001:db8:1001:* ff0e:db8:1001::1    0                     0
-------------------------------------------------------------------------------
Number of entries: 3
```

```
===============================================================================
*A:PE#
```

To show which ISIDs are local, the MFIB command will display ISIDs that are local and advertised. Static ISIDs are included in this display. However, ISID policy can override the ISIDs that are designated to use the default multicast tree and these do not show up in the MFIB. This is displayed on a B-VPLS control service.

```
*A:cses-B0102>show>service>id# mfib

===============================================================================
Multicast FIB, Service 510
===============================================================================
Source Address  Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               01:1E:83:00:01:F4     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:F5     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:F6     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:F7     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:F8     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:F9     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:FA     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:FB     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:FC     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:FD     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:FE     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:01:FF     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:02:00     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:02:01     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:02:02     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:02:03     b-sap:1/1/22:510         Local    Fwd
*               01:1E:83:00:02:04     b-sap:1/1/22:510         Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 21
===============================================================================
```

To show the ISID policy under a B-VPLS, the ISID policy is used.

```
*A:cses-B07>show>service>id# isid-policy

=======================================================
Isid Policy Range
=======================================================
Entry     Range               AdvLocal  UseDefMCTree
-------------------------------------------------------
2         1500-1600           Disabled  Enabled
=======================================================
```

The following example shows the MFIB for an EVPN-MPLS service.

```
*A:PE# show service id 1 mfib
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Sap/Sdp Id                   Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
```

```
*               235.0.0.1             sap:1/1/9:1                  Local    Fwd
                                      eMpls:1.1.1.2:262141         Local    Fwd
                                      eMpls:1.1.1.3:262141         Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
*A:PE#
```

Table 64 describes the command output fields.

Table 64 **Multicast FIB Fields**

Label	Description
Source Address	IPv4 unicast source address.
Group Address	IPv4 multicast group address.
SAP ID	Indicates the SAP/SDP to which the corresponding multicast stream will be forwarded/blocked.
Forwarding/Blocking	Indicates whether the corresponding multicast stream will be blocked/forwarded.
Number of Entries	Specifies the number of entries in the MFIB.
Forwarded Packets	Indicates the number of multicast packets forwarded for the corresponding source/group.
Forwarded Octets	Indicates the number of octets forwarded for the corresponding source/group.
Svc ID	Indicates the service to which the corresponding multicast stream will forwarded/blocked. Local means that the multicast stream will be forwarded/blocked to a SAP or SDP local to the service.

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>igmp-snooping

Description This command displays all multicast routers.

Parameters **detail** — Displays detailed information.

Output

Sample Output

```
*A:ala-427# show service id 1 igmp-snooping mrouters
===============================================================================
```

```
IGMP Snooping Multicast Routers for service 1
===============================================================================
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
10.10.1.1        1/1/5:1                0d 00:00:26    14s       3
10.20.1.6        1/1/2:1                0d 00:10:16    2s        3
-------------------------------------------------------------------------------
Number of mrouters: 2
===============================================================================
*A:ala-427#


*A:ala-427# show service id 1 igmp-snooping mrouters detail
===============================================================================
IGMP Snooping Multicast Routers for service 1
===============================================================================
MRouter 10.10.1.1
-------------------------------------------------------------------------------
Sap Id                  : 1/1/5:1
Expires                 : 17s
Up Time                 : 0d 00:00:32
Version                 : 3

General Query Interval  : 10s
Query Response Interval : 1.0s
Robust Count            : 2
-------------------------------------------------------------------------------
MRouter 10.20.1.6
-------------------------------------------------------------------------------
Sap Id                  : 1/1/2:1
Expires                 : 3s
Up Time                 : 0d 00:10:22
Version                 : 3

General Query Interval  : 2s
Query Response Interval : 1.0s
Robust Count            : 2
-------------------------------------------------------------------------------
Number of mrouters: 2
===============================================================================
*A:ala-427#
```

mvr

Syntax **mvr**

Context show>service>id>igmp-snooping

Description This command displays Multicast VPLS Registration (MVR) information.

Output

Sample Output

```
*A:ALA-1>show>service>id>snooping# mvr
```

```
===============================================================================
IGMP Snooping Multicast VPLS Registration info for service 10
===============================================================================
IGMP Snooping Admin State : Up
MVR Admin State           : Up
MVR Policy                : mvr-policy
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num Local
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
100         sap:1/1/10:10        Up        Local      100
100         sap:1/1/10:20        Up        Local      100
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=10)
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num MVR
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
20         sap:1/1/4:100          Up        10         100
30         sap:1/1/31:10.10       Up        10         100
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

Table 65 describes the show all service-id command output fields.

Table 65 **Show All Service-ID Fields**

Label	Description
MVR Admin State	Administrative state.
MVR Policy	Policy name.
Svc ID	The service identifier.
Sap/Sdp Id	Displays the SAP and SDP IDs of the service ID.
Oper State	Displays the operational state of the SAP and SDP IDs of the svcid.
Mrtr Port	Specifies if the port is a multicast router port.
From VPLS	Specifies from which VPLS the multicast streams corresponding to the groups learned via this SAP will be copied. If local, it is from its own VPLS.
Num Groups	Specifies the number of groups learned via this local SAP.

port-db

Syntax **port-db sap** *sap-id* [**detail**]

port-db sap *sap-id* **group** *grp-address*
port-db sdp *sdp-id*:*vc-id* [**detail**]
port-db sdp *sdp-id*:*vc-id* **group** *grp-address*
vxlan vtep *ip-address* **vni** *vni*

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping port database for the VPLS service.

Parameters *grp-ip-address* — Displays the IGMP snooping port database for a specific multicast group address

sap-id — Displays the IGMP snooping port database for a specific SAP

sdp-id — Displays only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information

Default For mesh SDPs only, all VC IDs

Values 1 to 4294967295

grp-address — Displays IGMP snooping statistics matching the specified group address. This parameter only applies to the 7450 ESS or 7750 SR.

ip-address — Displays IGMP snooping statistics matching one particular source within the multicast group. This parameter only applies to the 7450 ESS or 7750 SR.

vxlan vtep *ip-address* **vni** *<1..16777215>* — Displays the IGMP snooping entries associated with a specific VXLAN binding, given by the VXLAN Termination Endpoint (VTEP) and VXLAN Network Identifier (VNI). This parameter only applies to the 7450 ESS or 7750 SR.

vni — The VXLAN Network Identifier (VNI) for which to display information. This parameter only applies to the 7450 ESS or 7750 SR.

Values 1 to 16777215

Output

Sample Output

```
*A:ALA-1>show>service>id>snooping# port-db sap 1/1/2
===============================================================================
IGMP Snooping SAP 1/1/2 Port-DB for service 10
===============================================================================
Group Address    Mode     Type     Up Time        Expires   Num Sources
-------------------------------------------------------------------------------
225.0.0.1        include  dynamic  0d 00:04:44    0s        2
Group Address    Mode     Type     From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
225.0.0.1        include  dynamic  Local      0d 00:04:44    0s       2
```

```
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:ALA-1>show>service>id>snooping#


*A:ALA-1>show>service>id>snooping# port-db sap 1/1/2 detail
===============================================================================
IGMP Snooping SAP 1/1/2 Port-DB for service 10
===============================================================================
IGMP Group 225.0.0.1
-------------------------------------------------------------------------------
Mode             : include              Type             : dynamic
Up Time          : 0d 00:04:57          Expires          : 0s
Compat Mode      : IGMP Version 3
V1 Host Expires  : 0s                   V2 Host Expires  : 0s
-------------------------------------------------------
Source Address  Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
1.1.1.1         0d 00:04:57  20s      dynamic  Fwd
1.1.1.2         0d 00:04:57  20s      dynamic  Fwd
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

Table 66 describes the show output fields.

Table 66 **IGMP Snooping Port Database Fields**

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	Specifies the type of membership reports received on the interface for the group. In the include mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In the exclude mode, reception of packets sent to the specified multicast address is requested from all IP source addresses except those listed in the **source-list** parameter.
Type	Indicates how this group entry was learned. If this group entry was learned by IGMP, the value is set to dynamic. For statically configured groups, the value is set to static.

Table 66 **IGMP Snooping Port Database Fields (Continued)**

Label	Description
Compatibility mode	Specifies the IGMP mode. This is used for routers to be compatible with older-version routers. IGMPv3 hosts must operate in Version 1 and Version 2 compatibility modes. IGMPv3 hosts must keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the host compatibility mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of general queries heard on that interface as well as the older-version querier present timers for the interface.
V1 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 leave messages for this group that it receives on this interface.
V2 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 leave messages for this group that it receives on this interface.
Source address	The source address for which this entry contains information.
Up Time	The time since the source group entry was created.
Expires	The amount of time remaining before this entry will be aged out.
Number of sources	Indicates the number of IGMP group and source specific queries received on this SAP.
Forwarding/Blocking	Indicates whether this entry is on the forward list or block list.
Number of groups	Indicates the number of groups configured for this SAP.

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-address*

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping proxy reporting database for the VPLS service.

Parameters *grp-ip-address* — Displays the IGMP snooping proxy reporting database for a specific multicast group address.

Output

Sample Output

```
*A:ALA-1>show>service>id>snooping# proxy-db
===============================================================================
IGMP Snooping Proxy-reporting DB for service 10
===============================================================================
Group Address    Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
225.0.0.1        include  0d 00:05:40    2
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:ALA-1>show>service>id>snooping#


*A:ALA-1>show>service>id>snooping# proxy-db detail
===============================================================================
IGMP Snooping Proxy-reporting DB for service 10
-------------------------------------------------------------------------------
IGMP Group 225.0.0.1
-------------------------------------------------------------------------------
Up Time : 0d 00:05:54                  Mode : include
------------------------------
Source Address  Up Time
------------------------------
1.1.1.1         0d 00:05:54
1.1.1.2         0d 00:05:54
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

Table 67 describes the show output fields.

Table 67 **IGMP Snooping proxy Fields**

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	Specifies the type of membership report(s) received on the interface for the group. In the include mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report.

Table 67 **IGMP Snooping proxy Fields (Continued)**

Label	Description (Continued)
	In the “exclude” mode, reception of packets sent to the specified multicast address is requested from all IP source addresses except those listed in the source-list parameter.
Up Time	The total operational time in seconds.
Num Sources	Indicates the number of IGMP group and source specific queries received on this interface.
Number of groups	Number of IGMP groups.
Source Address	The source address for which this entry contains information.

querier

Syntax **querier**

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping queriers for the VPLS service.

Output

Sample Output

```
*A:ALA-1>show>service>id>snooping# querier
===============================================================================
IGMP Snooping Querier info for service 10
===============================================================================
Sap Id                  : 1/1/1
IP Address              : 10.10.10.1
Expires                 : 6s
Up Time                 : 0d 00:56:50
Version                 : 3

General Query Interval  : 5s
Query Response Interval : 2.0s
Robust Count            : 2
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

Table 68 describes the show output fields.

Table 68 **IGMP Snooping Queriers Fields**

Label	Description
SAP Id	Specifies the SAP ID of the service.
IP address	Specifies the IP address of the querier.
Expires	The time left, in seconds, that the query will expire.
Up time	The length of time the query has been enabled.
Version	The configured version of IGMP.
General Query Interval	The frequency at which host-query packets are transmitted.
Query Response Interval	The time to wait to receive a response to the host-query message from the host.
Robust Count	Specifies the value used to calculate several IGMP message intervals.

static

Syntax **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description This command displays information on static IGMP snooping source groups for the VPLS service.

Parameters **sap** *sap-id* — Displays static IGMP snooping source groups for a specific SAP.

sdp *sdp-id* — Displays the IGMP snooping source groups for a specific spoke or mesh SDP.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Default For mesh SDPs only, all VC IDs.

Values 1 to 4294967295

Output

Sample Output

```
*A:ALA-1>show>service>id>snooping# static
===============================================================================
IGMP Snooping Static Source Groups for SAP 1/1/2
-------------------------------------------------------------------------------
Source           Group
```

```
-------------------------------------------------------------------------------
*                 225.0.0.2
*                 225.0.0.3
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 2
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SDP 10:10
-------------------------------------------------------------------------------
Source            Group
-------------------------------------------------------------------------------
1.1.1.1           225.0.0.10
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 1
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

Table 69 describes the show output fields.

Table 69 **IGMP Snooping Source Groups Fields**

Label	Description
Source	Displays the IP source address used in IGMP queries.
Group	Displays the static IGMP snooping source groups for a specified SAP.

statistics

Syntax **statistics** [**evpn-mpls** | **sap** *sap-id* | **sdp** *sdp-id:vc-id* | **vxlan vtep** *ip-address* **vni** *vni*]

Context show>service>id>igmp-snooping

Description This command displays IGMP snooping statistics for the VPLS service.

Parameters **evpn-mpls** — Displays IGMP snooping statistics for EVPN-MPLS destinations

sap-id — Displays IGMP snooping statistics for a specific SAP

sdp-id — Displays the IGMP snooping statistics for a specific spoke or mesh SDP

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information

Default For mesh SDPs only, all VC IDs.

Values 1 to 4294967295

vxlan vtep *ip-address* **vni** *<1..16777215>* — Displays the IGMP snooping entries associated with a specific VXLAN binding, given by the VXLAN Termination Endpoint (VTEP) and VXLAN Network Identifier (VNI). This parameter only applies to the 7450 ESS or 7750 SR.

vni — The VXLAN Network Identifier (VNI) for which to display information. This parameter only applies to the 7450 ESS or 7750 SR.

Values 1 to 16777215

Output

Sample Output

```
*A:ALA-1>show>service>id>snooping# statistics
===============================================================================
IGMP Snooping Statistics for service 1
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         4             0             4
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

egress-replication

Syntax **egress-replication**

Context show

Description This command enables the context to display egress flooding information for a VPLS service context on a specified MDA. A VPLS service context supports both Layer 2 and Layer 3 flooding modes. The Layer 2 flooding mode is used for broadcast, Layer 2 multicast and unknown destination MAC addressed packets. All available interfaces (SAP, spoke-SDP and mesh-SDP) that reside on an egress forwarding complex are included in the egress list except for SAPs that are defined in a residential split horizon group (Layer 2 flooding is not permitted on residential SAPs). The Layer 3 flooding mode is used for VPLS interfaces participating in IGMP snooping and is represented by an IP multicast [s,g] record.

vpls

Syntax **vpls** *vpls-service-id* **mda** *card/slot*
vpls *vpls-service-id* **mda** *card/slot* [**igmp-record** *group ip-address* {**source** *ip-address* | **starg**}]

Context show>egress-replication

Description The **vpls** *vpls-service-id* **mda** *slot/mda* command displays the flooding list used by the Layer 2 flooding mode for the VPLS service on the specified MDA. The Layer 2 flooding list is limited to SAPs, spoke-SDP and mesh-SDP bindings that exist on the egress forwarding complex serviced by the specified MDA. The only VPLS interfaces that will not be included in the list are residential SAPs because Layer 2 replication is not permitted to a residential SAP. A packet processed by the egress Layer 2 flooding list may not be replicated to each destination. A packet will not be replicated to an interface on the Layer 2 flooding list because of the following:

The ingress interface is the same as egress interface (source squelching rule).

- The ingress interface split horizon group is the same as the egress interface (residential bridging rule).
- The egress interface is down or blocking.
- The packet matches a discard event while processing that destination interface.
- An egress MTU violation occurs for the destination interface.

Destination SAPs in the list may be displayed in a chain context representing common replication behavior. All SAPs in a single chain are processed a single time through the egress forwarding plane. If a discard decision is made for the first SAP in the chain, no replication processing is done for any of the chain members. If the forwarding plane decides to replicate the first SAP in the chain, it will replicate to all SAPs in the chain.

The **vpls** *vpls-service-id* **mda card/slot igmp-record** *grp-address* {**source** *source-ip-address* | **starg**} command displays the IGMP record based flooding list for the *vpls-service-id* on the specified MDA. Unlike the Layer 2 flooding list for the VPLS context, an IGMP record list may contain interfaces from other VPLS contexts due to MVR (Multicast VPLS Registration) events on the individual VPLS interfaces. VPLS interfaces in other VPLS

contexts become associated with the specified vpls-service-id based on the MVR from-vpls definition. Another difference between the VPLS Layer 2 flooding list and IGMP lists is that many IGMP lists may exist (each associated with a different [s,g] record) and the lists may contain residential SAPs. The SAP chaining and replication behavior is similar to the VPLS Layer 2 flooding list.

IP multicast packets ingressing the vpls-service-id must match either a [*,g] or [s,g] record to be associated with the record's egress IP multicast IGMP flooding list. A [*,g] record will match any ingress IP multicast packet destined for the class D destination IP address represented by "g". An [s,g] record will match any ingress IP multicast packet with a source IP address matching "s" and a destination IP address matching "g". In the case that a packet could match both a [*,g] and [s,g] record, the [s,g] record takes precedence. Each [*,g] and [s,g] record has its own IGMP flooding list. The list will only appear on an egress forwarding plane (MDA) when a member of the list (VPLS interface) exists on the forwarding plane.

Parameters *service-id* — Displays information about the specified service ID or service name

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters in length

slot/mda — Specifies a chassis and MDA slot

grp-ip-address — Specifies a multicast group address

src-ip-address — Specifies a source IP address

starg — Specifies a (*, G) record

mRouter — Specifies the (*,*) record.

ipv6 — Displays IPv6 information

grp-ipv6-address — ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D
multicast group IPv6 address

src-ipv6-address — ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D

3.8.2.3 IGMP Commands

group

Syntax **group** [*grp-ip-address*]

Context show>router>igmp

Description This command displays the multicast group and (s, g) addresses. If no *grp-ip-address* parameters are specified then all IGMP group, (*, g) and (s, g) addresses are displayed.

Parameters *grp-ip-address* — Displays specific multicast group addresses

Output

Sample Output

```
A:NYC# show router igmp group
===============================================================================
IGMP Groups
===============================================================================
(*,224.24.24.24)                        Up Time : 0d 05:21:38
    Fwd List  : nyc-vlc

(*,239.255.255.250)                     Up Time : 0d 05:21:38
    Fwd List  : nyc-vlc
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 2
===============================================================================
A:NYC#


A:NYC# show router igmp group 224.24.24.24
===============================================================================
IGMP Groups
===============================================================================
(*,224.24.24.24)                        Up Time : 0d 05:23:23
    Fwd List  : nyc-vlc
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 1
===============================================================================
A:NYC#
```

Table 70 describes the output fields for IGMP group information.

Table 70 **IGMP Group Output Fields**

Label	Description
IGMP Groups	Displays the IP multicast sources corresponding to the IP multicast groups which are statically configured.
Fwd List	Displays the list of interfaces in the forward list.
Blk List	Displays the list of interfaces in the block list.

ssm-translate

Syntax **ssm-translate**

Context show>router>igmp

Description This command displays IGMP SSM translate configuration information.

Output

Sample Output

```
A:ALA-48>config>router>igmp# show router igmp ssm-translate
=======================================================
IGMP SSM Translate Entries
=======================================================
Group Range                             Source
-------------------------------------------------------
<224.0.1.0 - 224.0.1.255>               1.1.1.1
<225.1.0.0 - 225.240.3.57>              2.2.2.2
<239.255.255.0 - 239.255.255.255>       3.3.3.3
-------------------------------------------------------
SSM Translate Entries : 3
=======================================================
A:ALA-48>config>router>igmp#
```

Table 71 provides IGMP SSM-Translate output field descriptions.

Table 71 **IGMP SSM-Translate Output Fields**

Label	Description
Group Range	Displays the address ranges of the multicast groups for which this router can be an RP.
Source	Displays the unicast address that sends data on an interface.
SSM Translate Entries	Displays the total number of SSM translate entries.

interface

Syntax **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-address*] [**detail**]

Context show>router>igmp

Description This command displays IGMP interface information.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name

ip-address — Only displays the information associated with the specified IP address

grp-address — Only displays IP multicast group address for which this entry contains information

detail — Displays detailed IP interface information along with the source group information learned on that interface

Output

Sample Output

```
A:BA# show router igmp interface
===============================================================================
IGMP Interfaces
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BA#

A:BA# show router 100 igmp interface IGMP_to_CE
===============================================================================
IGMP Interface IGMP_to_CE
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BA#

A:BA# show router 100 igmp interface 11.1.1.1
===============================================================================
IGMP Interface 11.1.1.1
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BA#


A:BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1
===============================================================================
IGMP Interface IGMP_to_CE
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
```

```
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
IGMP Group
-------------------------------------------------------------------------------
Group Address : 227.1.1.1               Up Time       : 0d 00:03:52
Interface     : IGMP_to_CE              Expires       : never
Last Reporter : 0.0.0.0                 Mode          : exclude
V1 Host Timer : Not running             Type          : static
V2 Host Timer : Not running             Compat Mode   : IGMP Version 3
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================


A:BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1 detail
===============================================================================
IGMP Interface IGMP_to_CE
===============================================================================
Interface          : IGMP_to_CE
Admin Status       : Up                 Oper Status        : Up
Querier            : 11.1.1.1           Querier Up Time    : 0d 00:04:01
Querier Expiry Time: N/A                Time for next query: 0d 00:13:42
Admin/Oper version : 1/1                Num Groups         : 3
Policy             : igmppol            Subnet Check       : Disabled
Max Groups Allowed : 16000              Max Groups Till Now: 3
MCAC Policy Name   :                    MCAC Const Adm St  : Enable
MCAC Max Unconst BW: no limit           MCAC Max Mand BW   : no limit
MCAC In use Mand BW: 0                  MCAC Avail Mand BW : unlimited
MCAC In use Opnl BW: 0                  MCAC Avail Opnl BW : unlimited
-------------------------------------------------------------------------------
IGMP Group
-------------------------------------------------------------------------------
Group Address : 227.1.1.1               Up Time       : 0d 00:04:02
Interface     : IGMP_to_CE              Expires       : never
Last Reporter : 0.0.0.0                 Mode          : exclude
V1 Host Timer : Not running             Type          : static
V2 Host Timer : Not running             Compat Mode   : IGMP Version 3
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BA#
```

Table 72 provides IGMP Interface output field descriptions.

Table 72 **IGMP Interface Output Fields**

Label	Description
Interface	Specifies the interfaces that participates in the IGMP protocol.
Adm Admin Status	Displays the administrative state for the IGMP protocol on this interface.
Oper Oper Status	Displays the current operational state of IGMP protocol on the interface.

Table 72 **IGMP Interface Output Fields (Continued)**

Label	Description (Continued)
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached.
Querier Up Time	Displays the time since the querier was last elected as querier.
Querier Expiry Timer	Displays the time remaining before the querier ages out. If the querier is the local interface address, the value will be zero.
Cfg/Opr Version Admin/Oper version	Cfg — The configured version of IGMP running on this interface. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN. Opr — The operational version of IGMP running on this interface. If the cfg value is 3 but all of the routers in the local subnet of this interface use IGMP version v1 or v2, the operational version will be v1 or v2.
Num Groups	The number of multicast groups which have been learned by the router on the interface.
Policy	Specifies the policy that is to be applied on the interface.
Group Address	Specifies the IP multicast group address for which this entry contains information.
Up Time	Specifies the time since this source group entry got created.
Last Reporter	Specifies the IP address of the source of the last membership report received for this IP Multicast group address on this interface. If no membership report has been received, this object has the value 0.0.0.0.
Mode	The mode is based on the type of membership report(s) received on the interface for the group. In the 'include' mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In 'exclude' mode, reception of packets sent to the specified multicast address is requested from all IP source addresses except those listed in the source-list parameter.
V1 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface.

Table 72 **IGMP Interface Output Fields (Continued)**

Label	Description (Continued)
V2 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 Leave messages for this group that it receives on this interface.
Type	Indicates how this group entry was learned. If this group entry was learned by IGMP, it will be set to 'dynamic'. For statically configured groups, the value will be set to 'static'.
Compat Mode	Used in order for routers to be compatible with older version routers. IGMPv3 hosts must operate in version 1 and version 2 compatibility modes. IGMPv3 hosts must keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the Host Compatibility Mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of General Queries heard on that interface as well as the Older Version Querier Present timers for the interface.

static

Syntax **static** [*ip-int-name* | *ip-addr*]

Context show>router>igmp

Description This command displays static IGMP, (*, G) (S, G) information.

Parameters *ip-int-name* — Displays the information associated with the specified IP interface name

ip-addr — Displays the information associated with the specified IP address

Output

Sample Output

```
A:BA# show router 100 igmp static
===============================================================
IGMP Static Group Source
===============================================================
Source           Group            Interface
---------------------------------------------------------------
11.11.11.11      226.136.22.3     IGMP_to_CE
*                227.1.1.1        IGMP_to_CE
```

```
22.22.22.22      239.255.255.255  IGMP_to_CE
-------------------------------------------------------------------
Static (*,G)/(S,G) Entries : 3
===================================================================
A:BA#
```

Table 73 provides static IGMP field descriptions.

Table 73 **Static IGMP Output Fields**

Label	Description
Source	Displays entries which represents a source address from which receivers are interested/not interested in receiving multicast traffic.
Group	Displays the IP multicast group address for which this entry contains information.
Interface	Displays the interface name.

statistics

Syntax **statistics** [*ip-int-name* | *ip-addr*ess]

Context show>router>igmp

Description This command displays IGMP statistics information.

Parameters *ip-int-name* — Displays the information associated with the specified IP interface name

ip-addr — Displays the information associated with the specified IP address.

Output

Sample Output

```
A:BA# show router 100 igmp statistics
===============================================
IGMP Interface Statistics
===============================================
Message Type        Received       Transmitted
-----------------------------------------------
Queries             0              5
Report V1           0              0
Report V2           0              0
Report V3           0              0
Leaves              0              0
-----------------------------------------------
General Interface Statistics
-----------------------------------------------
Bad Length        : 0
Bad Checksum      : 0
```

```
Unknown Type      : 0
Bad Receive If    : 0
Rx Non Local      : 0
Rx Wrong Version  : 0
Policy Drops      : 0
No Router Alert   : 0
Rx Bad Encodings  : 0
Rx Pkt Drops      : 0
-----------------------------------------------
Source Group Statistics
-----------------------------------------------
(S,G)             : 2
(*,G)             : 1
===============================================
A:BA#
```

Table 74 provides statistical IGMP field descriptions.

Table 74 **IGMP Statistics Output Fields**

Label	Description
IGMP Interface Statistics	The section listing the IGMP statistics for a particular interface.
Message Type	Queries — The number of IGMP general queries transmitted or received on this interface. Report — The total number of IGMP V1, V2, or V3 reports transmitted or received on this interface. Leaves — The total number of IGMP leaves transmitted on this interface.
Received	Column that displays the total number of IGMP packets received on this interface.
Transmitted	Column that displays the total number of IGMP packets transmitted from this interface.
General Interface Statistics	The section listing the general IGMP statistics.
Bad Length	Displays the total number of IGMP packets with bad length received on this interface.
Bad Checksum	Displays the total number of IGMP packets with bad checksum received on this interface.
Unknown Type	Displays the total number of IGMP packets with unknown type received on this interface.
Bad Receive If	Displays the total number of IGMP packets incorrectly received on this interface.

Table 74 **IGMP Statistics Output Fields (Continued)**

Label	Description (Continued)
Rx Non Local	Displays the total number of IGMP packets received from a non-local sender.
Rx Wrong Version	Displays the total number of IGMP packets with wrong versions received on this interface.
Policy Drops	Displays the number of times IGMP protocol instance matched the host IP address or group/source addresses in the import policy.
No Router Alert	Displays the total number of IGMPv3 packets received on this interface which did not have the router alert flag set.

status

Syntax **status**

Context show>router>igmp

Description This command displays IGMP status information.

If IGMP is not enabled, the following message appears:

```
A:NYC# show router igmp status
MINOR: CLI IGMP is not configured.
A:NYC#
```

Output

Sample Output

```
A:BA# show router 100 igmp status
===============================================================================
IGMP Status
===============================================================================
Admin State                      : Up
Oper State                       : Up
Query Interval                   : 1024
Last Member Query Interval       : 1024
Query Response Interval          : 1023
Robust Count                     : 10
===============================================================================
A:BA
```

Table 75 provides IGMP status field descriptions.

Table 75 **IGMP Status Output Fields**

Label	Description
Admin State	Displays the administrative status of IGMP.
Oper State	Displays the current operating state of this IGMP protocol instance on this router.
Query Interval	The frequency at which IGMP query packets are transmitted.
Last Member Query Interval	The maximum response time inserted into Group-Specific Queries sent in response to Leave Group messages, and is also the amount of time between Group-Specific Query messages.
Query Response Interval	The maximum query response time advertised in IGMPv2 queries.
Robust Count	Displays the number of times the router will retry a query.

bgp-evpn

Syntax **bgp-evpn**

Context show>service>id

Description This command displays the bgp-evpn configured parameters for a specified service, including the admin status of vxlan, the configuration for mac-advertisement and unknown-mac-route as well as the mac-duplication parameters. The command shows the duplicate mac addresses that mac-duplication has detected.

Output

Sample Output

```
*A:DutA# show service id 1 bgp-evpn
===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
VXLAN Admin Status : Enabled            Creation Origin    : manual
MAC Dup Detn Moves : 5                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 9                  Number of Dup MACs : 1
-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses             Time Detected
-------------------------------------------------------------------------------
00:12:12:12:12:00                            01/17/2014 16:01:02
-------------------------------------------------------------------------------
===============================================================================
```

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the specified service.

lease-state

Syntax **lease-state** [[**sap** *sap-id*] | [**sdp** *sdp-id:vc-id*] | [**interface** *interface-name*] | [**ip-address** *ip-address*]] [**detail**]

Context show>service>id>dhcp

Description This command displays DHCP lease state related information.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

sdp-id — The SDP identifier

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information

Values 1 to 4294967295

interface-name — Displays information for the specified IP interface

ip-address — Displays information associated with the specified IP address

detail — Displays detailed information

Output

Sample Output

```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
===============================================================================
DHCP lease state table, service 13
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
13.13.40.1      00:00:00:00:00:13 1/1/1:13            00h00m58s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#


A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
===============================================================================
DHCP lease states for service 13
===============================================================================
```

```
Service ID           : 13
IP Address           : 13.13.40.1
Mac Address          : 00:00:00:00:00:13
Interface            : ies-13-13.13.1.1
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h00m58s
Persistence Key      : N/A

Sub-Ident            : "TEST"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : 13.13.254.254
Secondary-Dns        : 13.13.254.253

ServerLeaseStart     : 12/24/2006 23:44:07
ServerLastRenew      : 12/24/2006 23:44:07
ServerLeaseEnd       : 12/24/2006 23:45:07
Session-Timeout      : 0d 00:01:00
DHCP Server Addr     : N/A

Persistent Relay Agent Information
  Circuit Id         : ancstb6_Dut-A|13|ies-13-13.13.1.1|0|13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

Routed CO Output Example

```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
===============================================================================
DHCP lease state table, service 13
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
13.13.40.1      00:00:00:00:00:13 1/1/1:13            00h00m58s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#


A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
===============================================================================
DHCP lease states for service 13
===============================================================================
Service ID           : 13
```

```
IP Address           : 13.13.40.1
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : ies-13-13.13.1.1
Group-interface      : intf-13
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h00m58s
Persistence Key      : N/A

Sub-Ident            : "TEST"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : 13.13.254.254
Secondary-Dns        : 13.13.254.253

ServerLeaseStart     : 12/24/2006 23:48:23
ServerLastRenew      : 12/24/2006 23:48:23
ServerLeaseEnd       : 12/24/2006 23:49:23
Session-Timeout      : 0d 00:01:00
DHCP Server Addr     : N/A

Persistent Relay Agent Information
  Circuit Id         : ancstb6_Dut-A|13|intf-13|0|13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

Wholesaler/Retailer Output Example

```
A:ALA-_Dut-A# show service id 2000 dhcp lease-state detail
===============================================================================
DHCP lease states for service 2000
-------------------------------------------------------------------------------
Wholesaler 1000 Leases
-------------------------------------------------------------------------------
Service ID           : 1000
IP Address           : 13.13.1.254
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : whole-sub
Group-interface      : intf-13
Retailer             : 2000
Retailer If          : retail-sub
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h09m59s
Persistence Key      : N/A

Sub-Ident            : "TEST"
Sub-Profile-String   : "ADSL GO"
```

```
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Retail DHCP
Strings origin       : Retail DHCP
Lease Info origin    : Retail DHCP

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : N/A
Secondary-Dns        : N/A

ServerLeaseStart     : 12/25/2006 00:29:41
ServerLastRenew      : 12/25/2006 00:29:41
ServerLeaseEnd       : 12/25/2006 00:39:41
Session-Timeout      : 0d 00:10:00
DHCP Server Addr     : 10.232.237.2

Persistent Relay Agent Information
  Circuit Id         : 1/1/1:13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

statistics

Syntax **statistics** [**sap** *sap-id*
statistics [**sdp** *sdp-id*:*vc-id*]
statistics [**interface** *interface-name*]

Context show>service>id>dhcp

Description Displays DHCP statistics information.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

sdp-id — The SDP identifier

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information

Values 1 to 4294967295

interface-name — Displays information for the specified IP interface

summary

Syntax **summary**

Context show>service>id>dhcp

Description Displays DHCP configuration summary information.

Output

Sample Output

```
A:ALA-49# show service id 88 dhcp summary
===============================================================================
DHCP Summary, service 88
===============================================================================
Interface Name                    Arp      Used/                Info    Admin
  SapId/Sdp                       Populate Provided             Option  State
-------------------------------------------------------------------------------
Sector A                          No       0/0                  Keep    Up
  sap:7/1/1.2.2                            0/0
  sap:7/1/1.2.2                            0/0
grp-if                            No       0/1                  Keep    Down
  sap:2/2/2:0                              0/1
  sap:2/2/2:0                              0/1
test                              No       0/0                  Keep    Up
  sap:10/1/2:0                             0/0
  sap:10/1/2:0                             0/0
-------------------------------------------------------------------------------
Interfaces: 3
===============================================================================
A:ALA-49#
```

Table 76 describes the output fields for DHCP summary.

Table 76 **Show DHCP Summary Output Fields**

Label	Description
Interface Name	Name of the router interface.
Arp Populate	Specifies whether or not ARP populate is enabled.
Used/Provided	Used — The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the ‘Provided’ field.
	Provided — The lease-populate value that is configured for a specific interface.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

3.8.2.3.1 Show Multi-Chassis Endpoint Commands

endpoint

Syntax **endpoint** [*endpoint-name*]

Context show>service>id

Description This command displays service endpoint information.

Parameters *endpoint-name* — Specifies an endpoint name created in the **config>service>vpls** context

Output

Sample Output

```
*A:Dut-B# show service id 1 endpoint
===============================================================================
Service 1 endpoints
===============================================================================
Endpoint name                : mcep-t1
Description                  : (Not Specified)
Revert time                  : 0
Act Hold Delay               : 0
Ignore Standby Signaling     : false
Suppress Standby Signaling   : false
Block On Mesh Fail           : true
Multi-Chassis Endpoint       : 1
MC Endpoint Peer Addr        : 3.1.1.3
Psv Mode Active              : No
Tx Active                    : 231:1
Tx Active Up Time            : 0d 00:06:57
Revert Time Count Down       : N/A
Tx Active Change Count       : 5
Last Tx Active Change        : 02/13/2009 22:08:33
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 221:1 Prec:1                          Oper Status: Up
Spoke-sdp: 231:1 Prec:2                          Oper Status: Up
===============================================================================
*A:Dut-B#
```

etree

Syntax **etree**

Context show>service>id

Description This command displays the same information shown in the show service ID base context, with the addition of the role of each object in the VPLS E-Tree service.

The following labels identify the configuration of the SAPs and SDP bindings:

- (L) indicates leaf-ac
- (RL) indicates root-leaf-tag

Output

Sample Output

```
*A:DutA# show service id 2005 etree
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2005                Vpn Id            : 0
Service Type      : VPLS
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 07/08/2014 01:12:43
Last Mgmt Change  : 07/08/2014 01:12:30
Etree Mode        : Enabled
Admin State       : Up                  Oper State        : Up
MTU               : 1514                Def. Mesh VC Id   : 2005
SAP Count         : 5                   SDP Bind Count    : 1
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1:2005 (L)                      q-tag        1518    1518    Up   Up
sap:1/1/7:2006.200 (RL)                 qinq         9000    9000    Up   Up
sap:1/1/7:0.*                           qinq         9000    9000    Up   Up
sap:1/1/7:2005.*                        qinq         9000    9000    Up   Up
sap:1/1/8:1                             q-tag        1518    1518    Up   Up
sdp:12:2005 (RL) S(192.0.0.72)          Spok         0       8974    Up   Up
-------------------------------------------------------------------------------
Legend: (L): Leaf-Ac, (RL): Root-Leaf-Tag
===========================
```

multi-chassis

Syntax **multi-chassis**

Context show>redundancy

Description This command enables the context to display multi-chassis information.

mc-endpoint

Syntax **mc-endpoint statistics**
mc-endpoint peer [*ip-address*] **statistics**
mc-endpoint endpoint [*mcep-id*] **statistics**
mc-endpoint peer [*ip-address*]

Context show>redundancy>multi-chassis

Description This command displays multi-chassis endpoint information.

Parameters **statistics** — Displays the global statistics for the MC endpoint

ip-address — Specifies the IP address of multi-chassis endpoint peer

mcep-id — Specifies the multi-chassis endpoint

Values 1 to 4294967295

Output

Sample Output

```
*A:Dut-B# show redundancy multi-chassis mc-endpoint statistics
===============================================================================
Multi-Chassis Endpoint Global Statistics
===============================================================================
Packets Rx                                : 533
Packets Rx Keepalive                      : 522
Packets Rx Config                         : 3
Packets Rx Peer Config                    : 1
Packets Rx State                          : 7
Packets Dropped Keep-Alive Task           : 7
Packets Dropped Too Short                 : 0
Packets Dropped Verify Failed             : 0
Packets Dropped Tlv Invalid Size          : 0
Packets Dropped Out Of Seq                : 0
Packets Dropped Unknown Tlv               : 0
Packets Dropped Tlv Invalid MC-Endpoint Id  : 0
Packets Dropped MD5                       : 0
Packets Dropped Unknown Peer              : 0
Packets Dropped MC Endpoint No Peer       : 0
Packets Tx                                : 26099
Packets Tx Keepalive                      : 8221
Packets Tx Config                         : 2
Packets Tx Peer Config                    : 17872
Packets Tx State                          : 4
Packets Tx Failed                         : 0
===============================================================================
*A:Dut-B#

*A:Dut-B# show redundancy multi-chassis mc-endpoint peer 3.1.1.3 statistics
```

```
===============================================================================
Multi-Chassis MC-Endpoint Statistics
===============================================================================
Peer Addr                       : 3.1.1.3
-------------------------------------------------------------------------------
Packets Rx                      : 597
Packets Rx Keepalive            : 586
Packets Rx Config               : 3
Packets Rx Peer Config          : 1
Packets Rx State                : 7
Packets Dropped State Disabled  : 0
Packets Dropped Packets Too Short : 0
Packets Dropped Tlv Invalid Size  : 0
Packets Dropped Tlv Invalid LagId : 0
Packets Dropped Out of Seq      : 0
Packets Dropped Unknown Tlv     : 0
Packets Dropped MD5             : 0
Packets Tx                      : 636
Packets Tx Keepalive            : 600
Packets Tx Peer Config          : 30
Packets Tx Failed               : 0
Packets Dropped No Peer         : 0
===============================================================================
*A:Dut-B#

*A:Dut-B# show redundancy multi-chassis mc-endpoint  endpoint 1 statistics
===============================================================================
Multi-Chassis Endpoint Statistics
===============================================================================
===============================================================================
MC-Endpoint Id 1
===============================================================================
Packets Rx Config               : 3
Packets Rx State                : 7
Packets Tx Config               : 2
Packets Tx State                : 4
Packets Tx Failed               : 0
===============================================================================
Number of Entries 1
===============================================================================
*A:Dut-B#

*A:Dut-B# tools dump redundancy multi-chassis mc-endpoint peer 3.1.1.3
===============================================================================
MC Endpoint Peer Info
    peer addr                   : 3.1.1.3
    peer name                   : Dut-C
    peer name refs              : 1
    src addr conf               : Yes
    source addr                 : 2.1.1.2
    num of mcep                 : 1
    num of non-mcep             : 0
    own sess num                : 58ba0d39
    mc admin state              : Up
    tlv own mc admin state      : Up
    tlv peer mc admin state     : Up
    reachable                   : Yes

    own sys priority            : 50
```

```
    own sys id                    : 00:03:fa:72:c3:c0
    peer sys priority             : 21
    peer sys id                   : 00:03:fa:c6:31:f8
    master                        : No

    conf boot timer               : 300
    boot timer active             : No
    conf ka intv                  : 10
    conf hold on num of fail      : 3
    tlv own ka intv               : 10
    tlv peer ka intv              : 10
    ka timeout tmr active         : Yes
    ka timeout tmr intvl          : 20
    ka timeout tmr time left      : 4
    peer ka intv                  : 10
    mc peer timed out             : No

    initial peer conf rx          : Yes
    peer-mc disabled              : No
    initial peer conf sync        : Yes
    peer conf sync                : Yes

    own passive mode              : Disable
    peer passive mode             : No

    retransmit pending            : No
    non-mcep retransmit pending   : No
    retransmit intvl              : 5
    last tx time                  : 1437130
    last rx time                  : 1437156

    own bfd                       : Enable
    peer bfd                      : Enable
    bfd vrtr if                   : 2
    bfd handle                    : 1
    bfd state                     : 3
    bfd code                      : 0
===============================================================================
*A:Dut-B#


*A:Dut-B#   tools dump service mc-endpoint 1
===============================================================================
MC Endpoint Info
    mc-endpoint id                : 1
    endpoint                      : mcep-t1
    service                       : 1
    peer ref type                 : peer-name
    peer                          : Dut-C
    mc sel logic                  : peer selected active
    selection master              : No
    retransmit pending            : No
    initial config sync           : Yes
    config sync                   : Yes
    peer not mcep                 : No
    peer acked non-mcep           : No
    config mismatch               : No
    initial state rx              : Yes
    initial state sync            : Yes
```

```
    state sync                  : Yes
    can aggregate               : Yes
    sel peer active             : No
    peer sel active             : Yes
    passive mode active         : No
    own eligible force          : No
    own eligible double active  : Yes
    own eligible pw status bits : 0
    own eligible precedence     : 2
    own eligible conf chg       : No
    own eligible revert wait    : No
    peer eligible force         : No
    peer eligible double active : Yes
    peer eligible pw status bits: 0
    peer eligible precedence    : 3
    peer eligible conf chg      : No
    peer eligible revert wait   : No
===============================================================================
*A:Dut-B#


*A:Dut-B# tools perform service id 1 endpoint mcep-t1 force-switchover 221:1
*A:Dut-B>show#
*A:Dut-B# show service  id  1 endpoint
===============================================================================
Service 1 endpoints
===============================================================================
Endpoint name               : mcep-t1
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Ignore Standby Signaling    : false
Suppress Standby Signaling  : false
Block On Mesh Fail          : true
Multi-Chassis Endpoint      : 1
MC Endpoint Peer Addr       : 3.1.1.3
Psv Mode Active             : No
Tx Active                   : 221:1(forced)
Tx Active Up Time           : 0d 00:00:17
Revert Time Count Down      : N/A
Tx Active Change Count      : 6
Last Tx Active Change       : 02/14/2009 00:17:32
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 221:1 Prec:1                          Oper Status: Up
Spoke-sdp: 231:1 Prec:2                          Oper Status: Up
===============================================================================
*A:Dut-B#
```

3.8.2.4 VPLS Clear Commands

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

Parameters *service-id* — Clears service ID information that uniquely identifies a service.

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters in length.

arp-host

Syntax **arp-host**
arp-host { **mac** *ieee-address* | **sap** *sap-id* | **ip-address** *ip-address*[/*mask*] }
arp-host [**port** ***port-id***] [**inter-dest-id** ***intermediate-destination-id*** | **no-inter-dest-id**]
arp-host statistics [**sap** *sap-id* | **interface** *interface-name*]

Context clear>service>id

Description This command clears ARP host data.

authentication

Syntax **authentication**

Context clear>service>id

Description This command enables the context to clear session authentication information.

capture-sap

Syntax **capture-sap** *sap-id* [*trigger*]

Context clear>service>id

Description This command clears the statistics for a particular capture SAP.

cem

Syntax **cem**

Context clear>service>id

Description This command clears CEM statistics for this service.

statistics

Syntax **statistics**

Context clear>service>stats
clear>service>id>authentication

Description This command clears session statistics for this service.

fdb

Syntax **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id*] | **mesh-sdp** *sdp-id*[**:***vc-id*] | **spoke-sdp** *sdp-id***:***vc-id*}

Context clear>service>id

Description This command clears FDB entries for the service.

Parameters **all** — Clears all FDB entries

ieee-address — Clears only FDB entries in the FDB table with the specified 48-bit address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sap-id — Clears the physical port identifier portion of the SAP definition

mesh-sdp — Clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional.

spoke-sdp — Clears only service FDB entries associated with the specified spoke-SDP ID. For a spoke-SDP, the VC ID must be specified.

sdp-id — Specifies the SDP ID for which the associated FDB entries will be cleared

vc-id — Specifies the virtual circuit ID on the SDP ID for which the associated FDB entries will be cleared

Values

sdp-id[:vc-id]	*sdp-id*	1 to 17407
	vc-id	1 to 4294967295
sdp-id:vc-id	sdp-id	1 to 17407
	vc-id	1 to 4294967295

mld-snooping

Syntax **mld-snooping**

Context clear>service>id

Description This command enables the context to clear MLD snooping-related data.

port-db

Syntax **port-db sap** *sap-id* [**group** *grp-ipv6-address*]
port-db sap *sap-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*
port-db sdp *sdp-id*:*vc-id* [**group** *grp-ipv6-address*]
port-db sdp *sdp-id*:*vc-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*

Context clear>service>id>mld-snooping

Description This command clears MLD snooping port-db group data.

querier

Syntax **querier**

Context clear>service>id>mld-snooping

Description This command clears MLD snooping querier information.

statistics

Syntax **statistics all**
statistics sap *sap-id*
statistics sdp *sdp-id*:*vc-id*

Context clear>service>id>mld-snooping

Description This command clears MLD snooping statistics.

msap

Syntax **msap** *msap-id*

Context clear>service>id

Description This command clears the managed SAP (MSAP).

Parameters *msap-id* — Specifies the MSAP ID

Values dot1q*port-id* | lag-*id*:qtag1
qinq*port-id* | lag-*id*::qtag1.qtag2
qtag1 0 to 4094
qtag20 to 4094

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[**:***vc-id*] **ingress-vc-label**

Context clear>service>id

Description This command clears and resets the mesh SDP bindings for the service.

Parameters *sdp-id* — Clears mesh SDP ID information

Values 1 to 17407

vc-id — Specifies virtual circuit ID on the SDP ID to be reset

Default All VC IDs on the SDP ID

Values 1 to 4294967295

pim-snooping

Syntax **pim-snooping**

Context clear>service>id

Description This command enables the context to clear PIM snooping information.

database

Syntax **database** [[**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**group** *grp-ip-address*] [**source** *src-ip-address*]] [*family*]

Context clear>service>id>pim-snooping

Description This command clears PIM snooping source group database information.

Parameters *sap-id* — Clears PIM snooping SAP information

sdp-id — Clears PIM snooping entries associated with the specified SDP. For a spoke-SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

grp-ip-address — Clears PIM snooping information matching the specified group address

src-ip-address — Clears PIM snooping information matching one particular source within the multicast group

family — Displays either IPv4 or IPv6 information

Values ipv4 or ipv6

neighbor

Syntax **neighbor** [*ip-address* | **sap** *sap-id* | **sdp** *sdp-id:vc-id*] [*family*]

Context clear>service>id>pim-snooping

Description This command clears PIM snooping neighbor information.

Parameters *ip-address* — Clears information for the neighbor with the specified IP address

sap *sap-id* — Clears PIM snooping entries associated with the specified SAP

sdp *sdp-id:vc-id* — Clears PIM entries associated with the specified SDP. For a spoke-SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

family — Displays either IPv4 or IPv6 information

Values ipv4 or ipv6

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [*family*]

Context clear>service>id>pim-snooping

Description This command clears PIM snooping statistics for the specified SAP or SDP.

Parameters *sap-id* — Clears PIM snooping statistics for the specified SAP

sdp-id:vc-id — Clears PIM statistics for the specified SDP. For a spoke-SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

family — Displays either IPv4 or IPv6 information

Values ipv4 or ipv6

proxy-arp

Syntax **proxy-arp**
proxy-arp duplicate [*ip-address*]

proxy-arp dynamic [*ip-address*]

Context clear>service>id

Description This command allows all the duplicate or dynamic proxy-ARP entries to be cleared from the table. Individual IP entries can also be specified.

proxy-nd

Syntax **proxy-nd**
proxy-nd duplicate [*ipv6-address*]
proxy-nd dynamic [*ipv6-address*]

Context clear>service>id

Description This command allows all the duplicate or dynamic proxy-ND entries to be cleared from the table. Individual IPv6 entries can also be specified.

spoke-sdp

Syntax **spoke-sdp** *sdp-id* [:*vc-id*] {**all** | **counters** | **stp** | **l2pt**}}

Context clear>service>id

Description This command clears and resets the spoke-SDP bindings for the service.

Parameters *sdp-id* — Specifies the spoke-SDP ID to be reset

Values 1 to 17407

vc-id — Specifies the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

all — Clears all queue statistics and STP statistics associated with the SDP

counters — Clears all queue statistics associated with the SDP

stp — Clears all STP statistics associated with the SDP

l2pt — Clears all L2PT statistics associated with the SDP

sap

Syntax **sap** *sap-id* {**all** | **cem** | **counters** | **l2pt** | **stp** | **mrp**}

Context clear>service>statistics

Description This command clears statistics for the SAP bound to the service.

Parameters *sap-id* — Specifies the SDP ID for which the statistics will be cleared

all — Clears all queue statistics and STP statistics associated with the SAP

cem — Clears all CEM statistics associated with the SAP. This parameter only applies to the 7450 ESS or 7750 SR.

counters — Clears all queue statistics associated with the SAP

l2pt — Clears all L2PT statistics associated with the SAP. This parameter only applies to the 7450 ESS or 7750 SR.

stp — Clears all STP statistics associated with the SAP. This parameter only applies to the 7450 ESS or 7750 SR.

mrp — Clears all MRP statistics associated with the SAP. This parameter only applies to the 7450 ESS or 7750 SR.

sdp

Syntax **sdp** *sdp-id* [**keep-alive**]

Context clear>service>statistics

Description This command clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — The SDP ID for which the statistics will be cleared

Values 1 to 17407

keep-alive — Clears the keep-alive history associated with this SDP ID

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

l2pt

Syntax **l2pt**

Context clear>service>statistics>id

Description This command clears the l2pt statistics for this service.

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[*:vc-id*] {**all** | **counters** | **stp** | **mrp**}

Context clear>service>statistics>id

Description This command clears the statistics for a particular mesh SDP bind.

Parameters *sdp-id* — The spoke-SDP ID for which the statistics will be cleared

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

all — Clears all queue statistics and STP statistics associated with the SDP

counters — Clears all queue statistics associated with the SDP

stp — Clears all STP statistics associated with the SDP

mrp — Clears all MRP statistics associated with the SDP

mrp

Syntax **mrp**

Context clear>service>statistics>id

Description This command clears all MRP statistics for the service ID.

pip

Syntax **pip**

Context clear>service>statistics>id

Description This command clears the Provider Internal Port statistics for this service.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] {**all** | **counters** | **stp**| **l2pt** | **mrp**}

Context clear>service>statistics>id

Description This command clears statistics for the spoke-SDP bound to the service.

Parameters *sdp-id* — Specifies the spoke-SDP ID for which the statistics will be cleared

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

all — Clears all queue statistics and STP statistics associated with the SDP

counters — Clears all queue statistics associated with the SDP

stp — Clears all STP statistics associated with the SDP

l2pt — Clears all L2PT statistics associated with the SDP

mrp — Clears all MRP statistics associated with the SDP. This parameter only applies to the 7450 ESS or 7750 SR.

stp

Syntax **stp**

Context clear>service>statistics>id

Description Clears all spanning tree statistics for the service ID.

detected-protocols

Syntax **detected-protocols** {**all** | **sap** *sap-id* | **spoke-sdp** *sdp-id*[*:vc-id*]}

Context clear>service>id>stp

Description RSTP automatically falls back to STP mode when it receives an STP BPDU. The **clear detected-protocols** command forces the system to revert to the default RSTP mode on the SAP or spoke-SDP.

Parameters **all** — Clears all detected protocol statistics

sap-id — Clears the specified lease state SAP information

sdp-id — The SDP ID to be cleared. This parameter only applies to the 7450 ESS or 7750 SR.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be cleared. This parameter only applies to the 7450 ESS or 7750 SR.

Values 1 to 4294967295

lease-state

Syntax **lease-state** [**no-dhcp-release**]
lease-state ip-address *ip-address* [**no-dhcp-release**]

lease-state mac *ieee-address* **no-dhcp-release**
lease-state sap *sap-id* [**no-dhcp-release**]
lease-state sdp *sdp-id:vc-id* [**no-dhcp-release**]

Context clear>service>id>dhcp

Description This command clears DHCP lease state information.

Parameters **no-dhcp-release** — Specifies that the node will clear the state without sending the DHCP release message

ip-address — Clears the DHCP IP address lease state information. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets).

ieee-address — Clears DHCP MAC address lease state information. The 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap-id — Clears DHCP SAP lease state information

sdp-id — The SDP ID to be cleared

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be cleared

Values 1 to 4294967295

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** [*sdp-id*[*:vc-id*] | **interface** [*ip-address* | *ip-int-name*]]

Context clear>service>id>dhcp

Description This command clears DHCP statistics for this service.

Parameters *sap-id* — Clears the specified SAP statistics

sdp-id — Specifies the SDP ID to be cleared

Values 1 to 17407

vc-id — Specifies the virtual circuit ID on the SDP ID to be cleared

Values 1 to 4294967295

ip-int-name — Clears the statistics for the IP interface with the specified name

ip-addr — Clears the statistics for the IP interface with the specified IP address

igmp-snooping

Syntax **igmp-snooping**

Context clear>service>id

Description This command enables the context to clear IGMP snooping-related data.

statistics

Syntax **statistics** {**evpn-mpls** | **all** | **sap** *sap-id* | **sdp** *sdp-id:vc-id* | **vxlan vtep** *ip-address* **vni** *vni-id* }]

Context clear>service>id>igmp-snooping

Description This command clears IGMP snooping statistics for the VPLS service.

Parameters **all** — Clears the IGMP snooping information for all port objects in the service

evpn-mpls — Clears IGMP snooping statistics for EVPN-MPLS destinations

sap-id — Clears the IGMP snooping information on the specified SAP

sdp-id — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — Clears statistics for the specified virtual circuit ID on the SDP ID

Default For mesh SDPs only, all VC IDs

Values 1 to 4294967295

vxlan vtep *ip-address* **vni** <*1..16777215*> — Clears the IGMP snooping statistics associated with a specific VXLAN destination given by the VXLAN Termination Endpoint (VTEP) and VXLAN Network Identifier (VNI). This parameter only applies to the 7450 ESS or 7750 SR.

vni-id — Displays information for the specified VXLAN Network Identifier (VNI). This parameter only applies to the 7450 ESS or 7750 SR.

Values 1 to 16777215

port-db

Syntax **port-db** [**sap** *sap-id*] [**group** *grp-address* [**source** *ip-address*]]
port-db sdp *sdp-id:vc-id* [**group** *grp-address* [**source** *ip-address*]]
port-db detail vxlan vtep *ip-address* **vni** *vni-id*

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping port database for the VPLS service.

Parameters *sap-id* — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value

sdp-id — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — Clears information for the specified virtual circuit ID on the SDP ID

Default For mesh SDPs only, all VC IDs

Values 1 to 4294967295

grp-address — Clears IGMP snooping statistics matching the specified group address

ip-address — Clears IGMP snooping statistics matching the specified particular source

vxlan vtep *ip-address* **vni** <*1..16777215*> — Clears the IGMP snooping statistics associated with a specific VXLAN destination given by the VXLAN Termination Endpoint (VTEP) and VXLAN Network Identifier (VNI).

vni-id — Displays information for the specified VXLAN Network Identifier (VNI)

Values 1 to 16777215

querier

Syntax **querier**

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping queriers for the VPLS service.

mfib

Syntax **mfib**

Context clear>service>id>

Description This command enables the context to clear multicast FIB info for the VPLS service.

statistics

Syntax **statistics** {**all** | **ipv4** | **ipv6** | **mac**}
statistics group *grp-address*

Context clear>service>id>mfib

Description This command clears multicast FIB statistics for the VPLS service.

Parameters **all** — Clears all statistics for the service ID

ipv4 — Clears IPv4 address statistics for the service ID

ipv6 — Clears IPv6 address statistics for the service ID

mac — Clears MAC address statistics for the service ID

grp-address — Specifies an IGMP multicast group address that receives data on an interface

statistics

Syntax **statistics** {**all** | **sap** *sap-id* | **sdp** *sdp-id:vc-id* | **vxlan vtep** *ip-address* **vni** *vni-id*}

Context clear>service>id>igmp-snooping

Description This command clears IGMP snooping statistics.

Parameters **all** — Clears all statistics for the service ID

sap-id — Clears statistics for the specified SAP ID

sdp-id:vc-id — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values *sdp-id*:1 to 17407
vc-id:1 to 4294967295

vxlan vtep *ip-address* **vni** <*1..16777215*> — Clears the IGMP snooping statistics associated with a specific VXLAN destination given by the VXLAN Termination Endpoint (VTEP) and VXLAN Network Identifier (VNI).

vni-id — Displays information for the specified VXLAN Network Identifier (VNI)

Values 1 to 16777215

dhcp

Syntax **dhcp**

Context clear>router

Description This command enables the context to clear and reset DHCP entities.

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description This command clears DHCP statistics.

Parameters *ip-int-name* — Clears the statistics for the IP interface with the specified name

ip-addr — Clears the statistics for the IP interface with the specified IP address.

3.8.2.5 VPLS Debug Commands

id

Syntax **id** *service-id*

Context debug>service

Description This command debugs commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters in length.

arp-host

Syntax [**no**] arp-host

Context debug>service>id

Description This command enables and configures ARP host debugging.

The **no** form of the command disables ARP host debugging.

igmp-snooping

Syntax [**no**] **igmp-snooping**

Context debug>service>id

Description This command enables and configures IGMP-snooping debugging.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>igmp

Description This command enables and configures the IGMP tracing detail level.

The **no** form of the command disables the IGMP tracing detail level.

evpn-mpls

Syntax [**no**] **evpn-mpls**

Context debug>service>id>igmp-snooping

Description This command shows IGMP packets for EVPN-MPLS destinations. The **no** form of the command disables the debugging for EVPN-MPLS destinations

mac

Syntax [**no**] **mac** *ieee-address*

Context debug>service>id>igmp

Description This command shows IGMP packets for the specified MAC address.

The **no** form of the command disables the MAC debugging.

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>igmp

Description This command enables and configures the IGMP tracing mode.

The **no** form of the command disables the configures the IGMP tracing mode.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>igmp

Description This command shows IGMP packets for a specific SAP.

The **no** form of the command disables the debugging for the SAP.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id>igmp

Description This command shows IGMP packets for a specific SDP.

The **no** form of the command disables the debugging for the SDP.

Parameters *sdp-id* — Displays only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — Displays information for the specified virtual circuit ID on the SDP ID

Values 1 to 4294967295

vxlan

Syntax [**no**] **vxlan vtep** *vtep* **vni** *vni-id*

Context debug>service>id>igmp-snooping

Description This command shows IGMP packets for a specific VXLAN binding.

The **no** form of the command disables the debugging for that VXLAN binding.

Parameters **vtep** — IP address of the VXLAN Termination Endpoint

vni — VXLAN Network Identifier of the VXLAN binding

Values 1 to 16777215

mld-snooping

Syntax [**no**] **mld-snooping**

Context debug>service>id

Description This command enables and configures MLD-snooping debugging.

The **no** form of the command disables MLD-snooping debugging

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>mld

Description This command enables and configures the MLD tracing detail level.

The **no** form of the command disables the MLD tracing detail level.

mac

Syntax [**no**] **mac** *ieee-address*

Context debug>service>id>mld

Description This command shows MLD packets for the specified MAC address.

The **no** form of the command disables the MAC debugging.

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>mld

Description This command enables and configures the MLD tracing mode.

The **no** form of the command disables the configures the MLD tracing mode.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>mld

Description This command shows MLD packets for a specific SAP.

The **no** form of the command disables the debugging for the SAP.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id>mld

Description This command shows MLD packets for a specific SDP.

The **no** form of the command disables the debugging for the SDP.

Parameters *sdp-id* — Displays only MLD entries associated with the specified mesh SDP or spoke-SDP

Values 1 to 17407

vc-id — Displays information for the specified virtual circuit ID on the SDP ID

Values 1 to 4294967295

mrp

Syntax [**no**] **mrp**

Context debug>service>id

Description This command enables and configures MRP debugging.

all-events

Syntax **all-events**

Context debug>service>id>mrp

Description This command enables MRP debugging for the applicant, leave all, periodic and registrant state machines and enables debugging of received and transmuted MRP PDUs.

applicant-sm

Syntax [**no**] **applicant-sm**

Context debug>service>id>mrp

Description This command enables debugging of the applicant state machine.

The **no** form of the command disables debugging of the applicant state machine.

leave-all-sm

Syntax [**no**] **leave-all-sm**

Context debug>service>id>mrp

Description This command enables debugging of the leave all state machine.

The **no** form of the command disables debugging of the leave all state machine.

mmrp-mac

Syntax [**no**] **mmrp-mac** *ieee-address*

Context debug>service>id>mrp

Description This command filters debug events and only shows events related to the MAC address specified.

The **no** form of the command removes the debug filter.

Parameters *ieee-address* — xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)

mrpdu

Syntax [**no**] **mrpdu**

Context debug>service>id>mrp

Description This command enables debugging of the MRP PDUs that are received or transmitted.

The **no** form of the command disables debugging of MRP PDUs.

periodic-sm

Syntax [**no**] **periodic-sm**

Context debug>service>id>mrp

Description This command enables debugging of the periodic state machine.

The **no** form of the command disables debugging of the periodic state machine.

registrant-sm

Syntax [**no**] **registrant-sm**

Context debug>service>id>mrp

Description This command enables debugging of the registrant state machine.

The **no** form of the command disables debugging of the registrant state machine.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>mrp

Description This command filters debug events and only shows events for the particular SAP.

The **no** form of the command removes the debug filter.

Parameters *sap-id* — Specifies the SAP ID to show filter and debug events

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id>mrp

Description This command filters debug events and only shows events for the particular SDP.

The **no** form of the command removes the debug filter.

Parameters *sdp-id* — Displays only MLD entries associated with the specified mesh SDP or spoke-SDP

Values 1 to 17407

vc-id — Displays information for the specified virtual circuit ID on the SDP ID

Values 1 to 4294967295

event-type

Syntax [**no**] **event-type** {**config-change** | **svc-oper-status-change** | **sap-oper-status-change** | **sdpbind-oper-status-change**}

Context debug>service>id

Description This command enables a particular debugging event type.
The **no** form of the command disables the event type debugging.

Parameters **config-change** — Debugs configuration change events

svc-oper-status-change — Debugs service operational status changes

sap-oper-status-change — Debugs SAP operational status changes

sdpbind-oper-status-change — Debugs SDP operational status changes

host-connectivity-verify

Syntax [**no**] **host-connectivity-verify**

Context debug>service>id

Description This command enables Subscriber Host Connectivity Verification (SHCV) debugging.

The **no** form of the command disables the SHCV debugging.

ip

Syntax [**no**] **ip** *ip-address*

Context debug>service>id>host-connectivity-verify

Description This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular IP address.

Parameters *ip-address* — Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets).

mac

Syntax [**no**] **mac** *ieee-address*

Context debug>service>id>host-connectivity-verify

Description This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular MAC address.

Parameters *mac-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

pim-snooping

Syntax [**no**] **pim-snooping**

Context debug>service>id

Description This command enables PIM-snooping debugging.

adjacency

Syntax [**no**] **adjacency**

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for PIM adjacencies.

all

Syntax **all** [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**] **no all**

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for all the PIM modules.

Parameters *grp-ip-address* — Debugs information associated with all PIM modules

Values multicast group address (IPv4 or IPv6)

ip-address — Debugs information associated with all PIM modules

Values IPv4 or IPv6 address

detail — Debugs detailed information on all PIM modules

database

Syntax **database** [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]
no database

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for the PIM database.

Parameters *grp-ip-address* — Debugs information associated with all PIM modules

Values multicast group address (IPv4 or IPv6) or zero

ip-address — Debugs information associated with the specified database

jp

Syntax **jp** [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no jp

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for the PIM Join-Prune mechanism.

Parameters *grp-ip-address* — Debugs information associated with the specified Join-Prune mechanism.

Values multicast group address (ipv4 or ipv6) or zero

ip-address — Debugs information associated with the specified Join-Prune mechanism

Values source IP address (IPv4 or IPv6)

detail — Debugs detailed Join-Prune mechanism information

mcs

Syntax **mcs** [**detail**]
no mcs

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for PIM snooping multi-chassis synchronization.

Parameters **detail** — Provides detailed debugging information

packet

Syntax **packet** [**hello** | **jp**] [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
no packet

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for PIM packets.

Parameters **hello | jp** — PIM packet types

sap-id — Debugs packets associated with the specified SAP

sdp-id:vc-id — Debugs packets associated with the specified SDP

port

Syntax **port** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**detail**]
no port

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for PIM ports.

Parameters *sap-id* — Only debugs packets associated with the specified SAP

sdp-id:vc-id — Only debugs packets associated with the specified SDP

detail — Provides detailed debugging information

red

Syntax **red** [**detail**]
no red

Context debug>service>id>pim-snooping

Description This command enables or disables debugging for PIM messages sent to the standby CPM.

Parameters **detail** — Displays detailed debugging information

proxy-arp

Syntax **proxy-arp** [**mac** [*ieee-address*]] [***ip*** [*ipaddr*] **all**]]

Context debug>service>id

Description This command enables the debug of the proxy-arp function for a specified service. Alternatively, the debug can be enabled only for certain entries given by their IP or MAC addresses.

proxy-nd

Syntax **proxy-nd** [**mac** [*ieee-address*]] [***ip*** [*ipaddr*] **all**]]

Context debug>service>id

Description This command enables the debug of the proxy-nd function for a specified service. Alternatively, the debug can be enabled only for certain entries given by their IPv6 or MAC addresses.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>host-connectivity-verify

Description This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id

Description This command enables debugging for a particular SAP.

Parameters *sap-id* — Specifies the SAP ID

stp

Syntax **stp**

Context debug>service>id

Description This command enables the context for debugging STP.

all-events

Syntax **all-events**

Context debug>service>id>stp

Description This command enables STP debugging for all events.

bpdu

Syntax [**no**] **bpdu**

Context debug>service>id>stp

Description This command enables STP debugging for received and transmitted BPDUs.

core-connectivity

Syntax [**no**] **core-connectivity**

Context debug>service>id>stp

Description This command enables STP debugging for core connectivity.

exception

Syntax [**no**] **exception**

Context debug>service>id>stp

Description This command enables STP debugging for exceptions.

fsm-state-changes

Syntax [**no**] **fsm-state-changes**

Context debug>service>id>stp

Description This command enables STP debugging for FSM state changes.

fsm-timers

Syntax [**no**] **fsm-timers**

Context debug>service>id>stp

Description This command enables STP debugging for FSM timer changes.

port-role

Syntax [**no**] **port-role**

Context debug>service>id>stp

Description This command enables STP debugging for changes in port roles.

port-state

Syntax [**no**] **port-state**

Context debug>service>id>stp

Description This command enables STP debugging for port states.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>stp

Description This command enables STP debugging for a specific SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>stp

Description This command enables STP debugging for a specific SDP.

interface

Syntax [**no**] **interface** [*ip-int-name* | *ip-address*]

Context debug>router>igmp

Description This command enables debugging on the IGMP interface.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name

ip-address — Only displays the information associated with the specified IP address

mcs

Syntax [**no**] **mcs** [*ip-int-name*]

Context debug>router>igmp

Description This command enables debugging for IGMP MCS.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name

misc

Syntax [**no**] **misc**

Context debug>router>igmp

Description This command enables debugging for IGMP miscellaneous.

packet

Syntax [**no**] **packet** [**query** |**v1-report** | **v2-report** | **v3-report** | **v2-leave**] [*ip-int-name* | *ip-address*]

Context debug>router>igmp

Description This command enables debugging for IGMP packets.

Parameters **query v1/v2/v3-report, v2-leave** — Selects the type of packet to debug

ip-int-name — Debugs the information associated with the specified IP interface name

ip-address — Debugs the information associated with the specified IP address

card-status

Syntax **card-status**

Context tools>dump>service>id>fdb

Description This command displays the following MAC address information for each line card in the system:

- the number of allocated MAC addresses
- the number of pending MAC address allocations
- the number of pending free MAC addresses

Output

Sample Output

```
*A:PE1# tools dump service id 1 fdb card-status
===============================================================================
VPLS FDB Card Status at 01/31/2017 08:44:38
===============================================================================
Card                Allocated           PendAlloc           PendFree
-------------------------------------------------------------------------------
1                   4                   0                   0
2                   4                   0                   0
5                   4                   0                   0
===============================================================================
*A:PE1#
```

mac-status

Syntax **mac-status** [**mac** *ieee-address*] [**card** *slot-id*] [**pending**]

Context tools>dump>service>id>fdb

Description This command displays the status of MAC addresses within the service, displaying the line cards on which FDB entries are allocated for the MAC addresses (if a MAC address has been allocated an entry on all cards provisioned in the system, it is displayed as "All") and those for which there are pending FDB entry updates (allocate, displayed as "PendAlloc", or free, displayed as "PendFree") for each MAC address. The MAC address status is displayed per service or line card and for a single MAC address. In addition, only MAC addresses with pending updates can be displayed.

Parameters *ieee-address* — The 48-bit MAC address for which the FDB entry will be displayed in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers.

slot-id — The slot ID of the card in the chassis. The maximum slot ID is platform-dependent. See the hardware installation guides for more information.

pending — Displays only those MAC address with pending FDB entry line card updates (allocate or free).

Output

Sample Output

```
*A:PE1# tools dump service id 1 fdb mac-status
===============================================================================
VPLS FDB MAC status at 01/31/2017 08:44:39
===============================================================================
MAC Address         Type                Status : Card list
-------------------------------------------------------------------------------
00:00:00:00:01:01   Select              Allocated : 5
00:00:00:00:01:02   Select              Allocated : 5
00:00:00:00:01:03   Global              Allocated : All
00:00:00:00:01:04   Global              Allocated : All
===============================================================================
*A:PE1#
```

provider-tunnels

Syntax **provider-tunnels** [**type** {**originating** | **terminating**}

Context tools>dump>service>id

Description This command dumps the inclusive provider tunnels based on type.

Output

Sample Output

```
*A:Dut-C>tools# dump service vpls 1001 provider-tunnels type terminating
===============================================================================
VPLS 1001 Inclusive Provider Tunnels Terminating

===============================================================================
ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------
                                            1001    61440   10.20.1.1

                                            1001    64944   10.20.1.2
-------------------------------------------------------------------------------

*A:Dut-C>tools# dump service vpls 1001 provider-tunnels type originating
===============================================================================
VPLS 1001 Inclusive Provider Tunnels Originating

===============================================================================
ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------

ipmsi-1001-73728                            1001    61440   10.20.1.3

 -------------------------------------------------------------------------------

*A:Dut-C>tools# dump service vpls 1001 provider-tunnels
===============================================================================

VPLS 1001 Inclusive Provider Tunnels Originating
===============================================================================
ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------

ipmsi-1001-73728                            1001    61440   10.20.1.3

-------------------------------------------------------------------------------


===============================================================================

VPLS 1001 Inclusive Provider Tunnels Terminating

===============================================================================

ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------

                                            1001    61440   10.20.1.1

                                            1001    64944   10.20.1.2

 -------------------------------------------------------------------------------
```

```
*A:Dut-C>tools# dump service vpls 1001 provider-tunnels type terminating
===============================================================================
VPLS 1001 Inclusive Provider Tunnels Terminating

===============================================================================
ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------
                                            1001    61440   10.20.1.1

                                            1001    64944   10.20.1.2
-------------------------------------------------------------------------------

*A:Dut-C>tools# dump service vpls 1001 provider-tunnels type originating
===============================================================================
VPLS 1001 Inclusive Provider Tunnels Originating

===============================================================================
ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------

ipmsi-1001-73728                            1001    61440   10.20.1.3

 -------------------------------------------------------------------------------

*A:Dut-C>tools# dump service vpls 1001 provider-tunnels
===============================================================================

VPLS 1001 Inclusive Provider Tunnels Originating
===============================================================================
ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------

ipmsi-1001-73728                            1001    61440   10.20.1.3

-------------------------------------------------------------------------------


===============================================================================

VPLS 1001 Inclusive Provider Tunnels Terminating

===============================================================================

ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID

-------------------------------------------------------------------------------

                                            1001    61440   10.20.1.1

                                            1001    64944   10.20.1.2

 -------------------------------------------------------------------------------
```

proxy-arp

Syntax **proxy-arp usage**

Context tools>dump>service

Description This command provides information about the usage and limit of the system-wide proxy-arp table for all the services. The command also shows if the limit has been exceeded and a trap raised.

Output

Sample Output

```
*A:Dut# tools dump service proxy-arp usage
Proxy arp Usage
Current Usage       :         10
System Limit        :     511999
High Usage Trap Raised:       No
High Usage Threshold:         95 percent
High Usage Clear Threshold:   90 percent
```

proxy-nd

Syntax **proxy-nd usage**

Context tools>dump>service

Description This command provides information about the usage and limit of the system-wide proxy-nd table for all the services. The command also shows if the limit has been exceeded and a trap raised.

Output

Sample Output

```
*A:Dut# tools dump service proxy-nd usage
Proxy nd Usage
Current Usage       :         0
System Limit        :     511999
High Usage Trap Raised:       No
High Usage Threshold:         95 percent
High Usage Clear Threshold:   90 percent
```

vpls-fdb-stats

Syntax **vpls-fdb-stats**

Context tools>dump>service

Description This command provides the VPLS FDB statistics for all services.

Output

Sample Output

```
*A:PE1# tools dump service vpls-fdb-stats
Service Manager VPLS FDB info at 01/31/2017 08:44:40:
Statistics last cleared at 01/31/2017 07:42:25
      Statistic           |    Count
--------------------------+-------------
          FdbEntriesInUse |          8
          TotalFdbEntries |     511999
       FdbMimDestIdxInUse |          0
TotalFdbMimDestIdxEntries |      32767
          FdbIsidIdxInUse |          0
TotalFdbMimIsidIdxEntries |     191999
              MacAddMsgs  |         38
           MacDeleteMsgs  |          0
            MacQueryMsgs  |          0
             UnknownMsgs  |          0
           MalformedMsgs  |          0
              FailedMsgs  |          0
          FdbHwTableFull  |          0
       FdbHwLimitExceeded |          0
            FdbTableFull  |          0
         FdbLimitExceeded |          0
   FdbMimDestIdxExhausted |          0
              MacAddReqs  |         30
           DupMacAddReqs  |         19
        DroppedMacAddReqs |          0
         FailedMacAddReqs |         19
              MacDelReqs  |          0
           DupMacDelReqs  |          0
        DroppedMacDelReqs |          0
         FailedMacDelReqs |          0
   FailedMacCmplxMapUpdts |          0
 RvplsFdbEntriesAllocated |          0
     RvplsFdbEntriesInUse |          0
EsBmacFdbEntriesAllocated |          0
    EsBmacFdbEntriesInUse |          0
*A:PE1#
```

vxlan

Syntax **vxlan** [**clear**]

Context tools>dump>service

Description This command displays the number of times a service could not add a VXLAN binding or <VTEP, Egress VNI> due to the following limits:

- The per System VTEP limit has been reached
- The per System <VTEP, Egress VNI> limit has been reached

- The per Service <VTEP, Egress VNI> limit has been reached
- The per System Bind limit: Total bind limit or vxlan bind limit has been reached.

The command adds a [clear] option to clear the above statistics.

Output

Sample Output

```
*A:PE63# tools dump service id 3 vxlan
VTEP, Egress VNI Failure statistics at 000 00:03:55.710:
statistics last cleared at 000 00:00:00.000:
      Statistic            |    Count
---------------------------+-------------
                      VTEP |          0
              Service Limit |          0
               System Limit |          0
   Egress Mcast List Limit |          0
Duplicate VTEP, Egress VNI |          1
```

dup-vtep-egrvni

Syntax **dup-vtep-egrvni** [**clear**]

Context tools>dump>service>vxlan

Description This command dumps the <VTEP, VNI> bindings that have been detected as duplicate attempts, that is, an attempt to add the same binding to more than one service. The commands provides a ‘clear’ option.

Output

Sample Output

```
*A:PE71# tools dump service vxlan dup-vtep-egrvni
Duplicate VTEP, Egress VNI usage attempts at 000 00:03:41.570:
1. 10.1.1.1:100
```

usage

Syntax **usage**

Context tools>dump>service>vxlan

Description This command displays the consumed VXLAN resources in the system.

Output

Sample Output

```
*A:PE71# tools dump service vxlan usage
VXLAN usage statistics at 001 17:46:11.170:

VTEP                          :      5/8191
VTEP, Egress VNI              :      5/131071
Sdp Bind + VTEP, Egress VNI   :     13/196607
RVPLS Egress VNI              :      0/40959
```

4 IEEE 802.1ah Provider Backbone Bridging

4.1 IEEE 802.1ah Provider Backbone Bridging (PBB) Overview

IEEE 802.1ah draft standard (IEEE802.1ah), also known as Provider Backbone Bridges (PBB), defines an architecture and bridge protocols for interconnection of multiple Provider Bridge Networks (PBNs - IEEE802.1ad QinQ networks). PBB is defined in IEEE as a connectionless technology based on multipoint VLAN tunnels. IEEE 802.1ah employs Provider MSTP as the core control plane for loop avoidance and load balancing. As a result, the coverage of the solution is limited by STP scale in the core of large service provider networks.

Virtual Private LAN Service (VPLS), RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, provides a solution for extending Ethernet LAN services using MPLS tunneling capabilities through a routed, traffic-engineered MPLS backbone without running (M)STP across the backbone. As a result, VPLS has been deployed on a large scale in service provider networks.

The Nokia implementation fully supports a native PBB deployment and an integrated PBB-VPLS model where desirable PBB features such as MAC hiding, service aggregation and the service provider fit of the initial VPLS model are combined to provide the best of both worlds.

4.2 PBB Features

This section provides information about PBB features.

4.2.1 Integrated PBB-VPLS Solution

HVPLS introduced a service-aware device in a central core location in order to provide efficient replication and controlled interaction at domain boundaries. The core network facing provider edge (N-PE) devices have knowledge of all VPLS services and customer MAC addresses for local and related remote regions resulting in potential scalability issues as depicted in Figure 106.

Figure 106 **Large HVPLS Deployment**

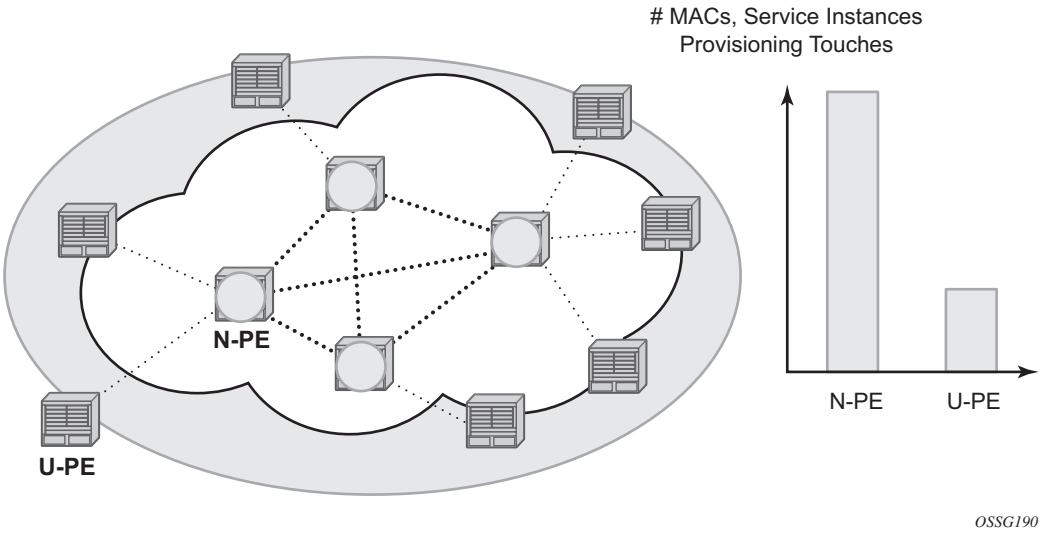

In a large VPLS deployment, it is important to improve the stability of the overall solution and to speed up service delivery. These goals are achieved by reducing the load on the N-PEs and respectively minimizing the number of provisioning touches on the N-PEs.

The integrated PBB-VPLS model introduces an additional PBB hierarchy in the VPLS network to address these goals as depicted in Figure 107.

Figure 107 **Large PBB-VPLS Deployment**

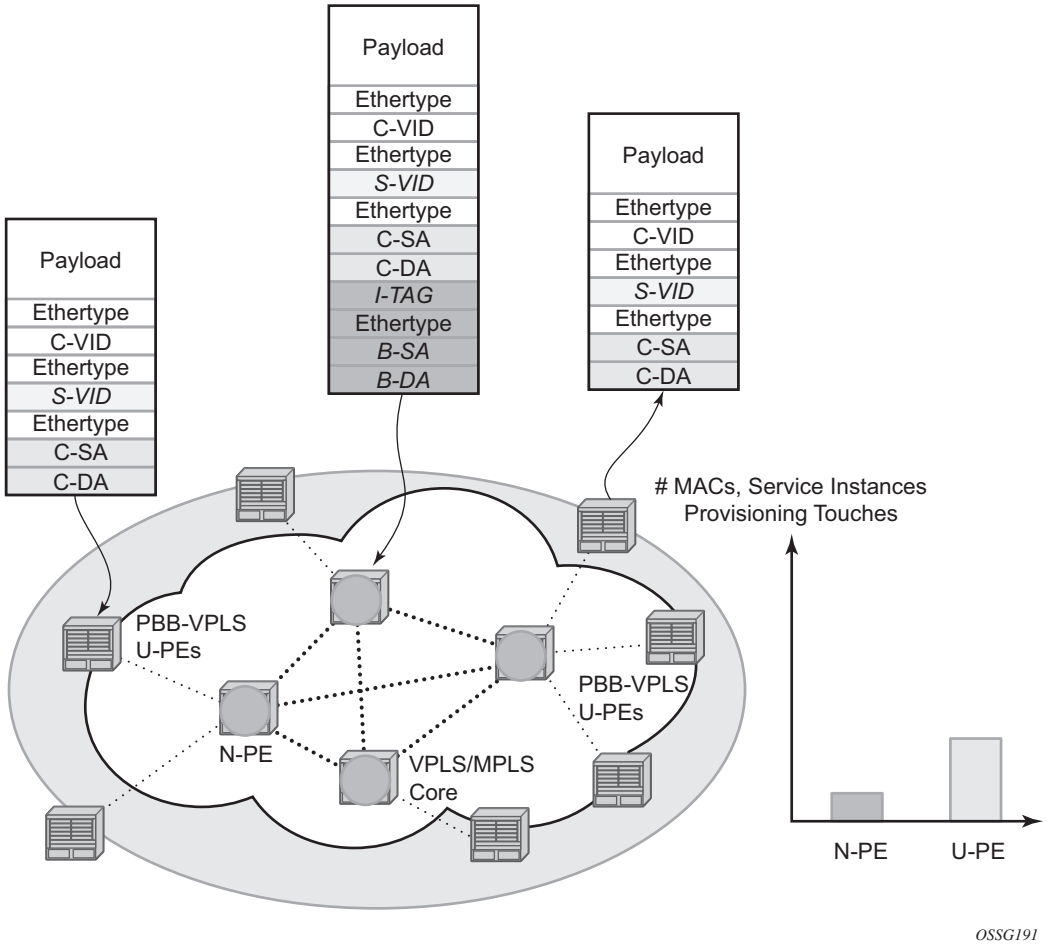

PBB encapsulation is added at the user facing PE (U-PE) to hide the customer MAC addressing and topology from the N-PE devices. The core N-PEs

need to only handle backbone MAC addressing and do not need to have visibility of each customer VPN. As a result, the integrated PBB-VPLS solution decreases the load in the N-PEs and improves the overall stability of the backbone.

The Nokia PBB-VPLS solution also provides automatic discovery of the customer VPNs through the implementation of IEEE 802.1ak MMRP minimizing the number of provisioning touches required at the N-PEs.

4.2.2 PBB Technology

IEEE 802.1ah specification encapsulates the customer or QinQ payload in a provider header as shown in Figure 108.

Figure 108 **QinQ Payload in Provider Header Example**

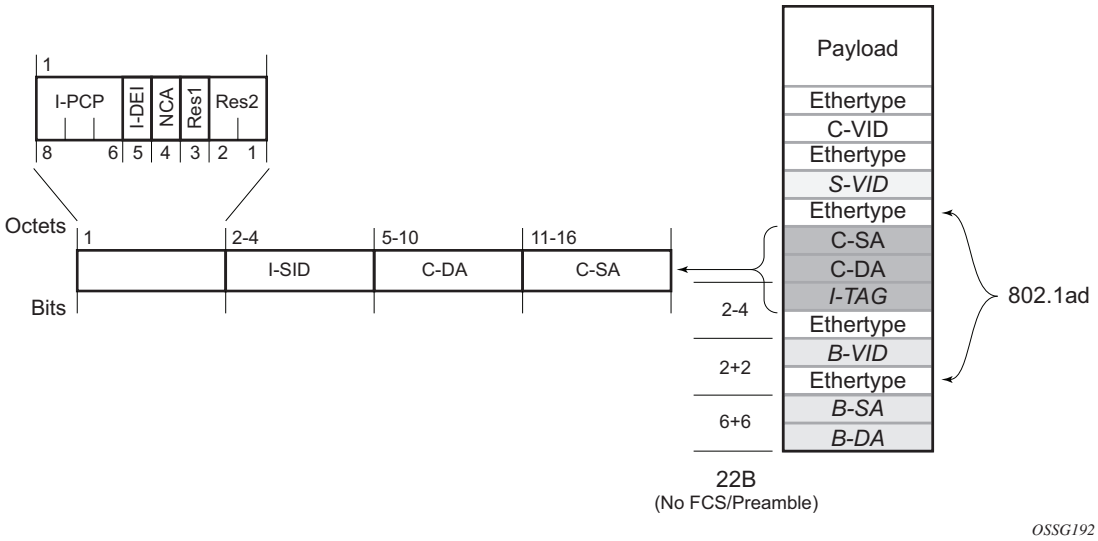

PBB adds a regular Ethernet header where the B-DA and B-SA are the backbone destination and respectively, source MACs of the edge U-PEs. The backbone MACs (B-MACs) are used by the core N-PE devices to switch the frame through the backbone.

A special group MAC is used for the backbone destination MAC (B-DA) when handling an unknown unicast, multicast or broadcast frame. This backbone group MAC is derived from the I-service instance identifier (ISID) using the rule: a standard group OUI (01-1E-83) followed by the 24 bit ISID coded in the last three bytes of the MAC address.

The BVID (backbone VLAN ID) field is a regular DOT1Q tag and controls the size of the backbone broadcast domain. When the PBB frame is sent over a VPLS pseudowire, this field may be omitted depending on the type of pseudowire used.

The following ITAG (standard Ether-type value of 0x88E7) has the role of identifying the customer VPN to which the frame is addressed through the 24 bit ISID. Support for service QoS is provided through the priority (3 bit I-PCP) and the DEI (1 bit) fields.

4.2.3 PBB Mapping to Existing VPLS Configurations

The IEEE model for PBB is organized around a B-component handling the provider backbone layer and an I-component concerned with the mapping of the customer/ provider bridge (QinQ) domain (MACs, VLANs) to the provider backbone (B-MACs, B-VLANs): for example, the I-component contains the boundary between the customer and backbone MAC domains.

The Nokia implementation is extending the IEEE model for PBB to allow support for MPLS pseudowires using a chain of two VPLS context linked together as depicted in Figure 109.

Figure 109 **PBB Mapping to VPLS Configurations**

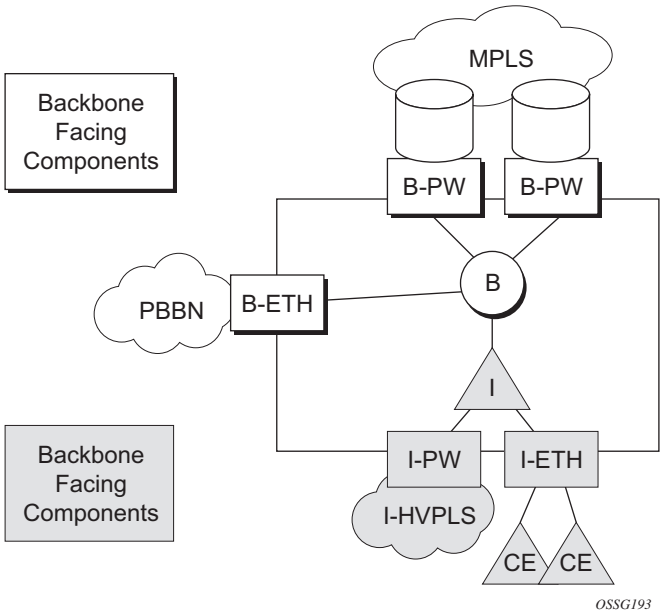

A VPLS context is used to provide the backbone switching component. The white circle marked B, referred to as backbone-VPLS (B-VPLS), operates on backbone MAC addresses providing a core multipoint infrastructure that may be used for one or multiple customer VPNs. The Nokia B-VPLS implementation allows the use of both native PBB and MPLS infrastructures.

Another VPLS context (I-VPLS) can be used to provide the multipoint I-component functionality emulating the E-LAN service (see the triangle marked “I” in Figure 109). Similar to B-VPLS, I-VPLS inherits from the regular VPLS the pseudowire (SDP bindings) and native Ethernet (SAPs) handoffs accommodating this way different types of access: for example, direct customer link, QinQ or HVPLS.

To support PBB E-Line (point-to-point service), the use of an Epipe as I-component is allowed. All Ethernet SAPs supported by a regular Epipe are also supported in the PBB Epipe.

4.2.4 SAP and SDP Support

This section provides information about SAP and SDP support.

4.2.4.1 PBB B-VPLS

- SAPs
  - Ethernet DOT1Q and QinQ are supported — This is applicable to most PBB use cases, for example, one backbone VLAN ID used for native Ethernet tunneling. In the case of QinQ, a single tag x is supported on a QinQ encapsulation port for example (1/1/1:x.* or 1/1/1:x.0).
  - Ethernet null is supported — This is supported for a direct connection between PBB PEs, for example, no BVID is required.
  - Default SAP types are blocked in the CLI for the B-VPLS SAP.
- The following rules apply to the SAP processing of PBB frames:
  - For “transit frames” (not destined for a local BMAC), there is no need to process the ITAG component of the PBB frames. Regular Ethernet SAP processing is applied to the backbone header (BMACs and BVID).
  - If a local I-VPLS instance is associated with the B-VPLS, “local frames” originated/terminated on local I-VPLS(s) are PBB encapsulated/de-encapsulated using the **pbb-etype** provisioned under the related port or SDP component.
- SDPs
  - For MPLS, both mesh and spoke-SDPs with split horizon groups are supported.
  - Similar to regular pseudowire, the outgoing PBB frame on an SDP (for example, B-pseudowire) contains a BVID qtag only if the pseudowire type is Ethernet VLAN. If the pseudowire type is ‘Ethernet’, the BVID qtag is stripped before the frame goes out.

4.2.4.2 PBB I-VPLS

- Port Level
  - All existing Ethernet encapsulation types are supported (for example, null, dot1q, qinq).
- SAPs
  - The I-VPLS SAPs can co-exist on the same port with SAPs for other business services, for example, VLL, VPLS SAPs.
  - All existing Ethernet encapsulation are supported: null, dot1q, qinq.
- SDPs
  - GRE and MPLS SDP are spoke-sdp only. Mesh SDPs can just be emulated by using the same split horizon group everywhere.

Existing SAP processing rules still apply for the I-VPLS case; the SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs to:

- Null encap defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP;
- dot1q encap defined on ingress — only first VLAN tag is considered;
- Qinq encap defined on ingress — both VLAN tags are considered; wildcard support for the inner VLAN tag
- For dot1q/qinq encapsulations, traffic encapsulated with VLAN tags for which there is no definition is discarded.
- Any VLAN tag used for service selection on the I-SAP is stripped before the PBB encapsulation is added. Appropriate VLAN tags are added at the remote PBB PE when sending the packet out on the egress SAP.

I-VPLS services do not support the forwarding of PBB encapsulated frames received on SAPs or Spoke-SDPs through their associated B-VPLS service. PBB frames are identified based on the configured PBB Ethertype (0x88e7 by default).

4.2.5 PBB Packet Walkthrough

This section describes the walkthrough for a packet that traverses the B-VPLS and I-VPLS instances using the example of a unicast frame between two customer stations as depicted in the following network diagram Figure 110.

Figure 110 **PBB Packet Walkthrough**

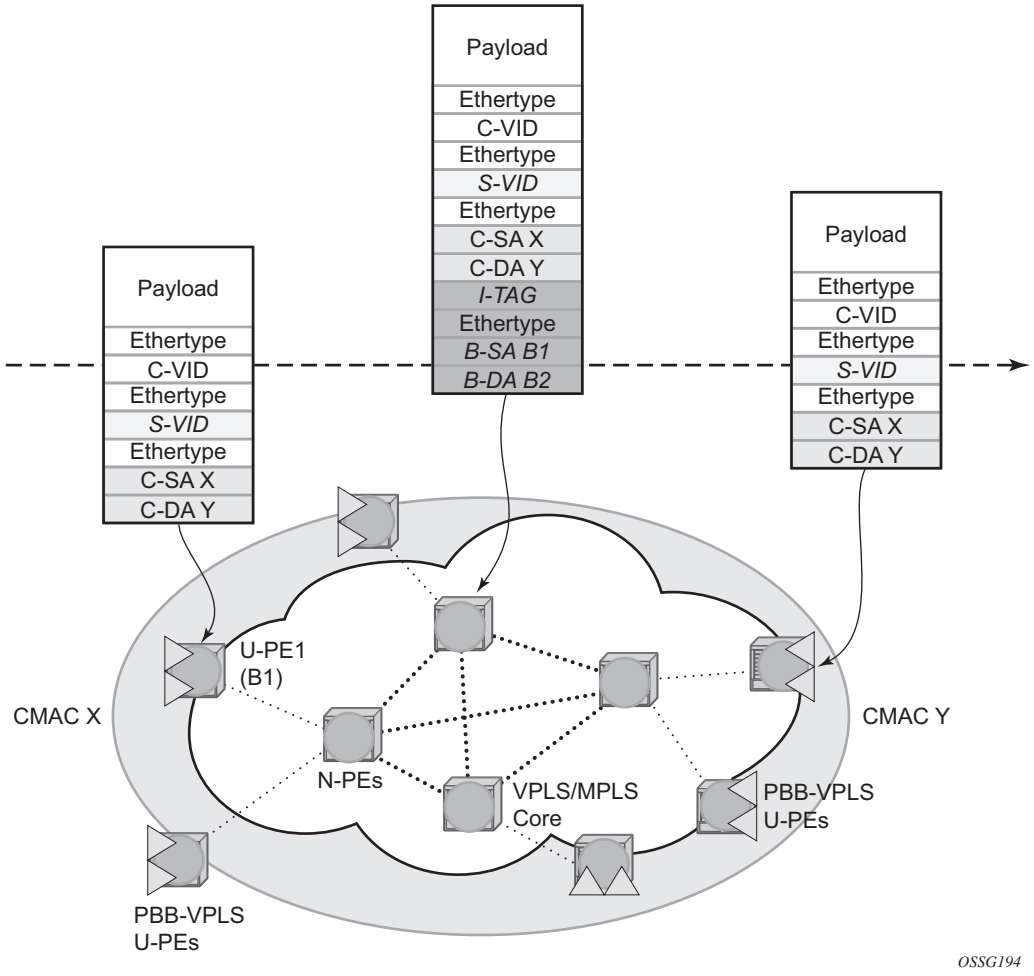

The station with CMAC (customer MAC) X wants to send a unicast frame to CMAC Y through the PBB-VPLS network. A customer frame arriving at PBB-VPLS U-PE1 is encapsulated with the PBB header. The local I-VPLS FDB on U-PE1 is consulted to determine the destination BMAC of the egress U-PE for CMAC Y. In our example, B2 is assumed to be known as the B-DA for Y. If CMAC Y is not present in the U-PE1 forwarding database, the PBB packet is sent in the B-VPLS using the standard group MAC address for the ISID associated with the customer VPN. If the uplink to the N-PE is a spoke pseudowire, the related PWE3 encapsulation is added in front of the B-DA.

Next, only the Backbone Header in green is used to switch the frame through the green B-VPLS/VPLS instances in the N-PEs. At the receiving U-PE2, the CMAC X is learned as being behind BMAC B1; then the PBB encapsulation is removed and the lookup for CMAC Y is performed. In the case where a pseudowire is used between N-PE and U-PE2, the pseudowire encapsulation is removed first.

4.2.5.1 PBB Control Planes

PBB technology can be deployed in a number of environments. Natively, PBB is an Ethernet data plane technology that offers service scalability and multicast efficiency.

Environment:

- MPLS (mesh and spoke-SDPs)
- Ethernet SAPs

Within these environments, SR OS offers a number of optional control planes:

- Shortest Path Bridging MAC (SPBM) (SAPs and spoke-SDPs); see Shortest Path Bridging MAC Mode (SPBM)
- Rapid Spanning Tree Protocol (RSTP) optionally with MMRP (SAPs and spoke-SDPs); see MMRP Support Over B-VPLS SAPs and SDPs.
- MSTP optionally with MMRP (SAPs and spoke-SDPs); see the *7450 ESS, 7750 SR, and 7950 XRS Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information.
- Multiple MAC registration Protocol (MMRP) alone (SAPs, spoke and mesh SDPs); see IEEE 802.1ak MMRP for Service Aggregation and Zero Touch Provisioning.

In general a control plane is required on Ethernet SAPs, or SDPs where there could be physical loops. Some network configurations of Mesh and Spoke SDPs can avoid physical loops and no control plane is required.

The choice of control plane is based on the requirement of the networks. SPBM for PBB offers a scalable link state control plane without BMAC flooding and learning or MMRP. RSTP and MSTP offer Spanning tree options based on BMAC flooding and learning. MMRP is used with flooding and learning to improve multicast.

4.2.6 Shortest Path Bridging MAC Mode (SPBM)

Shortest Path Bridging (SPB) enables a next generation control plane for PBB based on IS-IS that adds the stability and efficiency of link state to unicast and multicast services. Specifically this is an implementation of SPBM (SPB MAC mode). Current SR OS PBB B-VPLS offers point-to-point and multipoint to multipoint services with large scale. PBB B-VPLS is deployed in both Ethernet and MPLS networks supporting Ethernet VLL and VPLS services. SPB removes the flooding and learning mode from the PBB Backbone network and replaces MMRP for ISID Group Mac Registration providing flood containment. SR OS SPB provides true shortest path forwarding for unicast and efficient forwarding on a single tree for multicast. It supports selection of shortest path equal cost tie-breaking algorithms to enable diverse forwarding in an SPB network.

4.2.6.1 Flooding and Learning Versus Link State

SPB brings a link state capability that improves the scalability and performance for large networks over the xSTP flooding and learning models. Flooding and learning has two consequences. First, a message invoking a flush must be propagated, second the data plane is allowed to flood and relearn while flushing is happening. Message based operation over these data planes may experience congestion and packet loss.

Table 77 **B-VPLS Control Planes**

PBB B-VPLS Control Plane	Flooding and Learning	Multipath	Convergence time
xSTP	Yes	MSTP	xSTP + MMRP
G.8032	Yes	Multiple Ring instances Ring topologies only	Eth-OAM based + MMRP
SPB-M	No	Yes –ECT based	IS-IS link state (incremental)

Link state operates differently in that only the information that truly changes needs to be updated. Traffic that is not affected by a topology change does not have to be disturbed and does not experience congestion since there is no flooding. SPB is a link state mechanism that uses restoration to reestablish the paths affected by topology change. It is more deterministic and reliable than RSTP and MMRP mechanisms. SPB can handle any number of topology changes and as long as the network has some connectivity, SPB will not isolate any traffic.

4.2.6.2 SPB for B-VPLS

The SR OS model supports PBB Epipes and I-VPLS services on the B-VPLS. SPB is added to B-VPLS in place of other control planes (see Table 77). SPB runs in a separate instance of IS-IS. SPB is configured in a single service instance of B-VPLS that controls the SPB behavior (via IS-IS parameters) for the SPB IS-IS session between nodes. Up to four independent instances of SPB can be configured. Each SPB instance requires a separate control B-VPLS service. A typical SPB deployment uses a single control VPLS with zero, one or more user B-VPLS instances. SPB is multi-topology (MT) capable at the IS-IS LSP TLV definitions however logical instances offer the nearly the same capability as MT. The SR OS SPB implementation always uses MT topology instance zero. 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.

SPB Link State PDUs (LSPs) contains BMACs, I-SIDs (for multicast services) and link and metric information for an IS-IS database. Epipe I-SIDs are not distributed in SR OS SPB allowing high scalability of PBB Epipes. I-VPLS I-SIDs are distributed in SR OS SPB and the respective multicast group addresses are automatically populated in forwarding in a manner that provides automatic pruning of multicast to the subset of the multicast tree that supports I-VPLS with a common I-SID. This replaces the function of MMRP and is more efficient than MMRP so that in the future SPB will scale to a greater number of I-SIDs.

SPB on SR OS can leverage MPLS networks or Ethernet networks or combinations of both. SPB allows PBB to take advantage of multicast efficiency and at the same time leverage MPLS features such as resiliency.

4.2.6.3 Control B-VPLS and User B-VPLS

Control B-VPLS are required for the configuration of the SPB parameters and as a service to enable SPB. Control B-VPLS therefore must be configured everywhere SPB forwarding is expected to be active even if there are no terminating services. SPB uses the logical instance and a Forwarding ID (FID) to identify SPB locally on the node. The FID is used in place of the SPB VLAN identifier (Base VID) in IS-IS LSPs enabling a reference to exchange SPB topology and addresses. More specifically, SPB advertises B-MACs and I-SIDs in a B-VLAN context. Since the service model in SR OS separates the VLAN Tag used on the port for encapsulation from the VLAN ID used in SPB the SPB VLAN is a logical concept and is represented by configuring a FID. B-VPLS SAPs use VLAN Tags (SAPs with Ethernet encapsulation) that are independent of the FID value. The encapsulation is local to the link in SR/ESS so the SAP encapsulation has be configured the same between

neighboring switches. The FID for a specified instance of SPB between two neighbor switches must be the same. The independence of VID encapsulation is inherent to SR OS PBB B-VPLS. This also allows spoke-SDP bindings to be used between neighboring SPB instances without any VID tags. The one exception is mesh SDPs are not supported but arbitrary mesh topologies are supported by SR OS SPB.

Figure 111 illustrates two switches where an SPB control B-VPLS configured with FID 1 and uses a SAP with 1/1/1:5 therefore using a VLAN Tag 5 on the link. The SAP 1/1/1:1 could also have been be used but in SR OS the VID does not have to equal FID. Alternatively an MPLS PW (spoke-SDP binding) could be for some interfaces in place of the SAP. Figure 111 illustrates a control VPLS and two user B-VPLS. The User B-VPLS must share the same topology and are required to have interfaces on SAPs/Spoke SDPs on the same links or LAG groups as the B-VPLS. To allow services on different B-VPLS to use a path when there are multiple paths a different ECT algorithm can be configured on a B-VPLS instance. In this case, the user B-VPLS still fate shared the same topology but they may use different paths for data traffic; see Shortest Path and Single Tree.

Figure 111 **Control and User B-VPLS with FIDs**

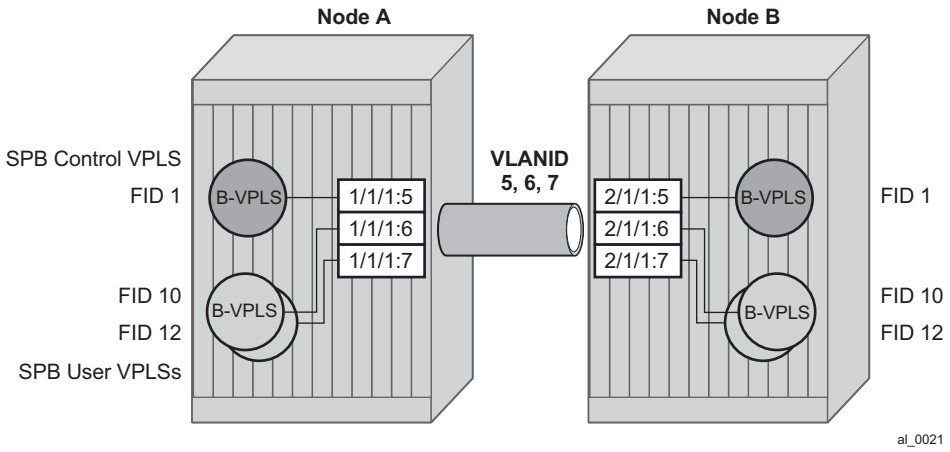

Each user BVPLS offers the same service capability as a control B-VPLS and are configured to “follow” or fate share with a control B-VPLS. User B-VPLS must be configured as active on the whole topology where control B-VPLS is configured and active. If there is a mismatch between the topology of a user B-VPLS and the control B-VPLS, only the user B-VPLS links and nodes that are in common with the control B-VPLS will function. The services on any B-VPLS are independent of a particular user B-VPLS so a mis-configuration of one of the user B-VPLS will not affect other B-VPLS. For example if a SAP or spoke-SDP is missing in the user B-VPLS any traffic from that user B-VPLS that would use that interface, will be missing forwarding information and traffic will be dropped only for that B-VPLS. The computation of paths is based only on the control B-VPLS topology.

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 short cut 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 FDB. If I-VPLS services are to be supported on a B-VPLS that B-VPLS must have an independent FID.

4.2.6.4 Shortest Path and Single Tree

IEEE 802.1aq standard SPB uses a source specific tree model. The standard model is more computationally intensive for multicast traffic since in addition to the SPF algorithm for unicast and multicast from a single node, an all pairs shorted path needs to be computed for other nodes in the network. In addition, the computation must be repeated for each ECT algorithm. While the standard yields efficient shortest paths, this computation is overhead for systems where multicast traffic volume is low. Ethernet VLL and VPLS unicast services are popular in PBB networks and the SR OS SPB design is optimized for unicast delivery using shortest paths. Ethernet supporting unicast and multicast services are commonly deployed in Ethernet transport networks. SR OS SPB Single tree multicast (also called shared tree or *,G) operates similarly today. The difference is that SPB multicast never floods unknown traffic.

The SR OS implementation of SPB with shortest path unicast and single tree multicast, requires only two SPF computations per topology change reducing the computation requirements. One computation is for unicast forwarding and the other computation is for multicast forwarding.

A single tree multicast requires selecting a root node much like RSTP. Bridge priority controls the choice of root node and alternate root nodes. The numerically lowest Bridge Priority is the criteria for choosing a root node. If multiple nodes have the same Bridge Priority, then the lowest Bridge Identifier (System Identifier) is the root.

In SPB the source-bmac can override the chassis-mac allowing independent control of tie breaking, The shortest path unicast forwarding does not require any special configuration other than selecting the ECT algorithm by configuring a B-VPLS use a FID with low-path-id algorithm or high-path-id algorithm to tie break between equal cost paths. Bridge priority allows some adjustment of paths. Configuring link metrics adjusts the number of equal paths.

To illustrate the behavior of the path algorithms a sample network is shown in Figure 112.

Figure 112 **Sample Partial Mesh network**

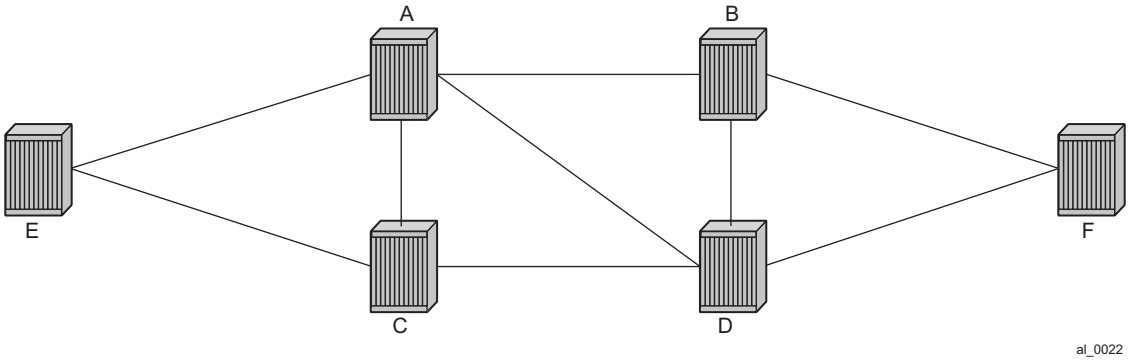

Assume that Node A is the lowest Bridge Identifier and the Multicast root node and all links have equal metrics. Also, assume that Bridge Identifiers are ordered such that Node A has a numerically lower Bridge identifier than Node B, and Node B has lower Bridge Identifier than Node C, etc. Unicast paths are configured to use shortest path tree (SPT). Figure 113 shows the shortest paths computed from Node A and Node E to Node F. There are only two shortest paths from A to F. A choice of low-path-id algorithm uses Node B as transit node and a path using high-path-id algorithm uses Node D as transit node. The reverse paths from Node F to A are the same (all unicast paths are reverse path congruent). For Node E to Node F there are three paths E-A-B-F, E-A-D-F, and E-C-D-F. The low-path-id algorithm uses path E-A-B-F and the high-path-id algorithm uses E-C-D-F. These paths are also disjoint and are reverse path congruent. Any nodes that are directly connected in this network have only one path between them (not shown for simplicity).

Figure 113 **Unicast Paths for Low-path-id and High-path-id**

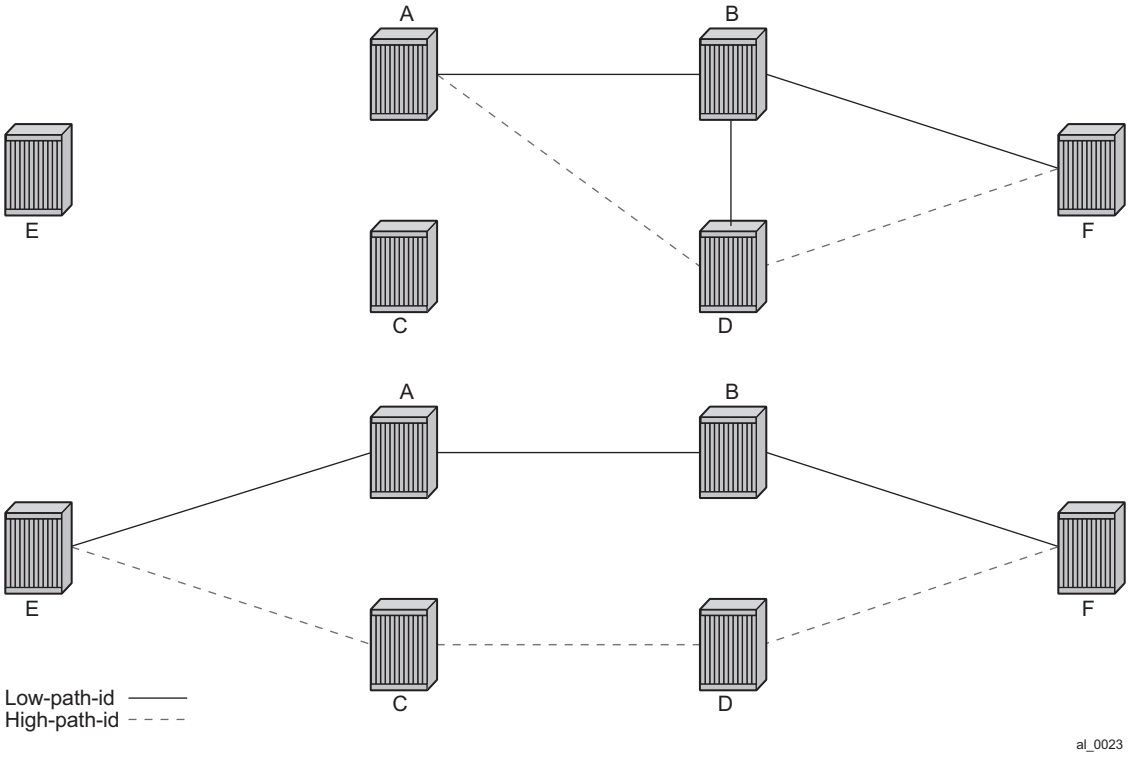

For Multicast paths the algorithms used are the same low-path-id or high-path-id but the tree is always a single tree using the root selected as described earlier (in this case Node A). Figure 114 illustrates the multicast paths for low-path-id and high-path-id algorithm.

Figure 114 **Multicast Paths for Low-path-id and High-path-id**

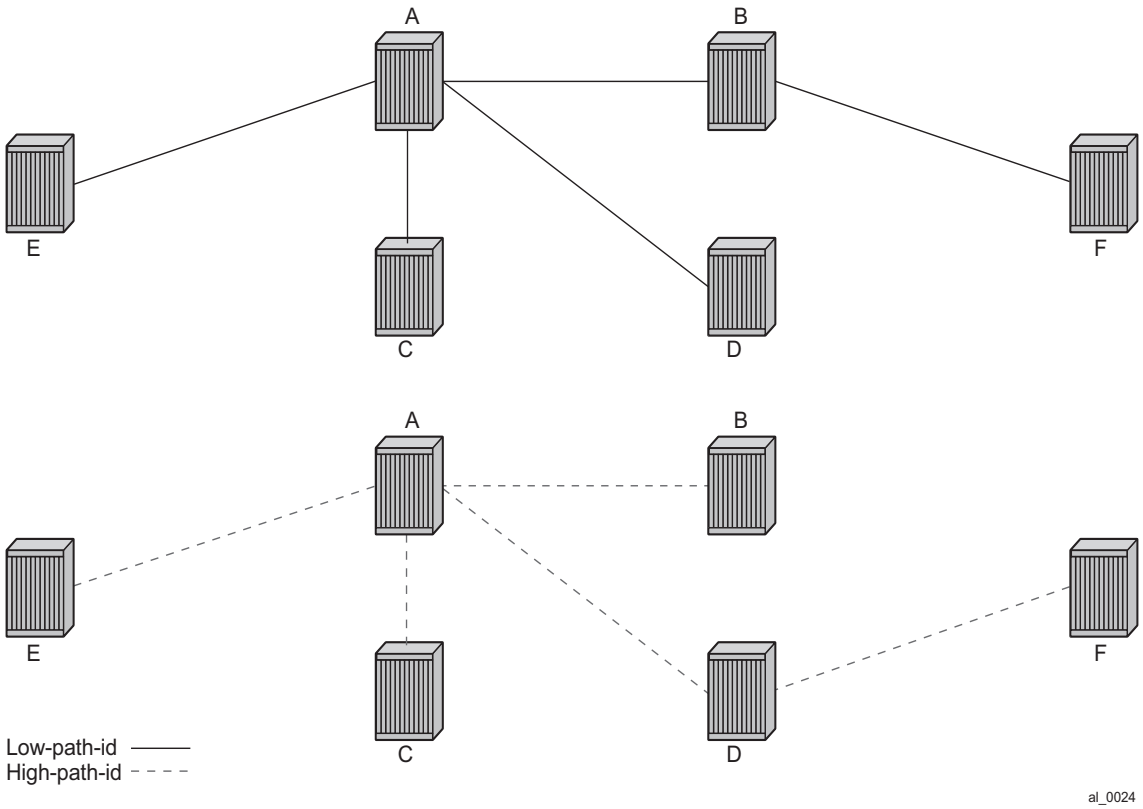

All nodes in this network use one of these trees. The path for multicast to/from Node A is the same as unicast traffic to/from Node A for both low-path-id and high-path-id. However, the multicast path for other nodes is now different from the unicast paths for some destinations. For example, Node E to Node F is now different for high-path-id since the path must transit the root Node A. In addition, the Node E multicast path to C is E-A-C even though E has a direct path to Node C. A rule of thumb is that the node chosen to be root should be a well-connected node and have available resources. In this example, Node A and Node D are the best choices for root nodes.

The distribution of I-SIDs allows efficient pruning of the multicast single tree on a per I-SID basis since only MFIB entries between nodes on the single tree are populated. For example, if Nodes A, B and F share an I-SID and they use the low-path–id algorithm only those three nodes would have multicast traffic for that I-SID. If the high-path-id algorithm is used traffic from Nodes A and B must go through D to get to Node F.

4.2.6.5 Data Path and Forwarding

The implementation of SPB on SR OS uses the PBB data plane. There is no flooding of BMAC based traffic. If a BMAC is not found in the FDB, traffic is dropped until the control plane populates that BMAC. Unicast BMAC addresses are populated in all FDBs regardless of I-SID membership. There is a unicast FDB per B-VPLS both control B-VPLS and user BVPLS. B-VPLS instances that do not have any I-VPLS, have only a default multicast tree and do not have any multicast MFIB entries.

The data plane supports an ingress check (reverse path forwarding check) for unicast and multicast frames on the respective trees. Ingress check is performed automatically. For unicast or multicast frames the BMAC of the source must be in the FDB and the interface must be valid for that BMAC or traffic is dropped. The PBB encapsulation (See PBB Technology) is unchanged from current SR OS. Multicast frames use the PBB Multicast Frame format and SPBM distributes I-VPLS I-SIDs which allows SPB to populate forwarding only to the relevant branches of the multicast tree. Therefore, SPB replaces both spanning tree control and MMRP functionality in one protocol.

By using a single tree for multicast the amount of MFIB space used for multicast is reduced. (Per source shortest path trees for multicast are not currently offered on SR OS.) In addition, a single tree reduces the amount of computation required when there is topology change.

4.2.6.6 SPB Ethernet OAM

Ethernet OAM works on Ethernet services and use a combination of unicast with learning and multicast addresses (REF to OAM section). SPB on SR OS supports both unicast and multicast forwarding, but with no learning and unicast and multicast may take different paths. In addition, SR OS SPB control plane offers a wide variety of show commands. The SPB IS-IS control plane takes the place of many Ethernet OAM functions. SPB IS-IS frames (Hello and PDU etc) are multicast but they are per SPB interface on the control B-VPLS interfaces and are not PBB encapsulated.

All Client Ethernet OAM is supported from I-VPLS interfaces and PBB Epipe interfaces across the SPB domain. Client OAM is the only true test of the PBB data plane. The only forms of Eth-OAM supported directly on SPB B-VPLS are Virtual MEPS (vMEPs). Only CCM is supported on these vMEPs; vMEPs use a S-TAG encapsulation and follow the SPB multicast tree for the specified B-VPLS. Each MEP has a unicast associated MAC to terminate various ETH-CFM tools. However, CCM

messages always use a destination Layer 2 multicast using 01:80:C2:00:00:3x (where x = 0 to 7). vMEPs terminate CCM with the multicast address. Unicast CCM can be configured for point to point associations or hub and spoke configuration but this would not be typical (when unicast addresses are configured on vMEPs they are automatically distributed by SPB in IS-IS).

Up MEPs on services (I-VPLS and PBB Epipes) are also supported and these behave as any service OAM. These OAM use the PBB encapsulation and follow the PBB path to the destination.

Link OAM or 802.1ah EFM is supported below SPB as standard. This strategy of SPB IS-IS and OAM gives coverage.

Table 78 **SPB Ethernet OAM Operation Summary**

OAM Origination	Data Plane Support	Comments
PBB-Epipe or Customer CFM on PBB Epipe. Up MEPs on PBB Epipe.	Fully Supported. Unicast PBB frames encapsulating unicast/multicast.	Transparent operation. Uses Encapsulated PBB with Unicast B-MAC address.
I-VPLS or Customer CFM on I-VPLS. Up MEPs on I-VPLS.	Fully Supported. Unicast/Multicast PBB frames determined by OAM type.	Transparent operation. Uses Encapsulated PBB frames with Multicast/Unicast BMAC address.
vMEP on B-VPLS Service.	CCM only. S-Tagged Multicast Frames.	Ethernet CCM only. Follows the Multicast tree. Unicast addresses may be configured for peer operation.

In summary SPB offers an automated control plane and optional Eth-CFM/Eth-EFM to allow monitoring of Ethernet Services using SPB. B-VPLS services PBB Epipes and I-VPLS services support the existing set of Ethernet capabilities.

4.2.6.7 SPB Levels

Levels are part of IS-IS. SPB supports Level 1 within a control B-VPLS. Future enhancements may make use of levels.

4.2.7 SPBM to Non-SPBM Interworking

By using static definitions of B-MACs and ISIDs interworking of PBB Epipes and I-VPLS between SPBM networks and non-SPBM PBB networks can be achieved.

4.2.7.1 Static MACs and Static ISIDs

To extend SPBM networks to other PBB networks, static MACs and ISIDs can be defined under SPBM SAPs/SDPs. The declaration of a static MAC in an SPBM context allows a non-SPBM PBB system to receive frames from an SPBM system. These static MACs are conditional on the SAP/SDP operational state. (Currently this is only supported for SPBM since SPBM can advertise these BMACs and ISIDs without any requirement for flushing.) The BMAC (and BMAC to ISID) must remain consistent when advertised in the IS-IS database.

The declaration of static-isids allows an efficient connection of ISID based services. The ISID is advertised as supported on the local nodal BMAC and the static BMACs which are the true destinations for the ISIDs are also advertised. When the I-VPLS learn the remote BMAC they will associated the ISID with the true destination BMAC. Therefore if redundancy is used the BMACs and ISIDs that are advertised must be the same on any redundant interfaces.

If the interface is an MC-LAG interface the static MAC and ISIDs on the SAPs/SDPs using that interface are only active when the associated MC-LAG interface is active. If the interface is a spoke-SDP on an active/ standby pseudo wire (PW) the ISIDs and BMACs are only active when the PW is active.

4.2.7.2 Epipe Static Configuration

For Epipe only, the BMACs need to be advertised. There is no multicast for PBB epipes. Unicast traffic will follow the unicast path shortest path or single tree. By configuring remote BMACs Epipes can be setup to non-SPBM systems. A special conditional static-mac is used for SPBM PBB B-VPLS SAPs/SDPs that are connected to a remote system. In the diagram ISID 500 is used for the PBB Epipe but only conditional MACs A and B are configured on the MC-LAG ports. The B-VPLS will advertise the static MAC either always or optionally based on a condition of the port forwarding.

Figure 115 **Static MACs Example**

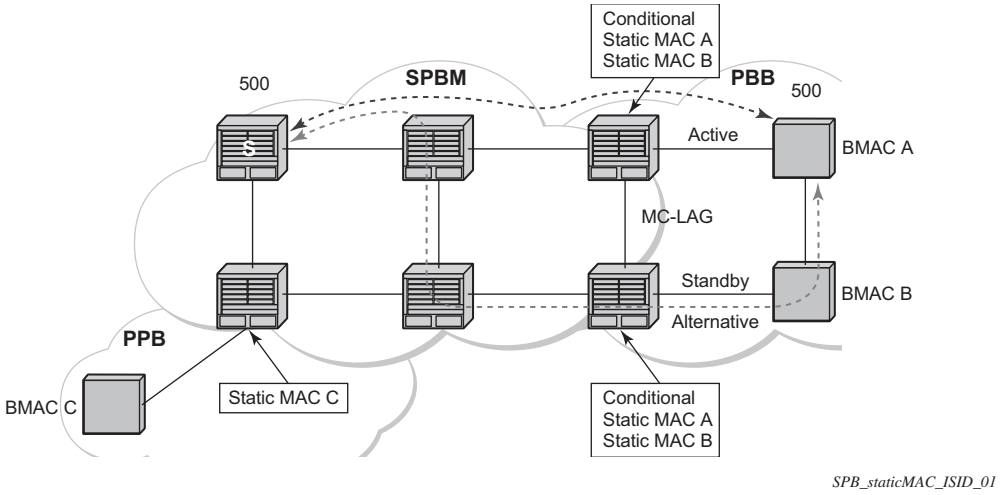

4.2.7.2.1 I-VPLS Static Config

I-VPLS static config consists of two components: static-mac and static ISIDs that represent a remote BMAC-ISID combination.

The static-MACs are configured as with Epipe, the special conditional static-mac is used for SPBM PBB B-VPLS SAPs/SDPs that are connected to a remote system. The B-VPLS will advertise the static MAC either always or optionally based on a condition of the port forwarding.

The static-isids are created under the B-VPLS SAP/SDPs that are connected to a non-SPBM system. These ISIDs are typically advertised but may be controlled by ISID policy.

For I-VPLS ISIDs the ISIDs are advertised and multicast MAC are automatically created using PBB-OUI and the ISID. SPBM supports the pruned multicast single tree. Unicast traffic will follow the unicast path shortest path or single tree. Multicast/ and unknown Unicast follow the pruned single tree for that ISID.

Figure 116 **Static ISIDs Example**

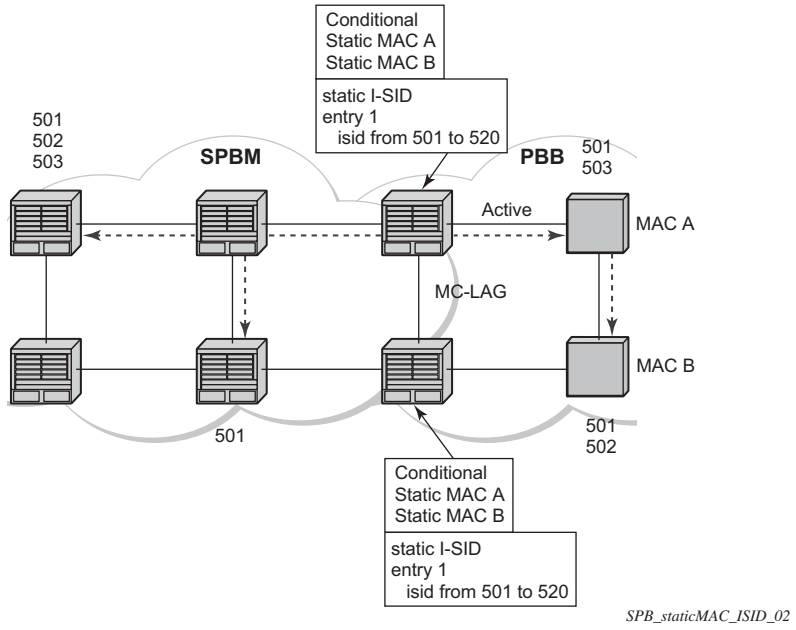

4.2.7.3 SPBM ISID Policies

ISID policies are an optional aspect of SPBM which allow additional control of ISIDs for I-VPLS. PBB services using SPBM automatically populate multicast for I-VPLS and static-isids. Improper use of isid-policy can create black holes or additional flooding of multicast.

To enable more flexible multicast, ISID policies control the amount of MFIB space used by ISIDs by trading off the default Multicast tree and the per ISID multicast tree. Occasionally customers want services that use I-VPLS that have multiple sites but use primarily unicast. The ISID policy can be used on any node where an I-VPLS is defined or static ISIDs are defined.

The typical use is to suppress the installation of the ISID in the MFIB using use-def-mcast and the distribution of the ISID in SPBM by using no advertise-local.

The use-def-mcast policy instructs SPBM to use the default B-VPLS multicast forwarding for the ISID range. The ISID multicast frame remains unchanged by the policy (the standard format with the PBB OUI and the ISID as the multicast destination address) but no MFIB entry is allocated. This causes the forwarding to use the default BVID multicast tree which is not pruned. When this policy is in place it only governs the forwarding locally on the current B-VPLS.

The advertise local policy ISID policies are applied to both static ISIDs and I-VPLS ISIDs. The policies define whether the ISIDs are advertised in SPBM and whether the use the local MFIB. When ISIDs are advertised they will use the MFIB in the remote nodes. Locally the use of the MFIB is controlled by the **use-def-mcast** policy.

The types of interfaces are summarized in Table 79.

Table 79 **SPBM ISID Policies Table**

Service Type	ISID Policy on B-VPLS	Notes
Epipe	No effect	PBB Epipe ISIDs are not advertised or in MFIB.
I-VPLS	None: Uses ISID Multicast tree. Advertised ISIDs of I-VPLS.	I-VPLS uses dedicated (pruned) multicast tree. ISIDs are advertised.
I-VPLS (for Unicast)	use-def-mcast no advertise-local	I-VPLS uses default Multicast. Policy only required where ISIDs are defined. ISIDs not advertised. must be consistently defined on all nodes with same ISIDs.
I-VPLS (for Unicast)	use-def-mcast advertise-local	I-VPLS uses default Multicast. Policy only required where ISIDs are defined. ISIDs advertised and pruned tree used elsewhere. May be inconsistent for an ISID.
Static ISIDs for I-VPLS interworking	None: (recommended) Uses ISID Multicast tree	I-VPLS uses dedicated (pruned) multicast tree. ISIDs are advertised.
Static ISIDs for I-VPLS interworking (defined locally)	use-def-mcast	I-VPLS uses default Multicast. Policy only required where ISIDs are configured or where I-VPLS is located.
No MFIB for any ISIDs Policy defined on all nodes	use-def-mcast no advertise-local	Each B-VPLS with the policy will not install MFIB. Policy defined on all switches ISIDs are defined. ISIDs advertised and pruned tree used elsewhere. May be inconsistent for an ISID.

4.2.8 ISID Policy Control

4.2.8.1 Static ISID Advertisement

Static ISIDs are advertised between using the SPBM Service Identifier and Unicast Address sub-TLV in IS-IS when there is no ISID policy. This TLV advertises the local B-MAC and one or more ISIDs. The BMAC used is the source-bmac of the Control/User VPLS. Typically remote BMACs (the ultimate source-bmac) and the associated ISIDs are configured as static under the SPBM interface. This allows all remote BMACs and all remote ISIDs can be configured once per interface.

4.2.8.2 I-VPLS for Unicast Service

If the service is using unicast only an I-VPLS still uses MFIB space and SPBM advertises the ISID. By using the default multicast tree locally, a node saves MFIB space. By using the no advertise-local SPBM will not advertise the ISIDs covered by the policy. Note the actual PBB multicast frames are the same regardless of policy. Unicast traffic is the not changed for the ISID policies.

The Static BMAC configuration is allowed under Multi-Chassis LAG (MC-LAG) based SAPs and active/standby PW SDPs.

Unicast traffic will follow the unicast path shortest path or single tree. By using the ISID policy Multicast/and unknown Unicast traffic (BUM) follows the default B-VPLS tree in the SPBM domain. This should be used sparingly for any high volume of multicast services.

Figure 117 **ISID Policy Example**

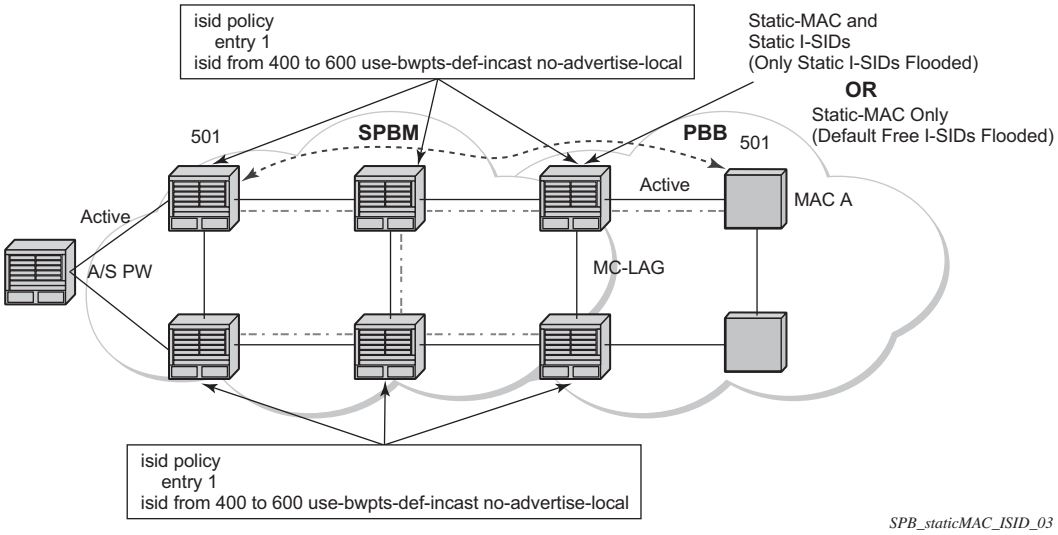

4.2.9 Default Behaviors

When static ISIDs are defined the default is to advertise the static ISIDs when the interface parent (SAP or SDP) is up.

If the advertisement is not needed, an ISID policy can be created to prevent advertising the ISID.

- use-def-mcast: If a policy is defined with use-def-mcast the local MFIB will not contain an Multicast MAC based on the PBB OUI+ ISID and the frame will be flooded out the local tree. This applies to any node where the policy is defined. On other nodes if the ISID is advertised the ISID will use the MFIB for that ISID.
- No advertise-local: If a policy of no advertise-local is defined the ISIDs in the policy will not be advertised. This combination should be used everywhere there is an I-VPLS with the ISID or where the Static ISID is defined to prevent black holes. If an ISID is to be moved from advertising to no advertising it is advisable to use **use-def-mcast** on all the nodes for that ISID which will allow the MFIB to not be installed and will start using the default multicast tree at each node with that policy. Then the no advertise-local option can be used.

Each Policy may be used alone or in combination.

4.2.10 Example Network Configuration

Figure 118 **Sample Network**

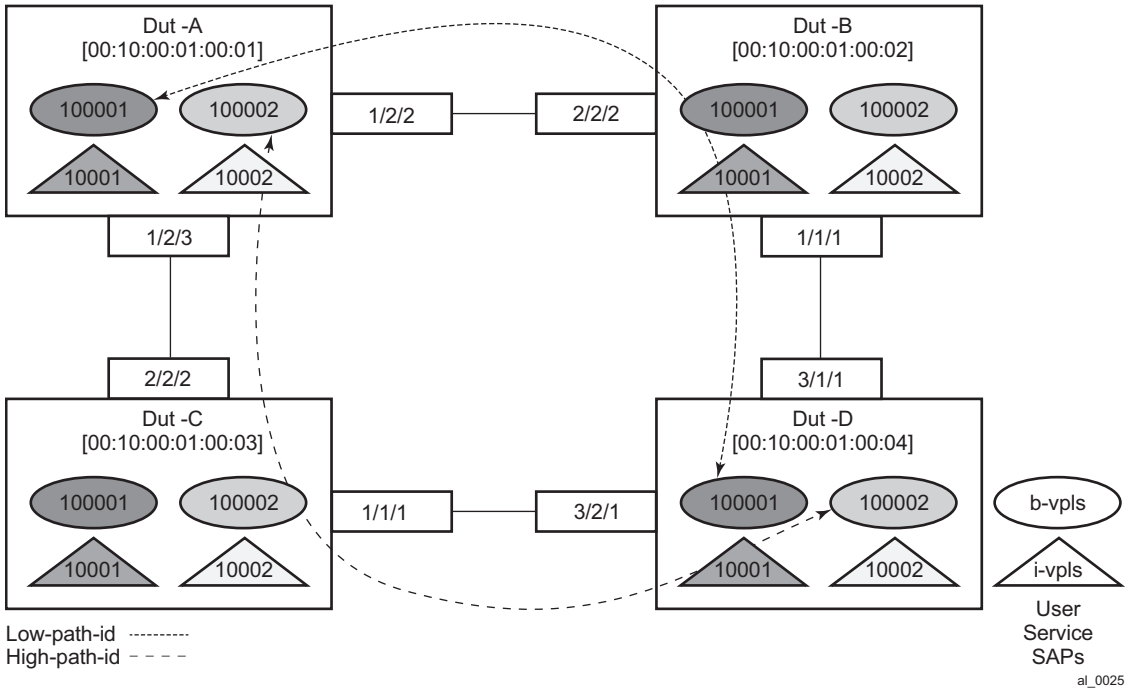

Figure 118 shows an example network showing four nodes with SPB B-VPLS. The SPB instance is configured on the B-VPLS 100001. B-VPLS 100001 uses FID 1 for SPB instance 1024. All BMACs and I-SIDs are learned in the context of B-VPLS 100001. B-VPLS 100001 has an i-vpls 10001 service, which also uses the I-SID 10001. B-VPLS 100001 is configured to use VID 1 on SAPs 1/2/2 and 1/2/3 and while the VID does not need to be the same as the FID the VID does however need to be the same on the other side (Dut-B and Dut-C).

A user B-VPLS service 100002 is configured and it uses B-VPLS 10001 to provide forwarding. It fate shares the control topology. In Figure 118, the control B-VPLS uses the low-path-id algorithm and the user B-VPLS uses high-path-id algorithm. Any B-VPLS can use any algorithm. The difference is illustrated in the path between Dut A and Dut D. The short dashed line through Dut-B is the low-path-id algorithm and the long dashed line thought Dut C is the high-path-id algorithm.

4.2.10.1 Sample Configuration for Dut-A

```
Dut-A:
Control B-VPLS:*A:Dut-A>config>service>vpls# pwc
```

```
-------------------------------------------------------------------------------
Present Working Context :
-------------------------------------------------------------------------------
 <root>
  configure
  service
  vpls "100001"
-------------------------------------------------------------------------------
*A:Dut-A>config>service>vpls# info
----------------------------------------------
            pbb
                source-bmac 00:10:00:01:00:01
            exit
            stp
                shutdown
            exit
            spb 1024 fid 1 create
                level 1
                    ect-algorithm fid-range 100-100 high-path-id
                exit
                no shutdown
            exit
            sap 1/2/2:1.1 create
                spb create
                    no shutdown
                exit
            exit
            sap 1/2/3:1.1 create
                spb create
                    no shutdown
                exit
            exit
            no shutdown
----------------------------------------------
User B-VPLS:
*A:Dut-A>config>service>vpls# pwc
-------------------------------------------------------------------------------
Present Working Context :
-------------------------------------------------------------------------------
 <root>
  configure
  service
  vpls "100002"
-------------------------------------------------------------------------------
*A:Dut-A>config>service>vpls# info
----------------------------------------------
            pbb
                source-bmac 00:10:00:02:00:01
            exit
            stp
                shutdown
            exit
            spbm-control-vpls 100001 fid 100
            sap 1/2/2:1.2 create
            exit
            sap 1/2/3:1.2 create
            exit
            no shutdown
----------------------------------------------
```

```
I-VPLS:
configure service
        vpls 10001 customer 1 i-vpls create
            service-mtu 1492
            pbb
                backbone-vpls 100001
                exit
            exit
            stp
                shutdown
            exit
            sap 1/2/1:1000.1 create
            exit
            no shutdown
        exit
        vpls 10002 customer 1 i-vpls create
            service-mtu 1492
            pbb
                backbone-vpls 100002
                exit
            exit
            stp
                shutdown
            exit
            sap 1/2/1:1000.2 create
            exit
            no shutdown
        exit
exit
```

4.2.10.1.1 Show Commands Outputs

The **show base** commands output a summary of the instance parameters under a control B-VPLS. The **show** command for a user B-VPLS indicates the control B-VPLS. The base parameters except for Bridge Priority and Bridge ID must match on neighbor nodes.

```
*A:Dut-A# show service id 100001 spb base
===============================================================================
Service SPB Information
===============================================================================
Admin State        : Up                 Oper State         : Up
ISIS Instance      : 1024               FID                : 1
Bridge Priority    : 8                  Fwd Tree Top Ucast : spf
Fwd Tree Top Mcast : st
Bridge Id          : 80:00.00:10:00:01:00:01
Mcast Desig Bridge : 80:00.00:10:00:01:00:01
===============================================================================
ISIS Interfaces
===============================================================================
Interface                      Level CircID  Oper State   L1/L2 Metric
-------------------------------------------------------------------------------
sap:1/2/2:1.1                  L1    65536   Up           10/-
sap:1/2/3:1.1                  L1    65537   Up           10/-
```

```
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
FID ranges using ECT Algorithm
-------------------------------------------------------------------------------
1-99      low-path-id
100-100   high-path-id
101-4095  low-path-id
===============================================================================
```

The **show adjacency** command displays the system ID of the connected SPB B-VPLS neighbors and the associated interfaces to connect those neighbors.

```
*A:Dut-A# show service id 100001 spb adjacency
===============================================================================
ISIS Adjacency
===============================================================================
System ID                Usage State Hold Interface                    MT Enab
-------------------------------------------------------------------------------
Dut-B                    L1    Up    19   sap:1/2/2:1.1                No
Dut-C                    L1    Up    21   sap:1/2/3:1.1                No
-------------------------------------------------------------------------------
Adjacencies : 2
===============================================================================
```

Details about the topology can be displayed with the **database** command. There is a detail option that displays the contents of the LSPs.

```
*A:Dut-A# show service id 100001 spb database
===============================================================================
ISIS Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------
Displaying Level 1 database
-------------------------------------------------------------------------------
Dut-A.00-00                           0xc       0xbaba   1103     L1
Dut-B.00-00                           0x13      0xe780   1117     L1
Dut-C.00-00                           0x13      0x85a    1117     L1
Dut-D.00-00                           0xe       0x174a   1119     L1
Level (1) LSP Count : 4
===============================================================================
```

The **show routes** command illustrates the next hop if for the MAC addresses both unicast and multicast. The path to 00:10:00:01:00:04 (Dut-D) illustrates the low-path-id algorithm id. For FID one the neighbor is Dut-B and for FID 100 the neighbor is Dut-C. Since Dut-A is the root of the multicast single tree the multicast forwarding is the same for Dut-A. However, unicast and multicast routes will differ on most other nodes. Also the I-SIDs exist on all of the nodes so I-SID base multicast follows the multicast tree exactly. If the I-SID had not existed on Dut-B or Dut-D then for FID 1 there would be no entry. Note only designated nodes (root nodes) show metrics. Non-designated nodes will not show metrics.

```
*A:Dut-A# show service id 100001 spb routes
===============================================================
MAC Route Table
===============================================================
Fid  MAC                                          Ver.   Metric
      NextHop If                    SysID
---------------------------------------------------------------
Fwd Tree: unicast
---------------------------------------------------------------
1    00:10:00:01:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
1    00:10:00:01:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
1    00:10:00:01:00:04                           10     20
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
100  00:10:00:02:00:04                           10     20
      sap:1/2/3:1.1                Dut-C

Fwd Tree: multicast
---------------------------------------------------------------
1    00:10:00:01:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
1    00:10:00:01:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
1    00:10:00:01:00:04                           10     20
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
100  00:10:00:02:00:04                           10     20
      sap:1/2/3:1.1                Dut-C
---------------------------------------------------------------
No. of MAC Routes: 12
===============================================================

===============================================================
ISID Route Table
===============================================================
Fid  ISID                                         Ver.
      NextHop If                    SysID
---------------------------------------------------------------
1    10001                                       10
      sap:1/2/2:1.1                Dut-B
      sap:1/2/3:1.1                Dut-C
100  10002                                       10
      sap:1/2/2:1.1                Dut-B
      sap:1/2/3:1.1                Dut-C
---------------------------------------------------------------
No. of ISID Routes: 2
===============================================================
```

The **show service spb fdb** command shows the programmed unicast and multicast source MACs in SPB-managed B-VPLS service.

```
*A:Dut-A# show service id 100001 spb fdb

===============================================================================
User service FDB information
===============================================================================
MacAddr           UCast Source          State   MCast Source          State
-------------------------------------------------------------------------------
00:10:00:01:00:02 1/2/2:1.1             ok      1/2/2:1.1             ok
00:10:00:01:00:03 1/2/3:1.1             ok      1/2/3:1.1             ok
00:10:00:01:00:04 1/2/2:1.1             ok      1/2/2:1.1             ok
-------------------------------------------------------------------------------
Entries found: 3
===============================================================================

*A:Dut-A# show service id 100002 spb fdb
===============================================================================
User service FDB information
===============================================================================
MacAddr           UCast Source          State   MCast Source          State
-------------------------------------------------------------------------------
00:10:00:02:00:02 1/2/2:1.2             ok      1/2/2:1.2             ok
00:10:00:02:00:03 1/2/3:1.2             ok      1/2/3:1.2             ok
00:10:00:02:00:04 1/2/3:1.2             ok      1/2/3:1.2             ok
-------------------------------------------------------------------------------
Entries found: 3
===============================================================================
```

The **show service spb mfib** command shows the programmed multicast ISID addresses Macs in SPB-managed B-VPLS service shows the multicast ISID pbb group mac addresses in SPB-managed B-VPLS. Other types of *,G multicast traffic is sent over the multicast tree and these MACs are not shown. OAM traffic that uses multicast (for example vMEP CCM) will take this path for example.

```
*A:Dut-A# show service id 100001 spb mfib
===============================================================================
User service MFIB information
===============================================================================
MacAddr           ISID     Status
-------------------------------------------------------------------------------
01:1E:83:00:27:11 10001    Ok
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
*A:Dut-A# show service id 100002 spb mfib
===============================================================================
User service MFIB information
===============================================================================
MacAddr           ISID     Status
-------------------------------------------------------------------------------
01:1E:83:00:27:12 10002    Ok
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
```

4.2.10.1.2 Debug Commands

- debug service id <svcId> spb
- debug service id <svcId> spb adjacency
- 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

4.2.10.1.3 Tools Commands

- tools perform service id <svcId> spb run-manual-spf
- tools dump service id spb
- tools dump service id spb default-multicast-list
- tools dump service id spb forwardingpath

4.2.10.1.4 Clear Commands

- clear service id <svcId> spb
- clear service id <svcId> spb adjacency
- clear service id <svcId> spb database
- clear service id <svcId> spb spf-log
- clear service id <svcId> spb statistics

4.2.11 IEEE 802.1ak MMRP for Service Aggregation and Zero Touch Provisioning

IEEE 802.1ah supports an M:1 model where multiple customer services, represented by ISIDs, are transported through a common infrastructure (B-component). The Nokia PBB implementation supports the M:1 model allowing for a service architecture where multiple customer services (I-VPLS or Epipe) can be transported through a common B-VPLS infrastructure as depicted in Figure 119.

Figure 119 **Customer Services Transported in 1 B-VPLS (M:1 Model)**

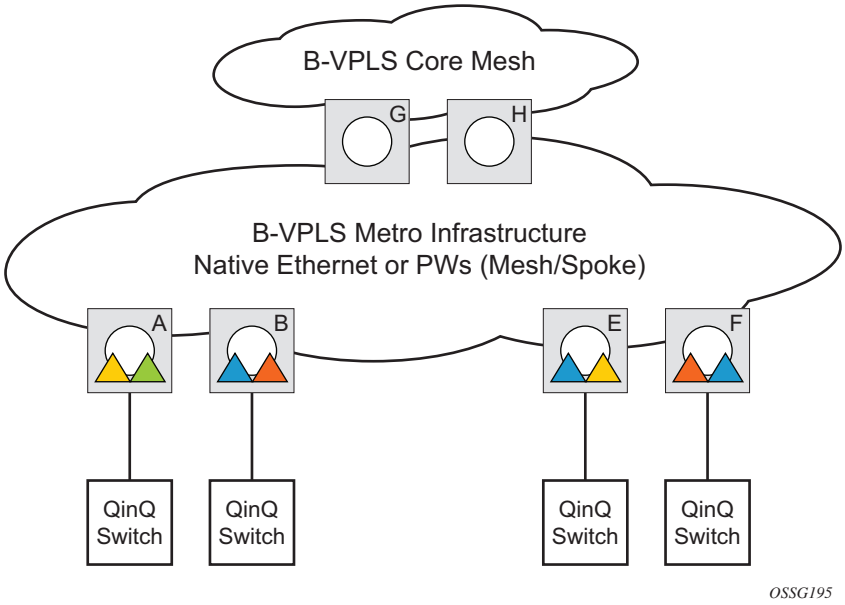

The B-VPLS infrastructure represented by the white circles is used to transport multiple customer services represented by the triangles of different colors. This service architecture minimizes the number of provisioning touches and reduces the load in the core PEs: for example, G and H use less VPLS instances and pseudowire.

In a real life deployment, different customer VPNs do not share the same community of interest – for example, VPN instances may be located on different PBB PEs. The M:1 model depicted in Figure 120 requires a per VPN flood containment mechanism so that VPN traffic is distributed just to the B-VPLS locations that have customer VPN sites: for example, flooded traffic originated in the blue I-VPLS should be distributed just to the PBB PEs where blue I-VPLS instances are present – PBB PE B, E and F.

Per customer VPN distribution trees need to be created dynamically throughout the BVPLS as new customer I-VPLS instances are added in the PBB PEs.

The Nokia PBB implementation employs the IEEE 802.1ak Multiple MAC Registration Protocol (MMRP) to dynamically build per I-VPLS distribution trees inside a certain B-VPLS infrastructure.

IEEE 802.1ak Multiple Registration Protocol (MRP) – Specifies changes to IEEE Std 802.1Q that provide a replacement for the GARP, GMRP and GVRP protocols. MMRP application of IEEE 802.1ak specifies the procedures that allow the registration/de-registration of MAC addresses over an Ethernet switched infrastructure.

In the PBB case, as I-VPLS instances are enabled in a certain PE, a group BMAC address is by default instantiated using the standard based PBB Group OUI and the ISID value associated with the I-VPLS.

When a new I-VPLS instance is configured in a PE, the IEEE 802.1ak MMRP application is automatically invoked to advertise the presence of the related group B-MAC on all active B-VPLS SAPs and SDP bindings.

When at least two I-VPLS instances with the same ISID value are present in a B-VPLS, an optimal distribution tree is built by MMRP in the related B-VPLS infrastructure as depicted in Figure 120.

Figure 120 **Flood Containment Requirement in M:1 Model**

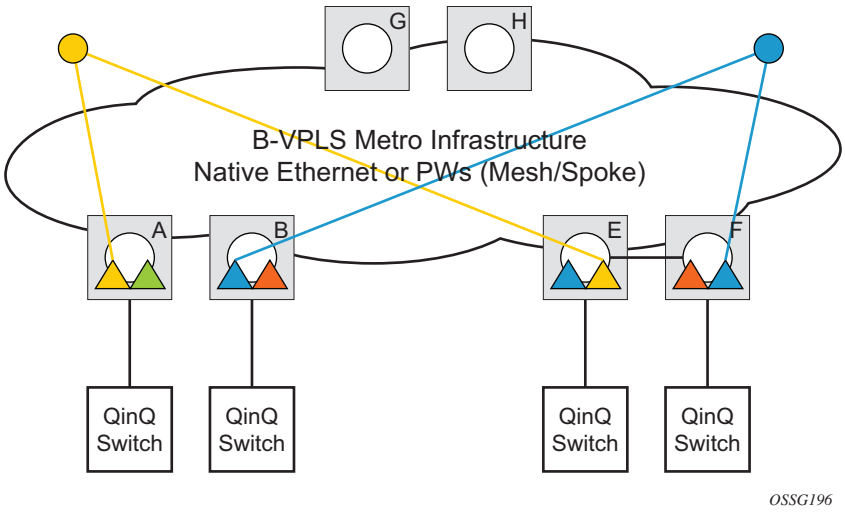

4.2.12 MMRP Support Over B-VPLS SAPs and SDPs

MMRP is supported in B-VPLS instances over all the supported BVPLS SAPs and SDPs, including the primary and standby pseudowire scheme implemented for VPLS resiliency.

When a B-VPLS with MMRP enabled receives a packet destined for a specific group BMAC, it checks its own MFIB entries and if the group BMAC does not exist, it floods it everywhere. This should never happen as this kind of packet will be generated at the I-VPLS/PBB PE when a registration was received for a local I-VPLS group BMAC.

4.2.12.1 I-VPLS Changes and Related MMRP Behavior

This section describes the MMRP behavior for different changes in IVPLS.

1. When an ISID is set for a certain I-VPLS and a link to a related B-VPLS is activated (for example, through the **config>service>vpls>backbone-vpls** *vpls id:isid* command), the group BMAC address is declared on all B-VPLS virtual ports (SAPs or SDPs).
2. When the ISID is changed from one value to a new one, the old group BMAC address is undeclared on all ports and the new group BMAC address is declared on all ports in the B-VPLS.
3. When the I-VPLS is disassociated with the B-VPLS, the old group BMAC is no longer advertised as a local attribute in the B-VPLS if no other peer B-VPLS PEs have it declared.
4. When an I-VPLS goes operationally down (either all SAPs/SDPs are down) or the I-VPLS is shutdown, the associated group BMAC is undeclared on all ports in the B-VPLS.
5. When the I-VPLS is deleted, the group BMAC should already be un-declared on all ports in the B-VPLS because the I-VPLS has to be shutdown in order to delete it.

4.2.12.2 Limiting the Number of MMRP Entries on a Per B-VPLS Basis

The MMRP exchanges create one entry per attribute (group BMAC) in the B-VPLS where MMRP protocol is running. When the first registration is received for an attribute, an MFIB entry is created for it.

The Nokia implementation allows the user to control the number of MMRP attributes (group BMACs) created on a per B-VPLS basis. Control over the number of related MFIB entries in the B-VPLS FDB is inherited from previous releases through the use of the **config>service>vpls>mfib-table-size** *table-size* command. This ensures that no B-VPLS will take up all the resources from the total pool.

4.2.12.3 Optimization for Improved Convergence Time

Assuming that MMRP is used in a certain B-VPLS, under failure conditions the time it takes for the B-VPLS forwarding to resume may depend on the data plane and control plane convergence plus the time it takes for MMRP exchanges to settle down the flooding trees on a per ISID basis.

To minimize the convergence time, the Nokia PBB implementation offers the selection of a mode where B-VPLS forwarding reverts for a short time to flooding so that MMRP has enough time to converge. This mode can be selected through configuration using **config>service>vpls> b-vpls>mrp>flood-time** *value* command where *value* represents the amount of time in seconds that flooding will be enabled. Refer to the PBB Configuration Command Reference for command syntax and usage.

If this behavior is selected, the forwarding plane reverts to B-VPLS flooding for a configurable time period, for example, for a few seconds, then it reverts back to the MFIB entries installed by MMRP.

The following B-VPLS events initiate the switch from per I-VPLS (MMRP) MFIB entries to “B-VPLS flooding”:

- Reception or local triggering of a TCN
- B-SAP failure
- Failure of a B-SDP binding
- Pseudowire activation in a primary/standby HVPLS resiliency solution
- SF/CPM switchover due to STP reconvergence

4.2.12.4 Controlling MRP Scope using MRP Policies

MMRP advertises the Group BMACs associated with ISIDs throughout the whole BVPLS context regardless of whether a specific IVPLS is present in one or all the related PEs or BEBs. When evaluating the overall scalability the resource consumption in both the control and data plane must be considered:

- Control plane - MMRP processing and number of attributes advertised
- Data plane – one tree is instantiated per ISID or Group BMAC attribute

In a multi-domain environment, for example multiple MANs interconnected through a WAN, the BVPLS and implicitly MMRP advertisement may span across domains. The MMRP attributes will be flooded throughout the BVPLS context indiscriminately, regardless of the distribution of IVPLS sites.

The solution described in this section limits the scope of MMRP control plane advertisements to a specific network domain using MRP Policy. ISID-based filters are also provided as a safety measure for BVPLS data plane.

Figure 121 **Inter-Domain Topology**

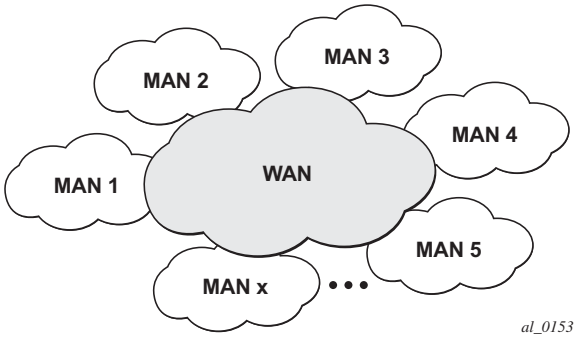

Figure 121 shows the case of an Inter-domain deployment where multiple metro domains (MANs) are interconnected through a wide area network (WAN). A BVPLS is configured across these domains running PBB M:1 model to provide infrastructure for multiple IVPLS services. MMRP is enabled in the BVPLS to build per IVPLS flooding trees. To limit the load in the core PEs or PBB BCBs, the local IVPLS instances must use MMRP and data plane resources only in the MAN regions where they have sites. A solution to the above requirements is depicted in Figure 122. The case of native PBB metro domains inter-connected via a MPLS core is used in this example. Other technology combinations are possible.

Figure 122 **Limiting the Scope of MMRP Advertisements**

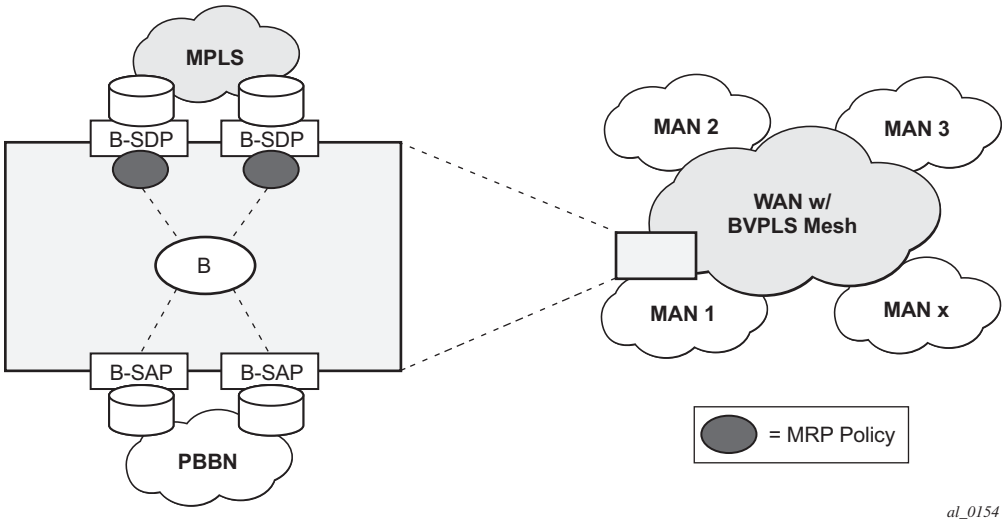

An MRP policy can be applied to the edge of MAN1 domain to restrict the MMRP advertisements for local ISIDs outside local domain. Or the MRP policy can specify the inter-domain ISIDs allowed to be advertised outside MAN1. The configuration of MRP policy is similar with the configuration of a filter. It can be specified as a template or exclusively for a specific endpoint under service mrp object. An ISID or a range of ISID(s) can be used to specify one or multiple match criteria that will be used to generate the list of Group MACs to be used as filters to control which MMRP attributes can be advertised. An example of a simple mrp-policy that allows the advertisement of Group BMACs associated with ISID range 100-150 is given below:

```
*A:ALA-7>config>service>mrp# info
----------------------------------------------
            mrp-policy "test" create
                default-action block
                entry 1 create
                    match
                        isid 100 to 150
                    exit
        action allow
                exit
            exit
----------------------------------------------
```

A special action end-station is available under mrp-policy entry object to allow the emulation on a specific SAP/PW of an MMRP end-station. This is usually required when the operator does not want to activate MRP in the WAN domain for interoperability reasons or if it prefers to manually specify which ISID will be interconnected over the WAN. In this case the MRP transmission will be shutdown on that SAP/PW and the configured ISIDs will be used the same way as an IVPLS connection into the BVPLS, emulating a static entry in the related BVPLS MFIB. Also if MRP is active in the BVPLS context, MMRP will declare the related GBMAC(s) continuously over all the other BVPLS SAP/PW(s) until the mrp-policy end-station action is removed from the mrp-policy assigned to that BVPLS context.

The MMRP usage of the mrp-policy will ensure automatically that traffic using Group BMAC will not be flooded between domains. There could be though small transitory periods when traffic originated from PBB BEB with unicast BMAC destination may be flooded in the BVPLS context as unknown unicast in the BVPLS context for both IVPLS and PBB Epipe. To restrict distribution of this traffic for local PBB services a new ISID match criteria is added to existing mac-filters. The mac-filter configured with ISID match criteria can be applied to the same interconnect endpoint(s), BVPLS SAP or PW, as the mrp-policy to restrict the egress transmission any type of frames that contain a local ISID. An example of this new configuration option is described below:

```
----------------------------------------------
A;ALA-7>config>filter# info
----------------------------------------------
mac-filter 90 create
description "filter-wan-man"
```

```
type isid
scope template
entry 1 create
description "drop-local-isids"
match
isid from 100 to 1000
exit
action drop
exit
-----------------------------------------------
```

These filters will be applied as required on a per B-SAP or B-PW basis just in the egress direction. The ISID match criteria is exclusive with any other criteria under mac-filter. A new mac-filter type attribute is defined to control the use of ISID match criteria and must be set to isid to allow the use of isid match criteria. The ISID tag is identified using the PBB ethertype provisioned under **config>port>ethernet>pbb-etype**.

4.2.13 PBB and BGP-AD

BGP auto-discovery is supported only in the BVPLS to automatically instantiate the BVPLS pseudowires and SDPs as described in the *7450 ESS, 7750 SR, and 7950 XRS Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

4.2.14 PBB E-Line Service

E-Line service is defined in PBB (IEEE 802.1ah) as a point-to-point service over the B-component infrastructure. The Nokia implementation offers support for PBB E-Line through the mapping of multiple Epipe services to a Backbone VPLS infrastructure.

The use of Epipe scales the E-Line services as no MAC switching, learning or replication is required in order to deliver the point-to-point service.

All packets ingressing the customer SAP/spoke-SDP are PBB encapsulated and unicasted through the B-VPLS “tunnel” using the backbone destination MAC of the remote PBB PE. The Epipe service does not support the forwarding of PBB encapsulated frames received on SAPs or Spoke-SDPs through their associated B-VPLS service. PBB frames are identified based on the configured PBB Ethertype (0x88e7 by default).

All the packets ingressing the B-VPLS destined for the Epipe are PBB de-encapsulated and forwarded to the customer SAP/spoke-SDP.

A PBB E-Line service support the configuration of a SAP or non-redundant spoke-SDP.

4.2.14.1 Non-Redundant PBB Epipe Spoke Termination

This feature provides the capability to use non-redundant pseudowire connections on the access side of a PBB Epipe, where previously only SAPs could be configured.

4.2.15 PBB Using G.8031 Protected Ethernet Tunnels

IEEE 802.1ah Provider Backbone Bridging (PBB) specification employs provider MSTP (PMSTP) to ensure loop avoidance in a resilient native Ethernet core. The usage of P-MSTP means failover times depend largely on the size and the connectivity model used in the network. The use of MPLS tunnels provides a way to scale the core while offering fast failover times using MPLS FRR. There are still service provider environments where Ethernet services are deployed using native Ethernet backbones. A solution based on native Ethernet backbone is required to achieve the same fast failover times as in the MPLS FRR case.

The Nokia PBB implementation offers the capability to use core Ethernet tunnels compliant with ITU-T G.8031 specification to achieve 50 ms resiliency for backbone failures. This is required to comply with the stringent SLAs provided by service providers in the current competitive environment. The implementation also allows a LAG-emulating Ethernet tunnel providing a complimentary native Ethernet E-LAN capability. The LAG-emulating Ethernet tunnels and G.8031 protected Ethernet tunnels operate independently.

The next section describes an applicability example where an Ethernet service provider using native PBB offers a carrier of carrier backhaul service for mobile operators.

4.2.15.1 Solution Overview

A simplified topology example for a PBB network offering a carrier of carrier service for wireless service providers is depicted in Figure 123.

Figure 123 **Mobile Backhaul Use Case**

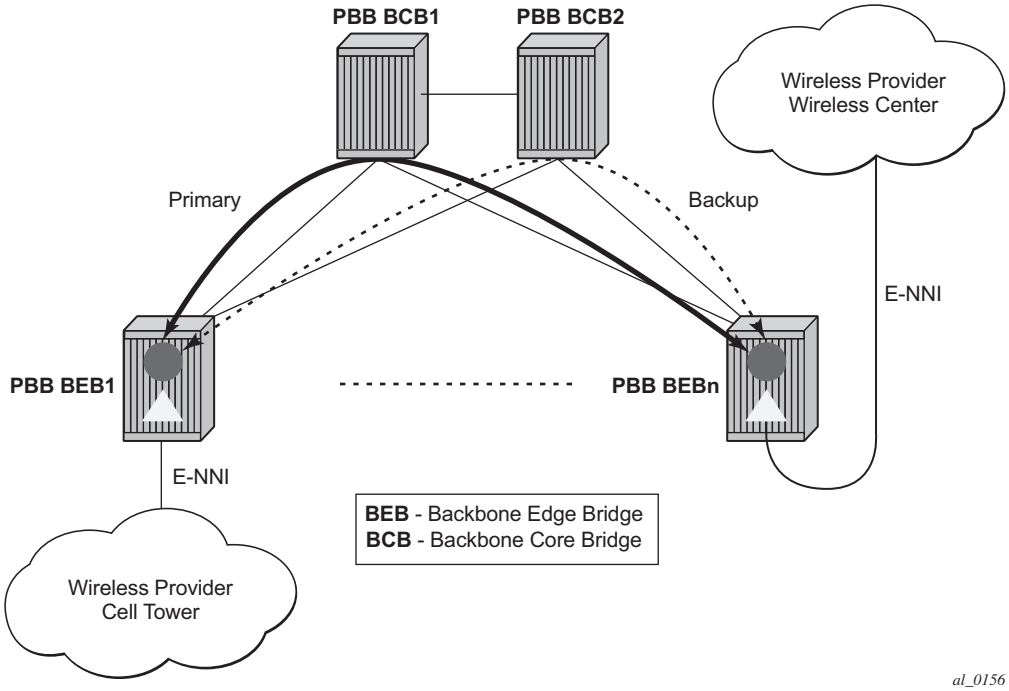

The wireless service provider in this example purchases an E-Line service between the ENNIs on PBB edge nodes, BEB1 and BEBn. PBB services are employing a type of Ethernet tunneling (Eth-tunnels) between BEBs where primary and backup member paths controlled by G.8031 1:1 protection are used to ensure faster backbone convergence. Ethernet CCMs based on IEEE 802.1ag specification may be used to monitor the liveliness for each individual member paths.

The Ethernet paths span a native Ethernet backbone where the BCBs are performing simple Ethernet switching between BEBs using an Epipe or a VPLS service.

Although the network diagram shows just the Epipe case, both PBB E-Line and E-LAN services are supported.

4.2.15.2 Detailed Solution Description

This section discusses the details of the Ethernet tunneling for PBB. The main solution components are depicted in Figure 124.

Figure 124 **PBB-Epipe with B-VPLS over Ethernet Tunnel**

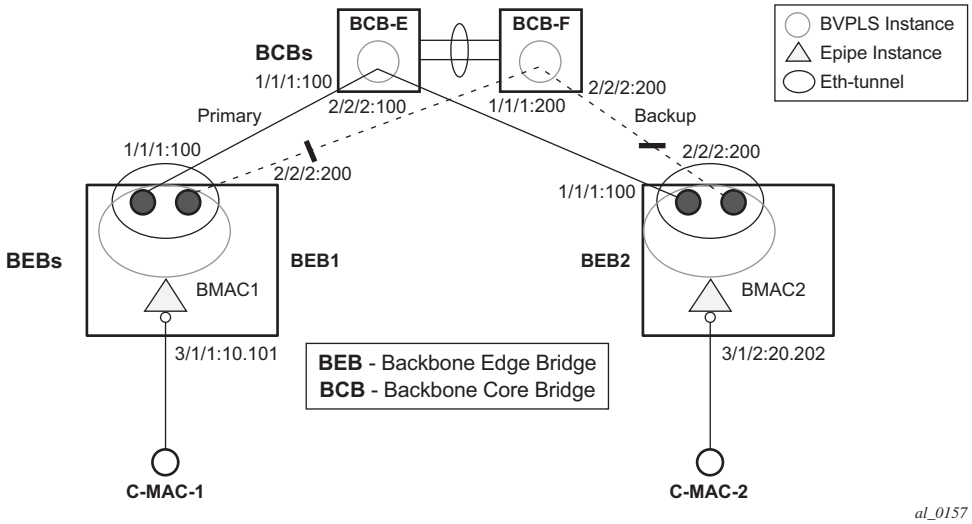

The PBB E-Line service is represented in the BEBs as a combination of an Epipe mapped to a BVPLS instance. A eth-tunnel object is used to group two possible paths defined by specifying a member port and a control tag. In our example, the blue-circle representing the eth-tunnel is associating in a protection group the two paths instantiated as (port, control-tag/bvid): a primary one of port 1/1/1, control-tag 100 and respectively a secondary one of port 2/2/2, control tag 200.

The BCBs devices will stitch each BVID between different BEB-BCB links using either a VPLS or Epipe service. Epipe instances are recommended as the preferred option due to increased tunnel scalability.

Fast failure detection on the primary and backup paths is provided using IEEE 802.1ag CCMs that can be configured to transmit at 10 msec interval. Alternatively, the link layer fault detection mechanisms like LoS/RDI or 802.3ah can be employed.

Path failover is controlled by an Ethernet protection module, based on standard G.8031 Ethernet Protection Switching. The Nokia implementation of Ethernet protection switching supports only the 1:1 model which is common practice for packet based services since it makes better use of available bandwidth. The following additional functions are provided by the protection module:

- Synchronization between BEBs such that both send and receive on the same Ethernet path in stable state.
- Revertive / non-revertive choices.
- Compliant G.8031 control plane.

The secondary path requires a MEP to exchange the G.8031 APS PDUs. The following Ethernet CFM configuration in the **eth-tunnel>path>eth-cfm>mep** context can be used to enable the G.8031 protection without activating the Ethernet CCMs:

- Create the domain (MD) in CFM.
- Create the association (MA) in CFM. NOTE: Do not put remote MEPs.
- Create the MEP.
- Configure control-mep and no shutdown on the MEP.
- The CCM transmission should stay disabled using the **no ccm-enable** command.

If a MEP is required for troubleshooting issues on the primary path, the configuration described above for the secondary path must be used to enable the use of Link Layer OAM on the primary path.

LAG loadsharing is offered to complement G.8031 protected Ethernet tunnels for situations where unprotected VLAN services are to be offered on some or all of the same native Ethernet links.

Figure 125 **G.8031 P2P Tunnels and LAG-Like Loadsharing Co-Existence**

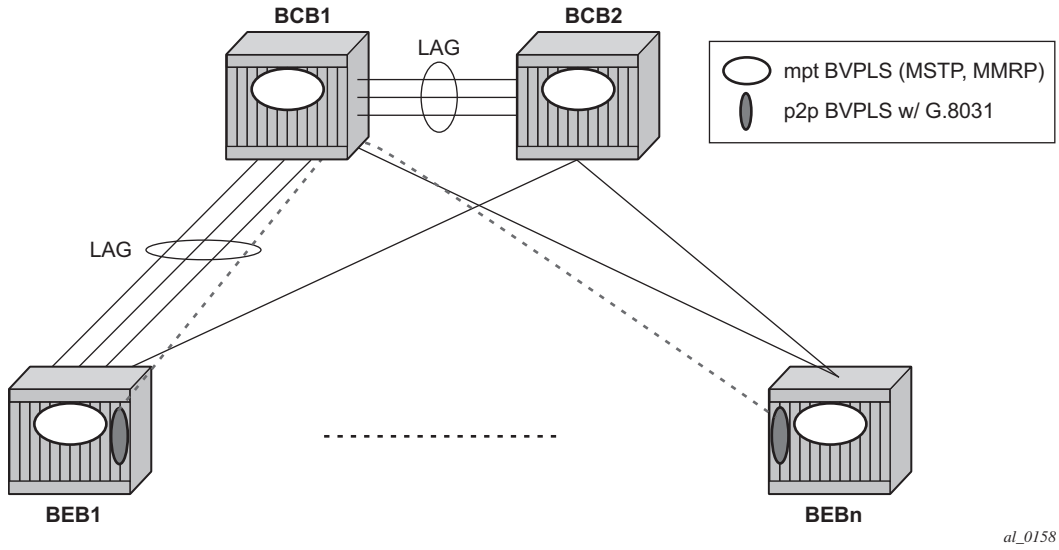

In Figure 125, the G.8031 Ethernet tunnels are used by the B-SAP(s) mapped to the green BVPLS entities supporting the E-Line services. A LAG-like loadsharing solution is provided for the Multipoint BVPLS (white circles) supporting the E-LAN (IVPLS) services. The green G.8031 tunnels co-exist with LAG-emulating Ethernet tunnels (loadsharing mode) on both BEB-BCB and BCB-BCB physical links.

The G.8031-controlled Ethernet tunnels will select an active tunnel based on G.8031 APS operation, while emulated-LAG Ethernet tunnels will hash traffic within the configured links. Upon failure of one of the links the emulated-LAG tunnels will rehash traffic within the remaining links and fail the tunnel when the number of links breaches the minimum required (independent of G.8031-controlled Ethernet tunnels on the links shared emulated-LAG).

4.2.15.3 Detailed PBB Emulated LAG Solution Description

This section discusses the details of the emulated LAG Ethernet tunnels for PBB. The main solution components are depicted in Figure 126 which overlays Ethernet Tunnels services on the network from Figure 124.

Figure 126 **Ethernet Tunnel Overlay**

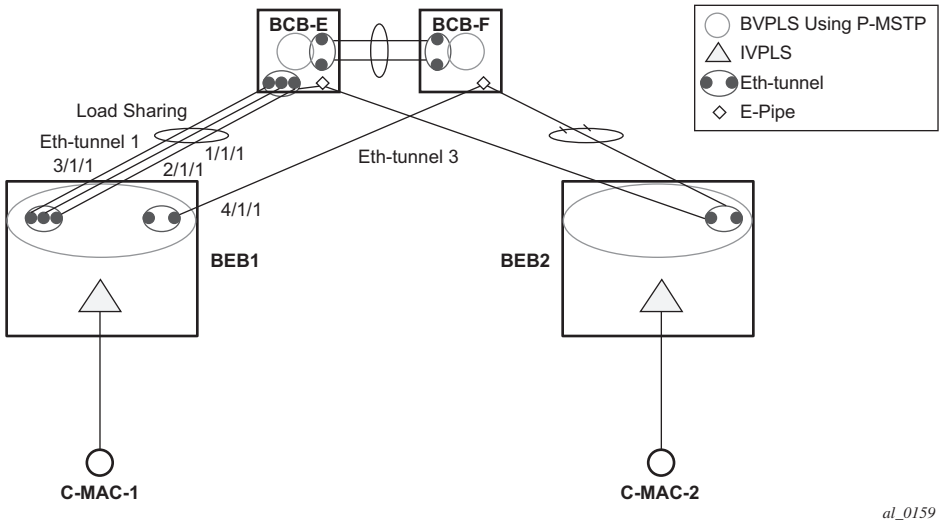

For a PBB Ethernet VLAN to make efficient use of an emulated LAG solution, a Management-VPLS (m-VPLS) is configured enabling Provider Multi-Instance Spanning Tree Protocol (P-MSTP). The m-VPLS is assigned to two SAPs; the eth-tunnels connecting BEB1 to BCB-E and BCB-F, respectively, reserving a range of VLANs for P-MSTP.

The PBB P-MSTP service is represented in the BEBs as a combination of an Epipe mapped to a BVPLS instance as before but now the PBB service is able to use the Ethernet tunnels under the P-MSTP control and load share traffic on the emulated LAN. In our example, the blue-circle representing the BVPLS is assigned to the SAPs which define two paths each. All paths are specified as primary precedence to load share the traffic.

A Management VPLS (m-VPLS) is first configured with a VLAN-range and assigned to the SAPs containing the path to the BCBs. The load shared eth-tunnel objects are defined by specifying a member ports and a control tag of zero. Then individual B-VPLS services can be assigned to the member paths of the emulated LAGs and defining the path encapsulation. Then individual services such as the IVPLS service can be assigned to the B-VPLS.

At the BCBs the tunnels are terminated the next BVPLS instance controlled by P-MSTP on the BCBs to forward the traffic.

In the event of link failure, the emulated LAG group will automatically adjust the number of paths. A threshold can be set whereby the LAG group is declared down. All emulated LAG operations are independent of any 8031-1to1 operation.

4.2.15.4 Support Service and Solution Combinations

The following considerations apply when Ethernet tunnels are configured under a VPLS service:

- Only ports in access or hybrid mode can be configured as eth-tunnel path members. The member ports can be located on the same or different IOMs, MDAs, XCMs, or XMAs.
- Dot1q and QinQ ports are supported as eth-tunnel path members.
- The same port cannot be used as member in both a LAG and an Ethernet-tunnel.
- A mix of regular and multiple eth-tunnel SAPs and PWs can be configured in the same BVPLS.
- Split horizon groups in BVPLS are supported on eth-tunnel SAPs. The use of split horizon groups allows the emulation of a VPLS model over the native Ethernet core, eliminating the need for P-MSTP.
- STP and MMRP are not supported in a BVPLS using eth-tunnel SAPs.
- Both PBB E-Line (Epipe) and E-LAN (IVPLS) services can be transported over a BVPLS using Ethernet-tunnel SAPs.
- MC-LAG access multi-homing into PBB services is supported in combination with Ethernet tunnels:
  - MC-LAG SAPs can be configured in IVPLS or Epipe instances mapped to a BVPLS that uses eth-tunnel SAPs
  - Blackhole Avoidance using native PBB MAC flush/MAC move solution is also supported

- Support is also provided for BVPLS with P-MSTP and MMRP control plane running as ships-in-the-night on the same links with the Ethernet tunneling which is mapped by a SAP to a different BVPLS.
  - Epipes must be used in the BCBs to support scalable point-to-point tunneling between the eth-tunnel endpoints when management VPLS is used.
- The following solutions or features are not supported in the current implementation for the 7450 ESS and 7750 SR and are blocked:
  - Capture SAP
  - Subscriber management
  - BSX
  - Eth-tunnels usage as a logical port in the **config>redundancy>multi-chassis>peer>sync>port** context

For further information, refer to the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide*.

4.2.16 Periodic MAC Notification

Virtual BMAC learning frames (for example, the frames sent with the source MAC set to the virtual BMAC) can be sent periodically, allowing all BCBs/BEBs to keep the virtual BMAC in their Layer 2 forwarding database.

This periodic mechanism is useful in the following cases:

- A new BEB is added after the current mac-notification method has stopped sending learning frames.
- When a new combination of [MC-LAG:SAP|A/S PW]+[PBB-Epipe]+[associated B-VPLS]+[at least one B-SDP|B-SAP] becomes active. The current mechanism only sends learning frames when the first such combination becomes active.
- A BEB containing the remote endpoint of a dual-homed PBB-epipe is rebooted.
- When traffic is not seen for the MAC aging timeout (assuming that the new periodic sending interval is less than the aging timeout).
- When there is uni-directional traffic.

In each of the above cases, all of the remote BEB/BCBs will learn the virtual MAC in the worse case after the next learning frame is sent.

In addition, this will allow all of the above when to be used in conjunction with discard-unknown in the B-VPLS. Currently, if discard-unknown is enabled in all related B-VPLSs (to avoid any traffic flooding), all above cases could experience an increased traffic interruption, or a permanent loss of traffic, as only traffic toward the dual homed PBB-epipe can restart bi-directional communication. For example, it will reduce the traffic outage when:

The PBB-Epipe virtual MAC is flushed on a remote BEB/BCB due to the failover of an MC-LAG or A/S pseudowires within the customer's access network, for example, in between the dual homed PBB-Epipe peers and their remote tunnel endpoint.

There is a failure in the PBB core causing the path between the two BEBs to pass through a different BCB.

It should be noted that this will not help in the case where the remote tunnel endpoint BEB fails. In this case traffic will be flooded when the remote BMAC ages out if discard-unknown is disabled. If discard-unknown is enabled, then the traffic will follow the path to the failed BEB but will eventually be dropped on the source BEB when the remote BMAC ages out on all systems.

To scale the implementation it is expected that the timescale for sending the periodic notification messages is much longer than that used for the current notification messages.

4.2.17 MAC Flush

4.2.17.1 PBB Resiliency for B-VPLS Over Pseudowire Infrastructure

The following VPLS resiliency mechanisms are also supported in PBB VPLS:

- Native Ethernet resiliency supported in both I-VPLS and B-VPLS contexts
- Distributed LAG, MC-LAG, RSTP
- MSTP in a management VPLS monitoring (B- or I-) SAPs and pseudowire
- BVPLS service resiliency, loop avoidance solutions – Mesh, active/standby pseudowires and multi-chassis endpoint
- IVPLS service resiliency, loop avoidance solutions – Mesh, active/standby pseudowires (PE-rs only role), BGP Multi-homing

To support these resiliency options, extensive support for blackhole avoidance mechanisms is required.

4.2.17.1.1 Porting existing VPLS LDP MAC Flush in PBB VPLS

Both the I-VPLS and B-VPLS components inherit the LDP MAC flush capabilities of a regular VPLS to fast age the related FDB entries for each domain: CMACs for I-VPLS and BMACs for B-VPLS. Both types of LDP MAC flush are supported for I-VPLS and B-VPLS domains:

- **flush-all-but-mine** - flush on positive event, for example:
  - Pseudowire activation — VPLS resiliency using active/standby pseudowire
  - Reception of a STP TCN
- **flush-all-from-me** - flush on negative event, for example:
  - SAP failure – link down or MC-LAG out-of-sync
  - Pseudowire or Endpoint failure

In addition, only for the B-VPLS domain, changing the backbone source MAC of a B-VPLS will trigger an LDP MAC flush-all-from-me to be sent in the related active topology. At the receiving PBB PE, a BMAC flush automatically triggers a flushing of the CMACs associated with the old source BMAC of the B-VPLS.

4.2.17.1.2 PBB Blackholing Issue

In the PBB VPLS solution, a B-VPLS may be used as infrastructure for one or more I-VPLS instances. B-VPLS control plane (LDP Signaling or P-MSTP) replaces I-VPLS control plane throughout the core. This is raising an additional challenge related to blackhole avoidance in the I-VPLS domain as described in this section.

PBB Blackholing Issue — Assuming that the link between PE A1 and node 5 is active, the remote PEs participating in the orange VPN (for example, PE D) will learn the CMAC X associated with backbone MAC A1. Under failure of the link between node 5 and PE A1 and activation of link to PE A2, the remote PEs (for example, PE D) will black-hole the traffic destined for customer MAC X to BMAC A1 until the aging timer expires or a packet flows from X to Y through the PE A2. This may take a long time (default aging timer is 5 minutes) and may affect a large number of flows across multiple I-VPLSs.

A similar issue will occur in the case where node 5 is connected to A1 and A2 I-VPLS using active/standby pseudowires. For example, when node 5 changes the active pseudowire, the remote PBB PE will keep sending to the old PBB PE.

Another case is when the QinQ access network dual-homed to a PBB PE uses RSTP or MVPLS with MSTP to provide loop avoidance at the interconnection between the PBB PEs and the QinQ SWs. In the case where the access topology changes, a TCN event will be generated and propagated throughout the access network. Similarly, this change needs to be propagated to the remote PBB PEs to avoid blackholing.

A solution is required to propagate the I-VPLS events through the backbone infrastructure (B-VPLS) in order to flush the customer MAC to BMAC entries in the remote PBB. As there are no IVPLS control plane exchanges across the PBB backbone, extensions to B-VPLS control plane are required to propagate the I-VPLS MAC flush events across the B-VPLS.

4.2.17.1.3 LDP MAC Flush Solution for PBB Blackholing

In the case of an MPLS core, B-VPLS uses T-LDP signaling to set up the pseudowire forwarding. The following I-VPLS events must be propagated across the core B-VPLS using LDP MAC **flush-all-but-mine** or **flush-all-from-me** indications:

For **flush-all-but-mine** indication (“positive flush”):

- TCN event in one or more of the I-VPLS or in the related M-VPLS for the MSTP use case.
- Pseudowire/SDP binding activation with Active/Standby pseudowire (standby, active or down, up)
- Reception of an LDP MAC withdraw “flush-all-but-mine” in the related I-VPLS

For **flush-all-from-me** indication (“negative flush”):

- MC-LAG failure - does not require send-flush-on-failure to be enabled in I-VPLS
- Failure of a local SAP – requires send-flush-on-failure to be enabled in I-VPLS
- Failure of a local pseudowires/SDP binding – requires send-flush-on-failure to be enabled in I-VPLS
- Reception of an LDP MAC withdraw flush-all-from-me in the related I-VPLS

To propagate the MAC flush indications triggered by the above events, the PE that originates the LDP MAC withdraw message must be identified. In regular VPLS “mine”/”me” is represented by the pseudowire associated with the FEC and the T-LDP session on which the LDP MAC withdraw was received. In PBB, this is achieved using the B-VPLS over which the signaling was propagated and the BMAC address of the originator PE.

Nokia PBB-VPLS solution addresses this requirement by inserting in the BVPLS LDP MAC withdraw message a new PBB-TLV (type-length-value) element. The new PBB TLV contains the source BMAC identifying the originator (“mine”/”me”) of the flush indication and the ISID list identifying the I-VPLS instances affected by the flush indication.

There are a number of advantages to this approach. Firstly, the PBB-TLV presence indicates this is a PBB MAC Flush. As a result, all PEs containing only the B-VPLS instance will automatically propagate the LDP MAC withdraw in the B-VPLS context respecting the split-horizon and active link topology. There is no flushing of the B-VPLS FDBs throughout the core PEs. Subsequently, the receiving PBB VPLS PEs uses the BMAC and ISID list information to identify the specific I-VPLS FDBs and the CMAC entries pointing to the source BMAC included in the PBB TLV.

An example of processing steps involved in PBB MAC Flush is depicted in Figure 127 for the case when a Topology Change Notification (TCN) is received on PBB PE 2 from a QinQ access in the I-VPLS domain.

Figure 127 **TCN Triggered PBB Flush-ALI-But-Mine Procedure**

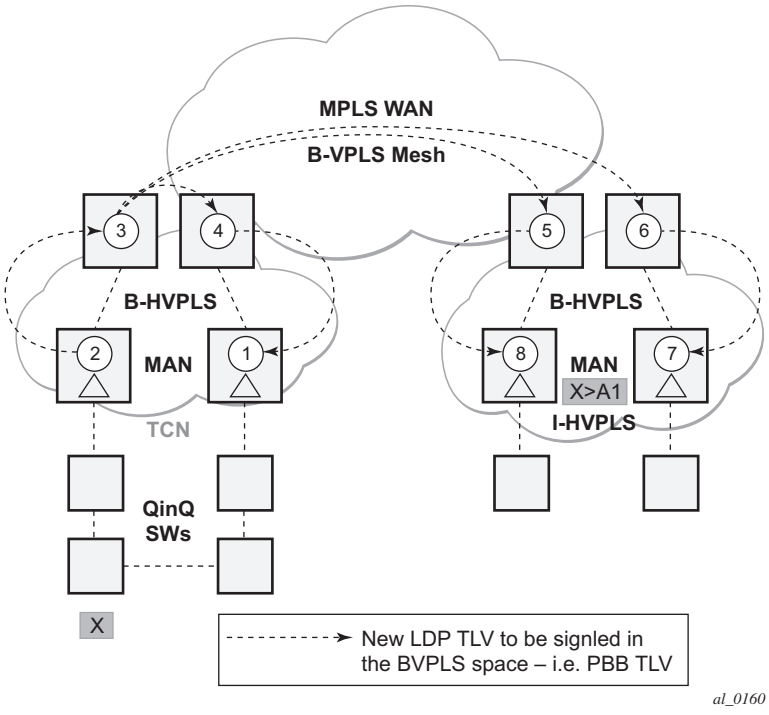

The received TCN may be related to one or more I-VPLS domains. This will generate a MAC Flush in the local I-VPLS instance(s) and if configured, it will originate a PBB MAC **flush-all-but-mine** throughout the related B-VPLS context(s) represented by the white circles 1 to 8 in our example.

A PBB-TLV is added by PE2 to the regular LDP MAC **flush-all-but-mine**. BMAC2, the source BMAC associated with B-VPLS on PE2 is carried inside the PBB TLV to indicate who “mine” is. The ISID list identifying the I-VPLS affected by the TCN is also included if the number of affected I-VPLS is 100 or less. No ISID list is included in the PBB-TLV if more than 100 ISIDs are affected. If no ISID list is included, then the receiving PBB PE will flush all the local I-VPLS instances associated with the B-VPLS context identified by the FEC TLV in the LDP MAC withdraw message. This is done to speed up delivery and processing of the message.

Recognizing the PBB MAC flush, the B-VPLS only PEs 3, 4, 5 and 6 refrain from flushing their B-VPLS FDB tables and propagate the MAC flush message regardless of their “propagate-mac-flush” setting.

When LDP MAC withdraw reaches the terminating PBB PEs 1 and 7, the PBB-TLV information is used to flush from the I-VPLS FDBs all CMAC entries except those associated with the originating BMAC BM2. If specific I-VPLS ISIDs are indicated in the PBB TLV, then the PBB PEs will flush only the CMAC entries from the specified I-VPLS except those mapped to the originating BMAC. Flush-all-but-mine indication is not propagated further in the I-VPLS context to avoid information loops.

The other events that trigger Flush-all-but-mine propagation in the B-VPLS (pseudowire/SDP binding activation, Reception of an LDP MAC Withdraw) are handled similarly. The generation of PBB MAC flush-all-but-mine in the B-VPLS must be activated explicitly on a per I-VPLS basis with the command **send-bvpls-flush all-but-mine**. The generation of PBB MAC flush-all-from-me in the B-VPLS must be activated explicitly on a per I-VPLS basis with the command **send-bvpls-flush all-from-me**.

4.2.18 Access Multi-Homing for Native PBB (B-VPLS over SAP Infrastructure)

Nokia PBB implementation allows the operator to use a native Ethernet infrastructure as the PBB core. Native Ethernet tunneling can be emulated using Ethernet SAPs to interconnect the related B-VPLS instances. This kind of solution might fit certain operational environments where Ethernet services was provided in the past using QinQ solution. The drawback is that no LDP signaling is available to provide support for Access Multi-homing for Epipe (pseudowire Active/Standby status) or I-VPLS services (LDP MAC Withdraw). An alternate solution is required.

A PBB network using Native Ethernet core is depicted in Figure 128. MC-LAG is used to multi-home a number of edge switches running QinQ to PBB BEBs.

Figure 128 **Access Dual-Homing into PBB BEBs - Topology View**

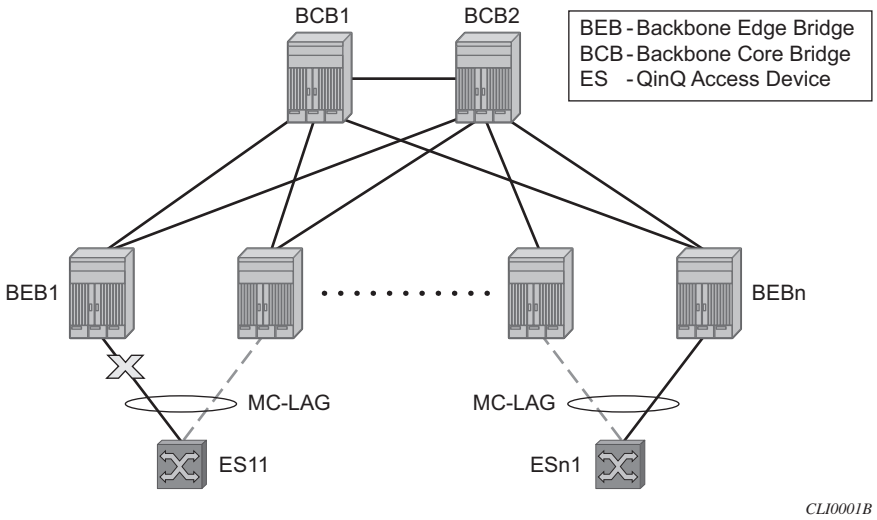

The interrupted line from the MC-LAG represents the standby, inactive link; the solid line is the active link. The BEBs are dual-homed to two core switches BCB1 and BCB2 using native Ethernet SAPs on the B-VPLS side. Multi-point B-VPLS with MSTP for loop avoidance can be used as the PBB core tunneling. Alternatively point-to-point, G.8031 protected Ethernet tunnels can be also used to interconnect B-VPLS instances in the BEBs as described in the PBB over G.8031 protected Ethernet tunnels.

Nokia implementation provides a solution for both PBB E-Line (Epipe) and E-LAN (IVPLS) services that avoids PBB blackholing when the active ES11-BEB1 link fails. It also provides a consistent behavior for both service type and for different backbone types: for example, native Ethernet, MPLS, or a combination. Only MC-LAG is supported initially as the Access-Multi-homing mechanism.

4.2.18.1 Solution Description for I-VPLS Over Native PBB Core

The use case described in the previous section is addressed by enhancing the existing native PBB solution to provide for blackhole avoidance.

The topology depicted in Figure 129 describes the details of the solution for the I-VPLS use case. Although the native PBB use case is used, the solution works the same for any other PBB infrastructure: for example, G.8031 Ethernet tunnels, pseudowire/MPLS, or a combination.

Figure 129 **PBB Active Topology and Access Multi-Homing**

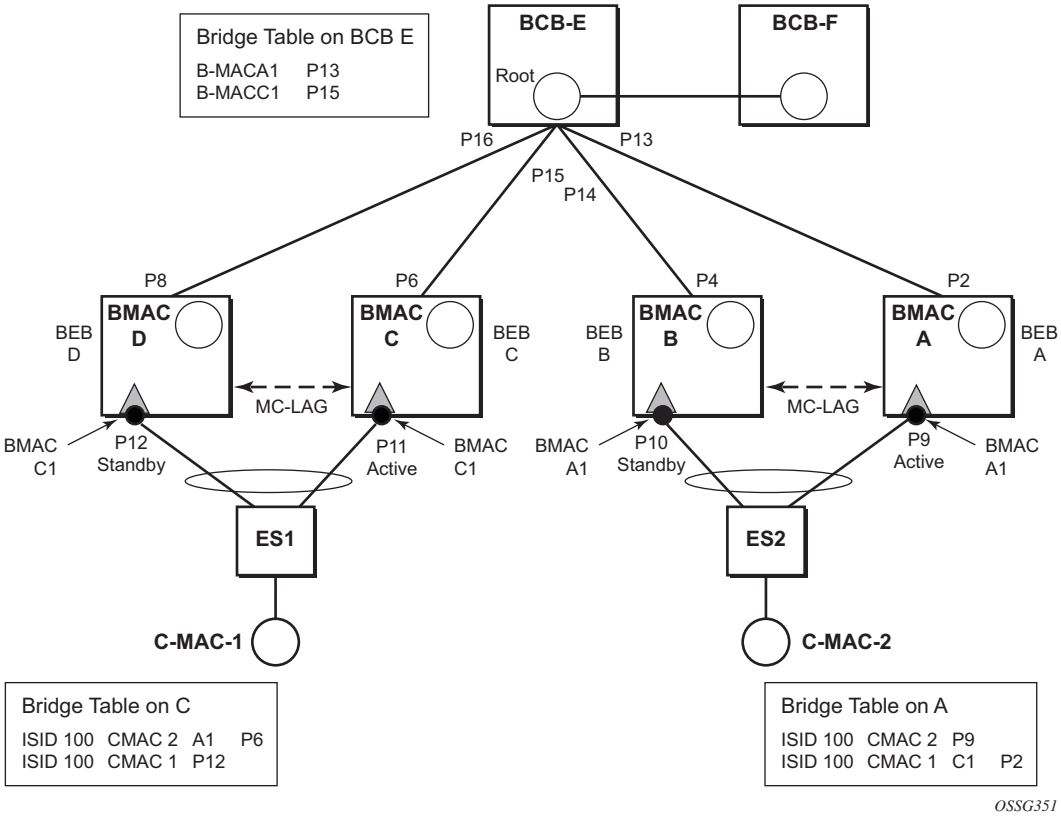

ES1 and ES2 are dual-homed using MC-LAG into two BEB devices: ES1 to BEB C and BEB D, ES2 to BEB A and BEB B. MC-LAG P11 on BEB C and P9 on BEB A are active on each side.

In the service context, the triangles are I-VPLS instances while the small circles are B-VPLS components with the related, per BVPLS source BMACs indicated next to each BVPLS instances. P-MSTP or RSTP may be used for loop avoidance in the multi-point BVPLS. For simplicity, only the active SAPs (BEB P2, P4, P6 and P8) are shown in the diagram.

In addition to the source BMAC associated with each BVPLS, there is an additional BMAC associated with each MC-LAG supporting multi-homed I-VPLS SAPs. The BEBs that are in a multi-homed MC-LAG configuration share a common B-MAC on the related MC-LAG interfaces. For example, a common BMAC C1 is associated in this example with ports P11 and P12 participating in the MC-LAG between BEB C and BEB D while BMAC A1 is associated with ports P9 and P10 in the MC-LAG between BEB A and BEB B. While BMAC C1 is associated through the I-VPLS SAPs with both BVPLS instances in BEB C and BEB D, it is actively used for forwarding to I-VPLS SAPs only on BEB C containing the active link P11.

MC-LAG protocol keeps track of which side (port or LAG) is active and which is standby for a specified MC-LAG grouping and activates the standby in case the active one fails. The source BMAC C1 and A1 are used for PBB encapsulation as traffic arrives at the IVPLS SAPs on P11 and P9, respectively. MAC Learning in the BVPLS instances installs MAC FDB entries in BCB-E and BEB A as depicted in Figure 129.

Active link (P11) or access node (BEB C) failures are activating through MC-LAG protocol the standby link (P12) participating in the MC-LAG on the pair MC-LAG device (BEB D).

Figure 130 shows the case of access link failure.

Figure 130 **Access Multi-Homing - Link Failure**

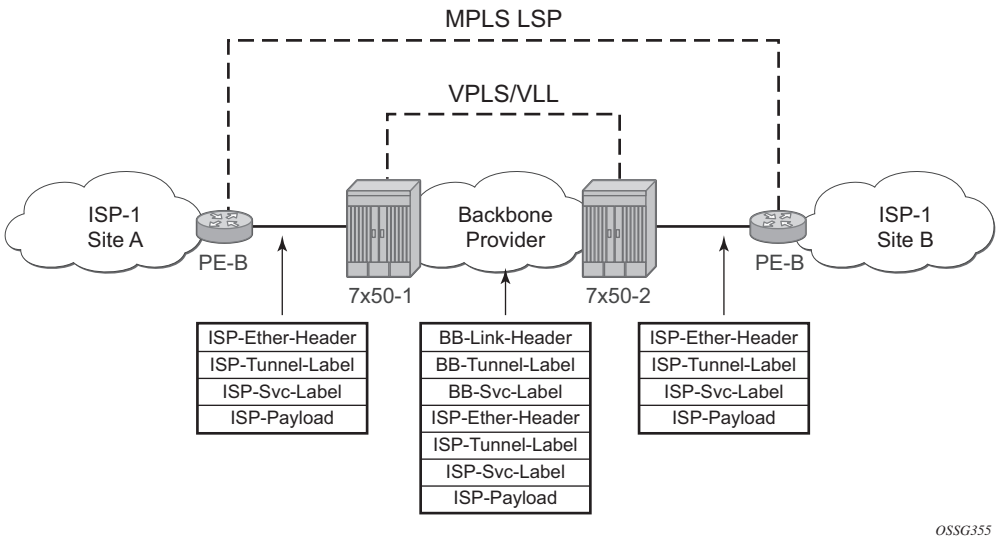

On failure of the active link P11 on BEB C the following processing steps apply:

- MC-LAG protocol activates the standby link P12 on the pair BEB D.
- BMAC C1 becomes active on BEB D and any traffic received on BEB D with destination BMAC C1 is forwarded on the corresponding I-VPLS SAPs on P12.
- BEB D determines the related B-VPLS instance(s) associated with all the I-VPLS SAP(s) mapped to P12, the newly activated MC-LAG link(s)/LAG component(s).
- Subsequently, BEB D floods in the related B-VPLS instance(s) an Ethernet CFM-like message using C1 as source BMAC. A vendor CFM opcode is used followed by an Nokia OUI.
- As a result, all the FDB entries in BCBs or BEBs along the path will be automatically updated to reflect the move of BMAC C1 to BEB D.

- In this particular configuration the entries on BEB A do not need to be updated saving MAC Flush operation.
- In other topologies, it is possible that the BMAC C1 FDB entries in the B-VPLS instance on the remote BEBs (like BEB A) will need to move between B-SAPs. This will involve a move of all CMAC using as next hop BMAC C1 and the new egress line card.

Identical procedure is used when the whole BEB C fails.

4.2.18.2 Solution Description for PBB Epipe over G.8031 Ethernet Tunnels

This section discusses the Access multi-homing solution for PBB E-Line over an infrastructure of G.8031 Ethernet tunnels. Although a specific use case is used, the solution works the same for any other PBB infrastructure: for example, native PBB, pseudowire/MPLS, or a combination.

The PBB E-Line service and the related BVPLS infrastructure are depicted in Figure 131.

Figure 131 **Access Multi-Homing Solution for PBB Epipe**

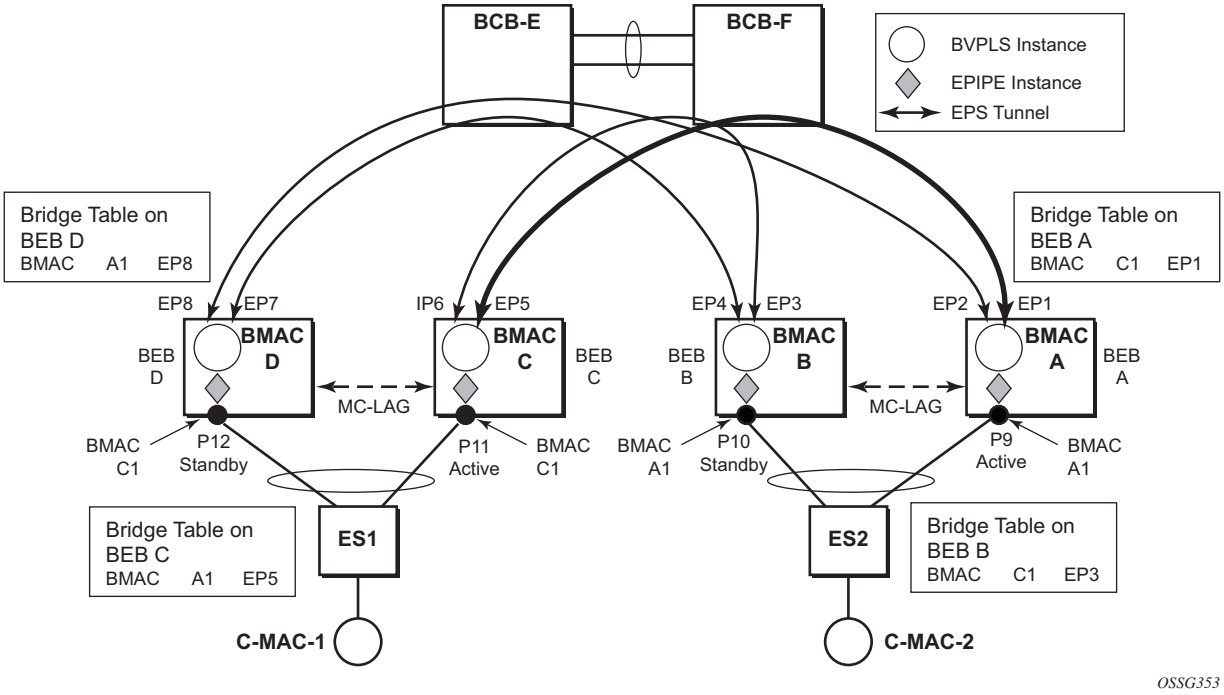

The E-Line instances are connected through the B-VPLS infrastructure. Each B-VPLS is interconnected to the BEBs in the remote pair using the G.8031, Ethernet Protection Switched (EPS) tunnels. Only the active Ethernet paths are shown in the network diagram to simplify the explanation. Split Horizon Groups may be used on EPS tunnels to avoid running MSTP/RSTP in the PBB core.

The same BMAC addressing scheme is used as in the E-LAN case: a BMAC per B-VPLS and additional BMACs associated with each MC-LAG connected to an Epipe SAP. The BMACs associated with the active MC-LAG are actively used for forwarding into B-VPLS the traffic ingressing related Epipe SAPs.

MC-LAG protocol keeps track of which side is active and which is standby for a specified MC-LAG grouping and activates the standby link in a failure scenario. The source BMACs C1 and A1 are used for PBB encapsulation as traffic arrives at the Epipe SAPs on P11 and P9, respectively. MAC Learning in the B-VPLS instances installs MAC FDB entries in BEB C and BEB A as depicted in Figure 131. The highlighted Ethernet tunnel (EPS) will be used to forward the traffic between BEB A and BEB C.

Active link (P11) or access node (BEB C) failures are activating through MC-LAG protocol, the standby link (P12) participating in the MC-LAG on the pair MC-LAG device (BEB D). The failure of BEB C is depicted in Figure 132. The same procedure applies for the link failure case.

Figure 132 **Access Dual-Homing for PBB E-Line - BEB Failure**

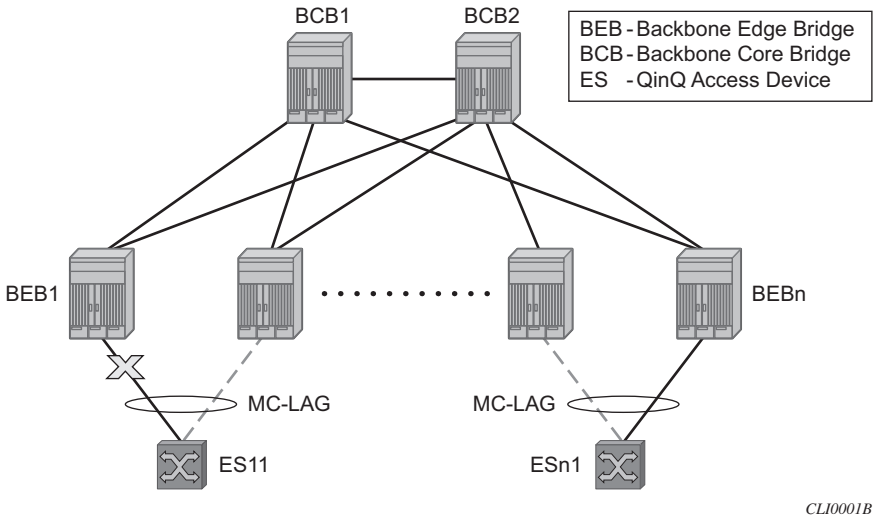

The following process steps apply:

- BEB D will lose MC-LAG communication with its peer BEB C - no more keep-alives from BEB C or next-hop tracking may kick in.

- BEB D assumes BEB C is down and activates all shared MC-LAG links, including P12.
- BMAC C1 becomes active on BEB D and any traffic received on BEB C with destination BMAC C1 is forwarded on the corresponding Epipe SAPs on P12.
- BEB D determines the related B-VPLS instance(s) associated with all the Epipe SAP(s) mapped to P12, the newly activated MC-LAG link(s)/LAG component(s).
- Subsequently, BEB D floods in the related B-VPLS instance(s) the same Ethernet CFM message using C1 as source BMAC.
- As a result, the FDB entries in BEB A and BEB B will be automatically updated to reflect the move of BMAC C1 from EP1 to EP2 and from EP3 to EP4, respectively.

The same process is executed for all the MC-LAGs affected by BEB C failure so BEB failure will be the worst case scenario.

4.2.18.2.1 Dual-Homing into PBB Epipe - Local Switching Use Case

When the service SAPs were mapped to MC-LAGs belonging to the same pair of BEBs in earlier releases, an IVPLS had to be configured even if there were just two SAPs active at any point in time. Since then, the PBB Epipe model has been enhanced to support configuring in the same Epipe instance two SAPs and a BVPLS uplink as depicted in Figure 133.

Figure 133 **Solution for Access Dual-Homing with Local Switching for PBB E-Line/Epipe**

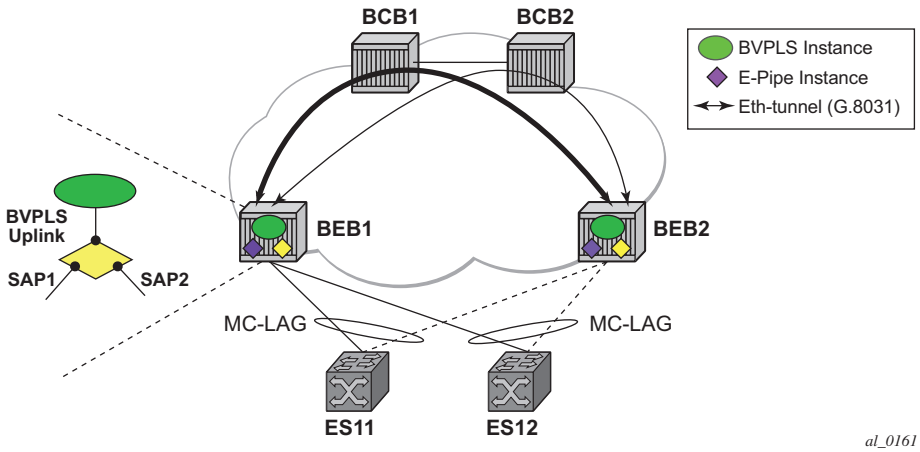

The PBB Epipe represented by the yellow diamond on BEB1 points through the BVPLS uplink to the BMAC associated with BEB2. The destination BMAC can be either the address associated with the green BVPLS on BEB2 or the BMAC of the SAP associated with the pair MC-LAG on BEB2 (preferred option).

The Epipe information model is expanded to accommodate the configuration of two SAPs (I-SAPs) and of a BVPLS uplink in the same time. For this configuration to work in an Epipe environment, only two of them will be active in the forwarding plane at any point in time, specifically:

- SAP1 and SAP2 when both MC-LAG links are active on the local BEB1 (see Figure 133)
- The Active SAP and the BVPLS uplink if one of the MC-LAG links is inactive on BEB1
  - PBB tunnel will be considered as a backup path only when the SAP is operationally down.
  - If the SAP is administratively down, then all traffic will be dropped.
- Although the CLI allows configuration of two SAPs and a BVPLS uplink in the same PBB Epipe, the BVPLS uplink is inactive as long as both SAPs are active.
  - Traffic received through PBB tunnel is dropped if BVPLS uplink is inactive.
- The same rules apply to BEB2.

4.2.19 BGP Multi-homing for I-VPLS

This section describes the application of BGP multi-homing to I-VPLS services. BGP multi-homing for I-VPLS uses the same mechanisms as those used when BGP multi-homing is configured in a non-PBB VPLS service, which are described in detail in this guide.

The multi-homed sites can be configured with either a SAP or spoke-SDP, and support both split horizon groups and fate-sharing by the use of oper-groups.

When the B-VPLS service is using LDP signaled pseudowires, blackhole protection is supported after a multi-homing failover event when **send-flush-on-failure** and **send-bvpls-flush flush-all-from-me** is configured within the I-VPLS. This causes the system on which the site object fails to send a MAC flush all-from-me message so that customer MACs are flushed on the remote backbone edge bridges using a flush-all-from-me message. The message sent includes a PBB TLV which contains the source BMAC identifying the originator (“mine”/“me”) of the flush indication and the ISID list identifying the I-VPLS instances affected by the flush indication, see section LDP MAC Flush Solution for PBB Blackholing.

The VPLS preference sent in BGP multi-homing updates will be always be set to zero, however, if a non-zero value is received in a valid BGP multi-homing update it will be used to influence the designated forwarder (DF) election.

4.2.20 Access Multi-Homing over MPLS for PBB Epipes

It is possible to connect backbone edge bridges (BEBs) configured with PBB Epipes to an edge device using active/standby pseudowires over an MPLS network. This is shown in Figure 134.

Figure 134 **Active/Standby PW into PBB Epipes**

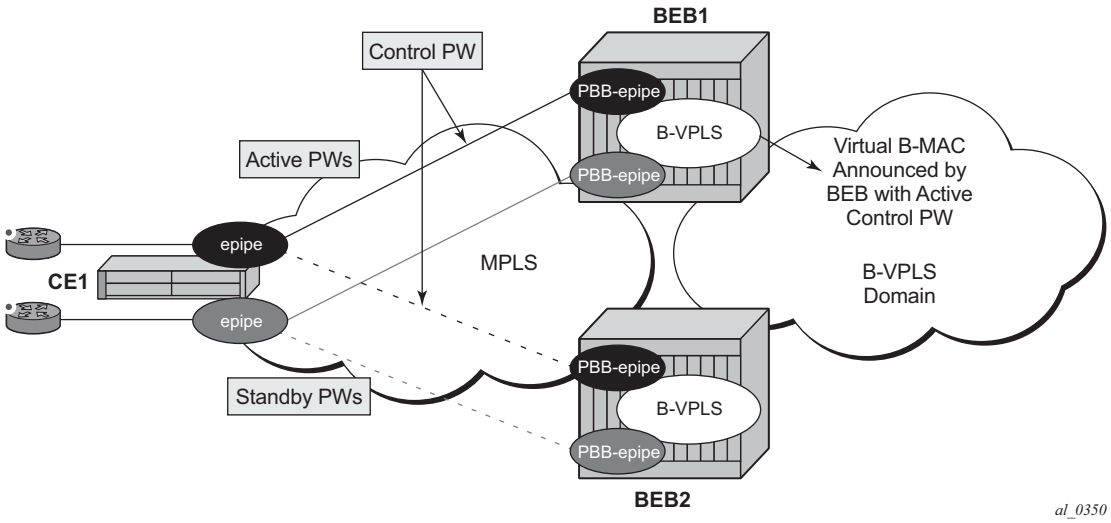

In this topology, the edge device (CE1) is configured with multiple Epipes to provide virtual lease line (VLL) connectivity across a PBB network. CE1 uses active/standby pseudowires (PWs) which terminate in PBB Epipe services on BEB1 and BEB2 and are signaled accordingly using the appropriate pseudowire status bits.

Traffic is sent from CE1 on the active pseudowires into the PBB epipe services, then onto the remote devices through the B-VPLS service. It is important that traffic sent to CE1 is directed to the BEB that is attached to the active pseudowire connected to CE1. To achieve this, a virtual backbone MAC (vBMAC) is associated with the services on CE1.

The vBMAC is announced into the PBB core by the BEB connected to the active pseudowire using SPBM configured in the B-VPLS services; hence SPBM is mandatory. In Figure 134, the vBMAC would be announced by BEB1; if the pseudowires failed over to BEB2, BEB1 would stop announcing the vBMAC and BEB2 will start announcing it.

The remote services are configured to use the vBMAC as the backbone destination MAC (backbone-dest-mac) which results in traffic being sent to the specified BEB.

The vBMAC is configured under the SDP used to connect to the edge device's active/standby pseudowires using the command source-bmac-lsb. This command defines a sixteen (16) bit value which overrides the sixteen least-significant-bits of source backbone MAC (source-bmac) to create the vBMAC. The operator must ensure that the vBMACs match on the two peering BEBs for a corresponding SDP.

The PBB Epipe pseudowires are identified to be connected to an edge device active/standby pseudowire using the spoke-sdp parameter use-sdp-bmac. Enabling this parameter will cause traffic forwarded from this spoke-SDP into the B-VPLS domain to use the vBMAC as its source MAC address when both this, and the control pseudowire, are in the active state on this BEB.

PBB Epipe pseudowires connected to edge device's non-active/standby pseudowires are still able to use the same SDP.

To cater for the case where there are multiple edge device active/standby pseudowires using a specified SDP, one pseudowire must be identified to be the control pseudowire (using the source-bmac-lsb parameter control-pw-vc-id). The state of the control pseudowire determines the announcing of the vBMAC by SPBM into the B-VPLS based on the following conditions:

- The source-bmac-lsb and control-pw-vc-id have both been configured.
- The spoke-SDP referenced by the control-pw-vc-id has use-sdp-bmac configured.
- The spoke-SDP referenced by the control-pw-vc-id is operationally up and the "Peer Pw Bits" do not include pwFwdingStandby.
- If multiple B-VPLS services are used with different SPBM Forward IDs (FIDs), the vBMAC is advertised into any FID which has a PBB Epipe with a spoke-SDP configured with use-sdp-bmac that is using an SDP with source-bmac-lsb configured (regardless of whether the PBB Epipe spoke-SDP defined as the control pseudowire is associated with the B-VPLS).

It is expected that pseudowires configured using an SDP with source-bmac-lsb and with the parameter use-sdp-bmac are in the same state (up, down, active, standby) as the control pseudowire. If this is not the case, the following scenarios are possible (based on Figure 134):

- If any non-control pseuodwires are active on BEB2 and standby on BEB1, then this will continue to allow bi-directional traffic for the related services as the return traffic to CE1 will be sent to BEB1, specifically to the BEB announcing the vBMAC. As the non-control PW is in standby state it will be used to send this traffic to the edge device. If this operation is not needed, it is possible to prevent traffic being sent on a standby PW using the standby-signaling-slave parameter under the spoke-SDP definition.
- If any non-control pseudowires are active on BEB2 but down on BEB1, then only uni-directional traffic is possible. The return traffic to CE1 will be sent to BEB1, as it is announcing the vBMAC but the pseudowire on BEB1 is down for this service.

Alarms are raised to track if, on the BEB with the control pseudowire in the standby/down state, any non-control pseudowires go active. Specifically, there will be an alarm when the first non-control pseudowire becomes active and another alarm when the last non-control pseudowire becomes standby/down.

If both control pseudowires are active (neither in standby) then both BEBs would announce the vBMAC – this would happen if the edge device was a 7450 ESS, 7750 SR, and 7950 XRS using an Epipe service without standby-signaling-master configured. Traffic from remote BEBs on any service related to the vBMAC would be sent to the nearest SPBM BEB and it would depend on the state of the pseudowires on each BEB as to whether it could reach the edge device. Similarly, the operator must ensure that the corresponding service pseudowires on each BEB are configured as the control pseudowire, otherwise SPBM might advertise the vBMAC from both BEBs resulting in the same consequences.

All traffic received from the edge device on a pseudowire into a PBB Epipe, on the BEB with the active control pseudowire, is forwarded by the B-VPLS using the vBMAC as the source backbone MAC, otherwise the source-bmac is used.

The control pseudowire can be changed dynamically without shutting down the spoke-SDPs, SDP or withdrawing the SPBM advertisement of the vBMAC; this allows a graceful change of the control pseudowire. Clearly, any change should be performed on both BEBs as closely in time as possible to avoid an asymmetric configuration, ensuring that the new control pseudowire is in the same state as the current control pseudowire on both BEBs during the change.

The following are not supported:

- Active/standby pseudowires within the PBB Epipe are not supported, consequently the following are not supported:
  - The configuration of endpoints.
  - The configuration of precedence under the spoke-SDP.
- The use of PW switching.

- BGP-MH support, namely configuring the pseuodwires to be part of a multi-homed site.
- Network-domains.
- Support for the following tunneling technologies
  - RFC 3107
  - GRE
  - L2TPv3

4.2.21 PBB and IGMP/MLD Snooping

The IGMP/MLD snooping feature provided for VPLS is supported similarly in the PBB I-VPLS context, in order to provide efficient multicast replication in the customer domain. The difference from regular VPLS is the handling of IGMP/MLD messages arriving from the B-VPLS side over a B-VPLS SAP or SDP.

The first IGMP/MLD join message received over the local B-VPLS adds all the B-VPLS SAP and SDP components into the related multicast table associated with the I-VPLS context. This is in line with the PBB model, where the B-VPLS infrastructure emulates a backbone LAN to which every I-VPLS is connected by one virtual link.

When the querier is connected to a remote I-VPLS instance, over the B-VPLS infrastructure, its location is identified by the B-VPLS SDP and SAP on which the query was received. It is also identified by the source BMAC address used in the PBB header for the query message. This is the BMAC associated with the B-VPLS instance on the remote PBB PE.

It is also possible to configure that a multicast router exists in a remote I-VPLS service. This can be achieved using the **mrouter-dest** command to specify the MAC name of the destination BMAC to be used to reach the remote I-VPLS service. This command is available in the VPLS service PBB IGMP and MLD snooping contexts.

The following are not supported in a PBB I-VPLS context with IGMP snooping or MLD snooping:

- multicast VPLS Registration (MVR)
- multicast CAC
- configuration under a default SAP

The following are not supported in a PBB I-VPLS context with MLD snooping:

- configuration of the maximum number of multicast group sources allowed per group
- configuration of the maximum number of multicast sources allowed per group

4.2.22 PBB and PIM Snooping

The PIM snooping feature for IPv4 is supported in the PBB I-VPLS context in order to provide efficient multicast replication in the customer domain. This is similar to PIM snooping for IPv4 in a regular VPLS with the difference being the handling of PIM messages arriving from the B-VPLS side over a B-VPLS SAP or SDP.

The first PIM join message received over the local B-VPLS adds all the B-VPLS SAP and SDP components into the related multicast table associated with the I-VPLS context, and the multicast for the join is flooded throughout the B-VPLS. This is in line with the PBB model, where the B-VPLS infrastructure emulates a backbone LAN to which every I-VPLS is connected by one virtual link.

When a neighbor is located on a remote I-VPLS instance over the B-VPLS infrastructure, its location is identified by the B-VPLS SDP and SAP on which the hello message was received. The neighbor is also identified by the source BMAC address used in the PBB header of the hello message. This is the BMAC associated with the B-VPLS instance on the remote PBB PE.

PIM snooping for IPv4 in an I-VPLS is not supported with the following forms of default SAP:

- :*
- *.null
- *.*

4.2.23 PBB QoS

For PBB encapsulation, the configuration used for DE and dot1p in SAP and SDP policies applies to the related bits in both backbone dot1q (BTAG) and ITAG fields.

The following QoS processing rules apply for PBB B-VPLS SAPs and SDPs:

B-VPLS SAP ingress

- If dot1p, DE based classification is enabled, the BTAG fields will be used by default to evaluate the internal forwarding class (fc) and discard profile if there is a BTAG field. The 802.1ah ITAG will be used only if the BTAG is absent (null SAP).
- If either one of the dot1p or DE based classification is not explicitly enabled or the packets are untagged then the default fc and profile is assigned.

B-VPLS SAP egress

- If the sap-egress policy for the SAP contains an fc to dot1p/de mapping, this entry is used to set the dot1p and DE bits from the BTAG of the frame going out from the SAP. The same applies for the ITAG on frames originated locally from an I-VPLS. The mapping does not have any effect on the ITAG of frames transiting the B-VPLS.
- If no explicit mapping exists, the related dot1p DE bits are set to zero on both ITAG and BTAG if the frame is originated locally from an I-VPLS. If the frame is transiting the B-VPLS the ITAG stays unchanged, the BTAG is set according to the type of ingress SAP.
  - If the ingress SAP is tagged, the values of the dot1p, DE bits are preserved in the BTAG going out on the egress SAP.
  - If the ingress SAP is untagged, the dot1p, DE bits are set to zero in the BTAG going out on the egress SAP.

B-VPLS SDP (network) ingress policy

- QoS policies for dot1p and DE bits apply only for the outer VLAN ID: this is the VLAN ID associated with the link layer and not the PBB BTAG. As a result, the dot1p DE bits will be checked if an outer VLAN ID exists in the packets ingressing the SDP. If that VLAN ID is absent, nothing above the pseudowire SL will be checked - for example, no dot1p bits in the BTAG or ITAG will be checked. It is expected that the EXP bits will be used to transport QoS information across the MPLS backbone and into the PEs.

B-VPLS SDP (network) egress policy

- When building PBB packets originating from a local I-VPLS, the BTAG and ITAG values (dot1p, DE bits) will be set according to the network egress policy. The same applies for newly added BTAG (VLAN mode pseudowires) in a packet transiting the B-VPLS (SAP/SDP to SDP). If either dot1p or DE based classification is not explicitly enabled in the CLI, the values from the default fc to dot1p, DE mapping are assumed.
- Dot1p, DE bits for existing BTAGs will remain unchanged - for example, applicable to packets transiting the B-VPLS and going out on SDP.

4.2.23.1 Transparency of Customer QoS Indication through PBB Backbone

Similar to PW transport, operators want to allow their customers to preserve all eight Ethernet CoS markings (three dot1p bits) and the discard eligibility indication (DE bit) while transiting through a PBB backbone.

This means any customer CoS marking on the packets inbound to the ingress SAP must be preserved when going out on the egress SAP at the remote PBB PE even if the customer VLAN tag is used for SAP identification at the ingress.

A solution to the above requirements is depicted in Figure 135.

Figure 135 **PCP, DE Bits Transparency in PBB**

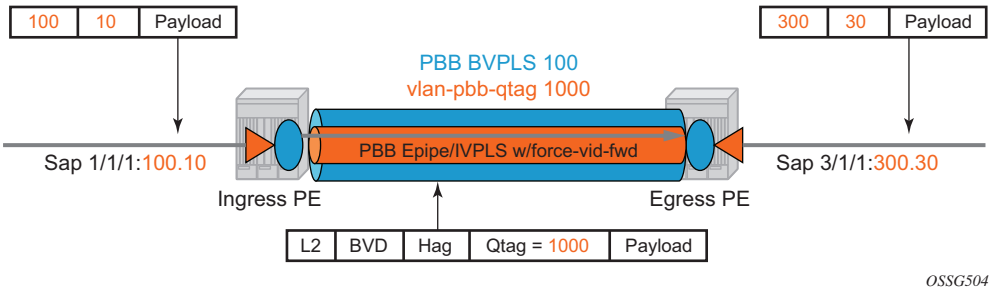

The PBB BVPLS is represented by the blue pipe in the middle with its associated CoS represented through both the service (I-tag) and tunnel CoS (BVID dot1p+DE or PW EXP bits).

The customer CoS is contained in the orange dot1q VLAN tags managed in the customer domains. There may be one (CVID) or two (CVID, SVID) tags used to provide service classification at the SAP. IVPLS or PBB Epipe instances (orange triangles) are used to provide a Carrier-of-Carrier service.

As the VLAN tags are stripped at the ingress SAP and added back at the egress SAP, the PBB implementation must provide a way to maintain the customer QoS marking. This is done using a force-qtag-forwarding configuration on a per IVPLS/ Epipe basis under the node specifying the uplink to the related BVPLS. When force-qtag-forwarding is enabled, a new VLAN tag is added right after the CMAC addresses using the configured QTAG. The dot1p, DE bits from the specified outer/ inner customer QTAG will be copied in the newly added tag.

When the force-qtag-forwarding is enabled in one IVPLS/PBB Epipe instance, it will be enabled in all of the related instances.

At the remote PBB PE/BEB on the egress SAPs or SDPs, the first QTAG after the CMAC addresses will be removed and its dot1p, DE bits will be copied in the newly added customer QTAGs.

4.2.23.1.1 Configuration Examples

This section gives usage examples for the new commands under PBB Epipe or IVPLS instances.

PBB IVPLS usage:

Example:
```
configure service vpls 100 ivpls
     sap 1/1/1:101
     pbb
               backbone-vpls 10 isid 100
               force-qtag-forwarding
```

PBB Epipe Usage:

Example:
```
configure service epipe 200
     sap 1/1/1:201
     pbb
               tunnel 10 backbone-dest-mac ab-bc-cd-ef-
                 01-01 isid 200
               force-qtag-forwarding
```

4.2.23.1.2 Details Solution Description

Figure 135 shows a specific use case. Keeping the same topology, an ingress PBB PE, a PBB core and an egress PBB PE, consider the generic use case where:

1. the packet arrives on the ingress PBB PE on an I-SAP or an I-SDP binding/PW and it is assigned to a PBB service instance (Epipe/IVPLS)
2. goes next through a PBB core (native Ethernet B-SAPs or PW/MPLS based B-SDP)
3. and finally, egresses at another PBB PE through a PBB service instance on either an I-SAP or I-SDP binding/PW.

Similar to the Ethernet-VLAN VC Type, the following packet processing steps apply for different scenarios.

- **Ingress PE, ingress I-SAP case** with force-qtag-forwarding enabled under PBB Epipe or IVPLS

The QTAG is inserted automatically right after CMAC addresses; an ethertype value of 8100 is used.

- **Case 1**: SAP type = null/dot1q default (1/1/1 or 1/1/1.*) so there is no service delimiting tag used and stripped on the ingress side.
  - VLAN and Dot1p+DE bits on the inserted QTAG are set to zero regardless of ingress QoS policy
- **Case 2**: SAP type = dot1q or qinq default (1/1/1.100 or 1/1/1.100.*) so there is a service delimiting tag used and stripped.
  - The service delimiting QTAG (dot1p + DE bits and VLAN) is copied as is in the inserted QTAG.
- **Case 3**: SAP type = qinq (1/1/1.100.10) so there are two service delimiting tags used and stripped.
  - The service delimiting QTAG (VLAN and dot1p + DE bits) is copied as is from the inner tag in the inserted QTAG.
- **Ingress PE, ingress I-SDP/PW case** with force-qtag-forwarding enabled under PBB Epipe or IVPLS

The QTAG is inserted automatically right after CMAC addresses; an ethertype value of 8100 is used.

- **Case 1**: SDP vc-type = Ethernet (force-vlan-vc-forwarding= not supported for I-PW) so there is no service delimiting tag stripped on the ingress side.
  - VLAN and Dot1p+DE bits on the inserted QTAG are set to zero regardless of ingress QoS policy
- **Case 2**: SDP vc-type = Ethernet VLAN so there is a service delimiting tag stripped.
  - VLAN and Dot1p + DE bits on the inserted QTAG are preserved from the service delimiting tag.

PBB packets are tunneled through the core the same way for native ETH/MPLS cases.

- **Egress PE, egress I-SAP case** with force-qtag-forwarding enabled under PBB Epipe or VPLS
  - The egress QoS policy (FC->dot1p+DE bits) is used to determine the QoS settings of the added QTAGs. If it required to preserve the ingress QoS, no egress policy should be added.
    - If QinQ SAP is used, at least qinq-mark-top-only option must be enabled to preserve the CTAG.
  - The “core QTAG” (core = received over the PBB core, 1st after CMAC addresses) is always removed after QoS information is extracted.
    - If no force-qtag-forwarding is used at egress PE, the inserted QTAG is maintained.

- If egress SAP is on the ingress PE, then the dot1p+DE value is read directly from the procedures described in Ingress PE, ingress I-SAP and Ingress PE, ingress I-SDP/PW cases. The use cases below still apply.
- **Case 1**: SAP type = null/dot1q default (2/2/2 or 2/2/2.*) so there is no service delimiting tag added on the egress side.
  - Dot1p+DE bits and the VLAN value contained in the QTAG are ignored.
- **Case 2**: SAP type = dot1q/qinq default (3/1/1.300 or 3/1/1.300.*) so a service delimiting tag is added on egress
  - The FC->dot1p, DE bit entries in the SAP egress QoS policy are applied.
  - If there are no such entries, then the values of the dot1p+DE bits from the stripped QTAG are used.
- **Case 3**: SAP type = qinq (3//1/1.300.30) so two service delimiting tags are added on egress
  - The FC->dot1p, DE bit entries in the SAP egress QoS policy are applied.
  - If the **qinq-mark-top-only** command under **vpls>sap>egress** is not enabled (default), the policy is applied to both service delimiting tags.
  - If the qinq-mark-top-only command is enabled, the policy is applied only to the outer service delimiting tag.
  - On the tags where the egress QoS policies do not apply the values of the dot1p+DE bits from the stripped QTAG are used.

- **Egress PE, egress I-SDP case** with force-qtag-forwarding enabled under PBB Epipe or IVPLS
  - **Case 1**: I-SDP vc-type = Ethernet VLAN so there is service delimiting tag added after PW encapsulation.
    - The dot1p+DE bits from the QTAG received over the PBB core side are copied to the QTAG added on the I-SDP.
    - The VLAN value in the QTAG might change to match the provisioned value for the I-SDP configuration.
  - **Case 2**: I-SDP vc-type = Ethernet (force-vlan-vc-forwarding=not supported for I-SDPs) so there is no service delimiting tag added on egress PW
    - The QTAG received over the PBB core is stripped and the QoS information is lost.

4.2.24 Egress B-SAP per ISID Shaping

This feature allows users to perform egress data path shaping of packets forwarded within a B-VPLS SAP. The shaping is performed within a more granular context within the SAP. The context for a B-SAP is an ISID.

4.2.24.1 B-SAP Egress ISID Shaping Configuration

Users can enable the per-ISID shaping on the egress context of a B-VPLS SAP by configuring an encapsulation group, referred to as **encap-group** in CLI, under the QoS sub-context, referred to as **encap-defined-qos**.

config>service>vpls>sap>egress>encap-defined-qos>encap-group *group-name* [**type** *group-type*] [**qos-per-member**] [**create**]

The group name is unique across all member types. The **isid** type is currently the only option.

The user adds or removes members to the **encap-group**, one at a time or as a range of contiguous values. However, when the **qos-per-member** option is enabled, members must be added or removed one at a time. These members are also referred to as ISID contexts.

config>service> vpls>sap>egress>encap-defined-qos>encap-group

[**no**] **member** *encap-id* [**to** *encap-id*]

The user can configure one or more encap-groups in the egress context of the same B-SAP, defining different ISID values and applying each a different SAP egress QoS policy, and optionally a different scheduler policy/agg-rate-limit. ISID values are unique within the context of a B-SAP. The same ISID value cannot be re-used in another encap-group under the same B-SAP but can be re-used in an encap-group under a different B-SAP. Finally, if the user adds to an encap-group an ISID value which is already a member of this encap-group, the command causes no effect. The same if the user attempts to remove an ISID value which is not a member of this encap-group.

When a group is created, the user assigns a SAP egress QoS policy, and optionally a scheduler policy or aggregate rate limit, using the following commands:

config>service> vpls>sap>egress>encap-defined-qos>encap-group>qos *sap-egress-policy-id*

config>service> vpls>sap>egress>encap-defined-qos>encap-group>scheduler-policy *scheduler-policy-name*

config>service> vpls>sap>egress>encap-defined-qos>encap-group>agg-rate-limit *kilobits-per-second*

A SAP egress QoS policy must first be assigned to the created encap-group before the user can add members to this group. Conversely, the user cannot perform the **no qos** command until all members are deleted from the **encap-group**.

An explicit or the default SAP egress QoS policy will continue to be applied to the entire B-SAP but this will serve to create the set of egress queues which will be used to store and forward a packet which does not match any of the defined ISID values in any of the encap-groups for this SAP.

Only the queue definition and fc-to-queue mapping from the encap-group SAP egress QoS policy is applied to the ISID members. All other parameters configurable in a SAP egress QoS policy must be inherited from egress QoS policy applied to the B-SAP.

Furthermore, any other CLI option configured in the egress context of the B-SAP will continue to apply to packets matching a member of any encap-group defined in this B-SAP.

Note also that the SAP egress QoS policy must not contain an active policer or an active queue-group queue or the application of the policy to the encap-group will be failed. A policer or a queue-group queue is referred to as active if one or more FC map to it in the QoS policy or the policer is referenced within the action statement of an IP or IPv6 criteria statement. Conversely, the user will not be allowed to assign a FC to a policer or a queue-group queue, or reference a policer within the action statement of an IP or IPv6 criteria statement, once the QoS policy is applied to an encap-group.

The **qos-per-member** keyword allows the user to specify that a separate queue set instance and scheduler/agg-rate-limit instance will be created for each ISID value in the encap-group. By default, shared instances will be created for the entire encap-group.

When the B-SAP is configured on a LAG port, the ISID queue instances defined by all the encap-groups applied to the egress context of the SAP will be replicated on each member link of the LAG. The set of scheduler/agg-rate-limit instances will be replicated per link or per IOM or XMA depending if the adapt-qos option is set to link/port-fair mode or distribute mode. This is the same behavior as that applied to the entire B-SAP in the current implementation.

4.2.24.2 Provisioning Model

The main objective of this proposed provisioning model is to separate the definition of the QoS attributes from the definition of the membership of an encap-group. The user can apply the same SAP egress QoS policy to a large number of ISID members without having to configure the QoS attributes for each member.

The following are conditions of the provisioning model:

- A SAP egress policy ID must be assigned to an **encap-group** before any member can be added regardless of the setting of the **qos-per-member** option.
- When **qos-per-member** is specified in the **encap-group** creation, the user must add or remove ISID members one at a time. The command is failed if a range is entered.
- When **qos-per-member** is specified in the **encap-**group creation, the sap-egress QoS policy ID and the scheduler policy name cannot be changed unless the group membership is empty. However, the **agg-rate-limit** parameter value can be changed or the command removed (**no agg-rate-limit**).
- When **qos-per-member** is not specified in the **encap-group** creation, the user may add or remove ISID members as a singleton or as a range of contiguous values.
- When **qos-per-member** is not specified in the **encap-group** creation, the sap-egress QoS policy ID and the scheduler policy name or **agg-rate-limit** parameter value may be changed at anytime. Note however that the user cannot still remove the SAP egress QoS policy (**no qos**) while there are members defined in the encap-group.
- The QoS policy or the scheduler policy itself may be edited and modified while members are associated with the policy.
- There will be a maximum number of ISID members allowed in the lifetime of an encap-group.

Operationally, the provisioning consists of the following steps:

1. Create an encap-group.
2. Define and assign a SAP egress QoS policy to the encap-group. This step is mandatory else the user is allowed to add members to the **encap-group**.
3. Manage membership for the encap-group using the **member** command (or SNMP equivalent).
   - Supports both range and singleton ISIDs
   - Cannot add an ISID if it already exists on the SAP in another encap-group
     - The **member** command is all-or-nothing. No ISID in a range is added if one fails

- It the first ISID that fails in the error message is identified.
- Must first remove the ISID using **no member** command.

– Specifying an ISID in a group that already exists within the group is a no-op (no failure)

– If insufficient queues or scheduler policies or FC-to-Queue lookup table space exist to support a new member or a modified membership range, the entire member command is failed

4. Define and assign a scheduling policy or agg-rate-limit for the encap-group. This step is optional.

Logically, the encap-group membership operation can be viewed as three distinct functions:

1. Creation or deletion of new queue sets and optionally scheduler/agg-rate-limit at QoS policy association time.
2. Mapping or un-mapping the member ISID to either the group queue set and scheduler (group QoS) or the ISID specific queue set and scheduler (**qos-per-member**).
3. Modifying the groups objective membership based on newly created or expanded ranges or singletons based on the membership operation.

4.2.24.3 Egress Queue Scheduling

Figure 136 **Egress Queue Scheduling**

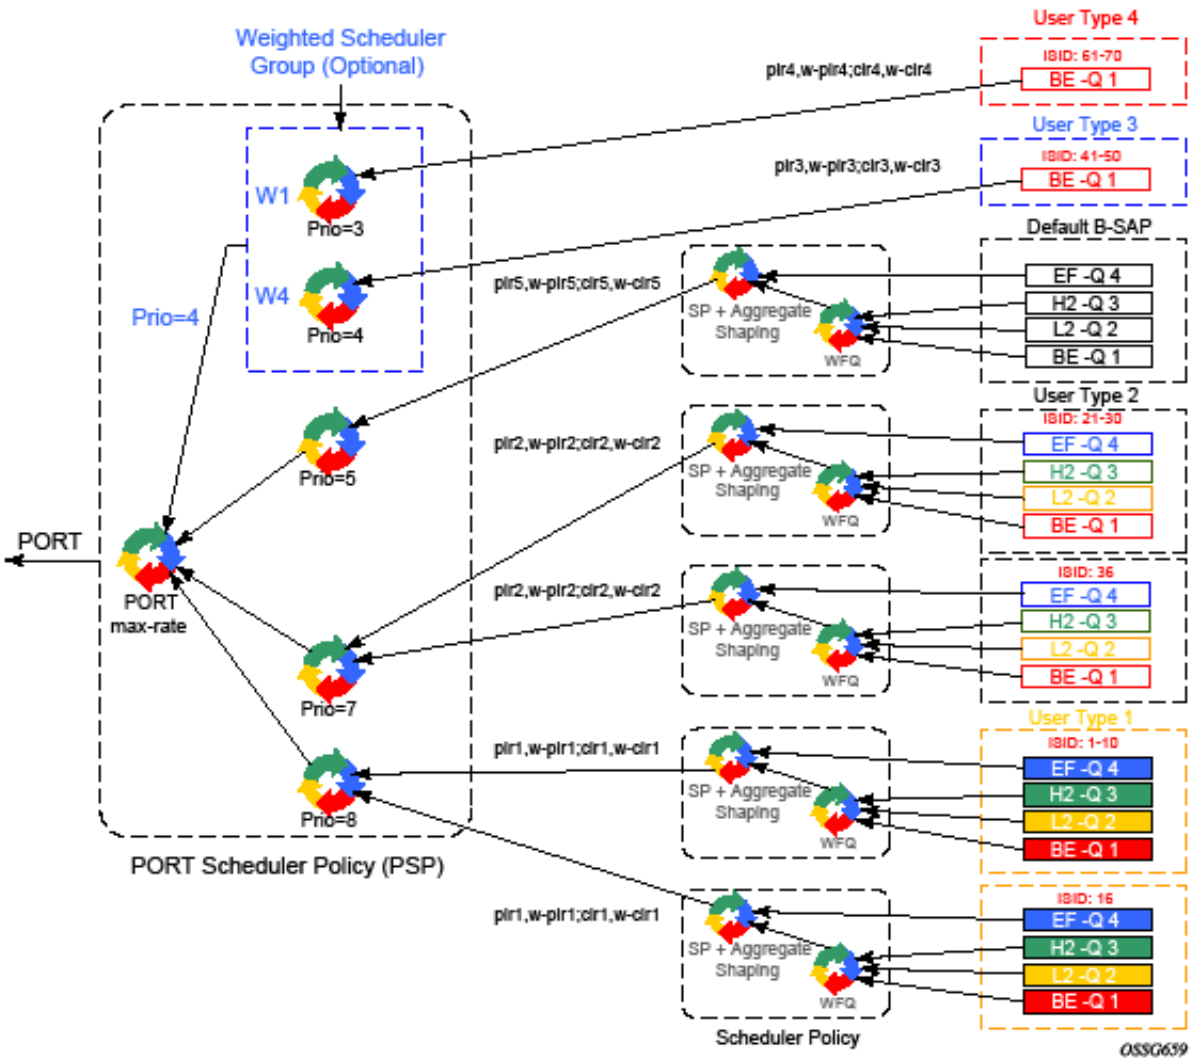

Figure 136 displays an example of egress queue scheduling.

The queuing and scheduling re-uses existing scheduler policies and port scheduler policy with the difference that a separate set of FC queues are created for each defined ISID context according to the encap-group configured under the egress context of the B-SAP. This is in addition to the set of queues defined in the SAP egress QoS policy applied to the egress of the entire SAP.

The user type in Figure 136 maps to a specific encap-group defined for the B-SAP in CLI. The operator has the flexibility of scheduling many user types by assigning different scheduling parameters as follows:

- A specific scheduler policy to each encap-group with a root scheduler which shapes the aggregate rate of all queues in the ISID context of the encap-group and provides strict priority scheduling to its children.

  A second tier scheduler can be used as a WFQ scheduler to aggregate a subset of the ISID context FC queues. Alternatively, the operator can apply an aggregate rate limit to the ISID context instead of a scheduler policy.
- A specific priority level when parenting the ISID queues or the root of the scheduler policy serving the ISID queues to the port scheduler.
- Ability to use the weighted scheduler group to further distribute the bandwidth to the queues or root schedulers within the same priority level according to configured weights.

To make the shaping of the ISID context reflect the SLA associated with each user type, it is required to subtract the operator's PBB overhead from the Ethernet frame size. For that purpose, a **packet byte-offset** parameter is added to the context of a queue.

config>qos>sap-egress>queue>packet-byte-offset {**add bytes** | **subtract bytes**}

When a packet-byte-offset value is applied to a queue instance, it adjusts the immediate packet size. This means that the queue rates, like the operational PIR and CIR, and queue bucket updates use the adjusted packet size. In addition, the queue statistics will also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size.

The port scheduler **max-rate** and **priority level** rates and weights, if a Weighted Scheduler Group is used, are always "on-the-wire" rates and thus use the actual frame size. The same applies to the agg-rate-limit on a SAP, a subscriber, or a Multi-Service Site (MSS) when the queue is port-parented.

When the user enables **frame-based-accounting** in a scheduler policy or **queue-frame-based**-accounting with agg-rate-limit in a port scheduler policy, the queue rate is capped to a user- configured "on-the-wire" rate and the **packet-byte-offset** is not included; however, the packet-byte-offset is applied to the statistics..

4.2.24.4 B-SAP per-ISID Shaping Configuration Example

The following CLI configuration for B-SAP per-ISID shaping achieves the specific use case shown in Egress Queue Scheduling.

```
config
 qos
       port-scheduler-policy "bvpls-backbone-port-scheduler"
       group scheduler-group1 create
       rate 1000
       level 3 rate 1000 group scheduler-group1 weight w1
       level 4 rate 1000 group scheduler-group1 weight w4
       level 5 rate 1000 cir-rate 100
       level 7 rate 5000 cir-rate 5000
       level 8 rate 500 cir-rate 500
exit

       scheduler-policy "user-type1"
       tier 1
       scheduler root
port-parent level 8 rate pir1 weight w-pir1 cir-level 8 cir-rate cir1
cir-weight w-cir1
            exit
       tier 3
       scheduler wfq
           rate pir1
       parent root
            exit
        exit
exit

       scheduler-policy "user-type2"
       tier 1
       scheduler root
port-parent level 7 rate pir2 weight w-pir2 cir-level 7 cir-rate cir2
cir-weight w-cir2
            exit
       tier 3
       scheduler wfq
           rate pir2
       parent root
            exit
        exit
exit

       scheduler-policy "b-sap"
       tier 1
       scheduler root
port-parent level 5 rate pir5 weight w-pir5 cir-level 1 cir-rate cir5 cir-weight
w-cir5
            exit
       tier 3
       scheduler wfq
           rate pir5
       parent root
            exit
        exit
```

```
exit

    sap-egress 100 // user type 1 QoS policy
    queue 1
                parent wfq weight x level 3 cir-weight x cir-level 3
            packet-byte-offset subtract bytes 22
 queue 2
            packet-byte-offset subtract bytes 22
                parent wfq weight y level 3 cir-weight y cir-level 3
 queue 3
            packet-byte-offset subtract bytes 22
                parent wfq weight z level 3 cir-weight z cir-level 3
 queue 4
                parent root level 8 cir-level 8
            packet-byte-offset subtract bytes 22
 fc be queue 1
 fc l2 queue 2
 fc h2 queue 3
 fc ef queue 4
exit

       sap-egress 200 // user type 2 QoS policy
 queue 1
                parent wfq weight x level 3 cir-weight x cir-level 3
            packet-byte-offset subtract bytes 26
 queue 2
                parent wfq weight y level 3 cir-weight y cir-level 3
            packet-byte-offset subtract bytes 26
queue 3
                parent wfq weight z level 3 cir-weight z cir-level 3
            packet-byte-offset subtract bytes 26
queue 4
                parent root level 8 cir-level 8
            packet-byte-offset subtract bytes 26
 fc be queue 1
 fc l2 queue 2
 fc h2 queue 3
 fc ef queue 4
exit

       sap-egress 300 // User type 3 QoS policy
 queue 1
                port-parent level 4 rate pir3 weight w-pir3 cir-level
       4 cir-rate cir3 cir-weight w-cir3
            packet-byte-offset subtract bytes 22
 fc be queue 1
exit

       sap-egress 400 // User type 4 QoS policy
 queue 1
                port-parent level 3 rate pir4 weight w-pir4 cir-level
       3 cir-rate cir4 cir-weight w-cir4
            packet-byte-offset subtract bytes 22
 fc be queue 1
exit

       sap-egress 500 // B-SAP default QoS policy
queue 1
                parent wfq weight x level 3 cir-weight x cir-level 3
```

```
queue 2
                parent wfq weight y level 3 cir-weight y cir-level 3
 queue 3
                parent wfq weight z level 3 cir-weight z cir-level 3
 queue 4
                parent root level 8 cir-level 8
 fc be queue 1
 fc l2 queue 2
 fc h2 queue 3
 fc ef queue 4
exit
 exit
exit

config
       service
       vpls 100 bvpls
            sap 1/1/1:100
                egress
                    encap-defined-qos
                        encap-group type1-grouped type isid
       member 1 to 10
                                qos 100
       scheduler-policy user-type1
                        exit
encap-group type1-separate type isid qos-per-member
       member 16
                                qos 100
       scheduler-policy user-type1
       exit
       encap-group type2-grouped type isid
       member 21 to 30
                                qos 200
       scheduler-policy user-type2
                        exit
encap-group type2-separate type isid qos-per-member
       member 36
                                qos 200
       scheduler-policy user-type2
       exit
       encap-group type3-grouped type isid
       member 41 to 50
                                qos 300
                        exit
                           encap-group type4-grouped type isid
       member 61 to 70
                                qos 400
       exit
                    qos 500
       scheduler-policy b-sap
       exit
            exit
        exit
    exit
exit
```

4.2.25 PBB OAM

The Nokia PBB implementation supports both MPLS and native Ethernet tunneling. In the case of an MPLS, SDP bindings are used as the B-VPLS infrastructure while T-LDP is used for signaling. As a result, the existing VPLS, MPLS diagnostic tools are supported in both I-VPLS and B-VPLS domains as depicted in Figure 137.

Figure 137 **PBB OAM View for MPLS Infrastructure**

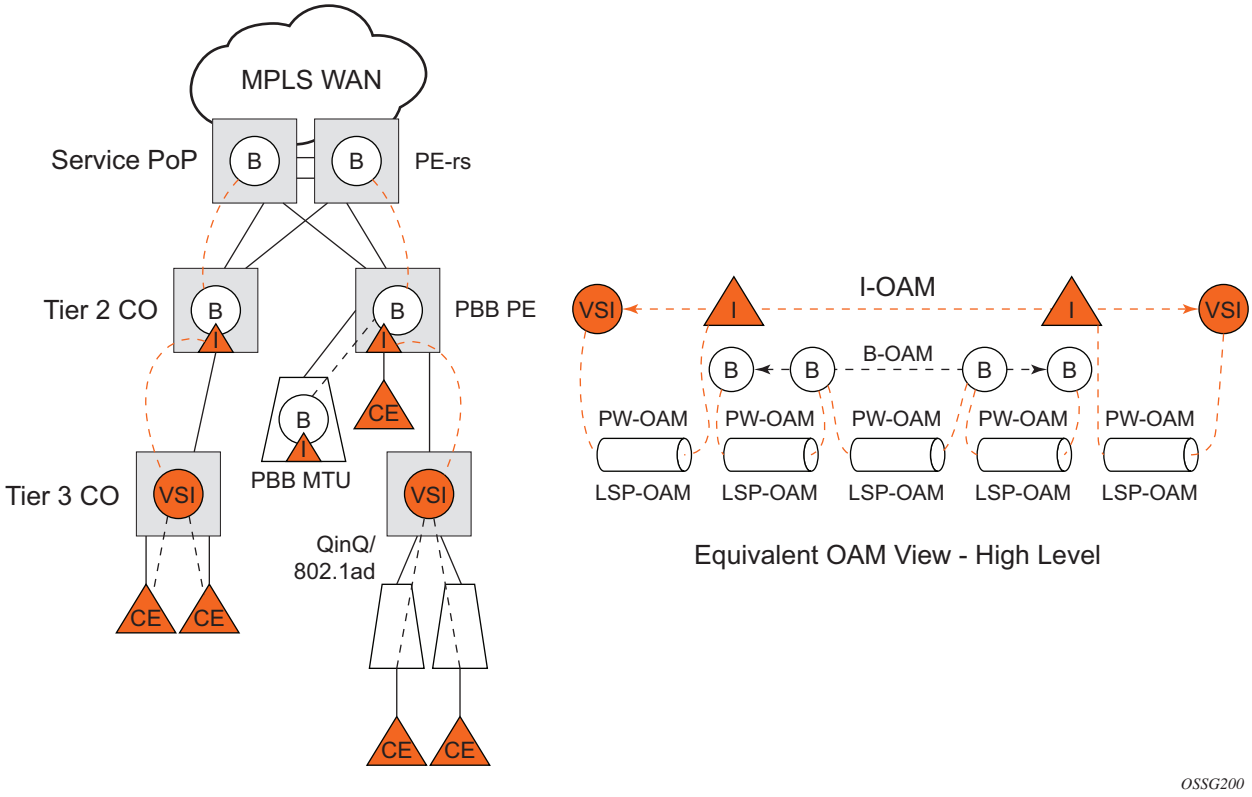

When an Ethernet switching backbone is used for aggregation between PBB PEs, a SAP is used as the B-VPLS uplink instead of an SDP. No T-LDP signaling is available.

The existing IEEE 802.1ag implemented for regular VPLS SAPs may be used to troubleshoot connectivity at I-VPLS and B-VPLS layers.

4.2.25.1 Mirroring

There are no restrictions for mirroring in I-VPLS or B-VPLS.

4.2.25.2 OAM Commands

All VPLS OAM commands may be used in both I-VPLS and B-VPLS instances.

I-VPLS

- The following OAM commands are meaningful only toward another I-VPLS service instance (spoke-SDP in I-VPLS):
  - LSP-ping, LSP-trace, SDP-ping, SDP-MTU
- The following I-VPLS OAM exchanges are transparently transported over the B-VPLS core:
  - SVC-ping, MAC-ping, MAC-trace, MAC-populate, MAC-purge, CPE-ping (toward customer CPE), 802.3ah EFM, SAA
- PBB uplinks using MPLS/SPP: there are no PBB specific OAM commands.

B-VPLS

- In case of Ethernet switching backbone (B-SAPs on B-VPLS), 802.1ag OAM is supported on B-SAP, operating on:
  - The customer level (C-SA/C-DA and C-type layer)
  - The tunnel level (B-SA/B-DA and B-type layer)

4.2.25.3 CFM Support

There is no special 802.1ag CFM (Connectivity Fault Management) support for PBB. B-component and I-components run their own maintenance domain and levels. CFM for I-components run transparently over the PBB network and will appear as directly connected.

4.3 Configuration Examples

Use the CLI syntax displayed to configure PBB.

4.3.1 PBB using G.8031 Protected Ethernet Tunnels

The following displays PBB configuration examples:

Ethernet links on BEB1:
BEB1 to BEB1 L1:
BEB1 to BCB1 L1: 1/1/1 – Member port of LAG-emulation ET1, terminate ET3
BEB1 to BCB1 L2: 2/1/1 – Member port of LAG-emulation ET1
BEB1 to BCB1 L3: 3/1/1 - Member port of LAG-emulation ET1
BEB1 to BCB2: 4/1/1 – terminate ET3

```
*A:7750_ALU>config>eth-tunnel 1
        description "LAG-emulation to BCB1 ET1"
        protection-type loadsharing
        ethernet
            mac 00:11:11:11:11:12
            encap-type dot1q
        exit
        ccm-hold-time down 5 up 10 // 50 ms down, 1 sec up
        lag-emulation
           access adapt-qos distribute
           path-threshold 1
        exit
        path 1
            member 1/1/1
            control-tag 0
            eth-cfm
                …
            exit
            no shutdown
        exit
        path 2
            member 2/1/1
            control-tag 0
            eth-cfm
                …
            exit
            no shutdown
        exit
        path 3
            member 3/1/1
            control-tag 0
            eth-cfm
                …
```

```
            exit
            no shutdown
        exit
        no shutdown
-------------------------------------------------
*A:7750_ALU>config>eth-tunnel 3
        description "G.8031 tunnel ET3"
        protection-type 8031_1to1
        ethernet
            mac 00:11:11:11:11:11
            encap-type dot1q
        exit
        ccm-hold-time down 5 // 50 ms down, no up hold-down
        path 1
            member 1/1/1
            control-tag 5
            precedence primary
            eth-cfm
                mep 2 domain 1 association 1
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        path 2
            member 4/1/1
            control-tag 5
            eth-cfm
                mep 2 domain 1 association 2
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        no shutdown
-------------------------------------------------
# Service config
-------------------------------------------------
*A:7750_ALU>config>service vpls 1 customer 1 m-vpls b-vpls create
   description "m-VPLS for multipoint traffic"
   stp
      mst-name  "BVPLS"
      mode p-mstp
      mst-instance 10
         mst-priority 4096
         vlan-range 100-199
      exit
      mst-instance 20
         mst-priority 8192
         vlan-range 200-299
      exit
      no shutdown
   exit

   sap eth-tunnel-1 create // BSAP0 to BCB E
```

```
   sap 4/1/1:0 create // physical link to BCB F (NOTE 0 or 0.*)
                      // indicate untagged for m-VPLS)
   exit
   no shutdown
-------------------------------------------------
# Service config: one of the same-fate SAP over
# loadsharing tunnel
-------------------------------------------------
A:7750_ALU>config service vpls 100 customer 1 b-vpls create
   sap eth-tunnel-1:1 create //to BCB E
      // must specify tags for each path for loadsharing
      eth-tunnel
path 1 tag 100
            path 2 tag 100
            path 3 tag 100
   exit
   no shutdown   …
   sap 3/1/1:200 // to BCBF
   …

A:7750_ALU>config service vpls 1000 customer 1 i-vpls create
   pbb backbone-vpls 100 isid 1000
   sap 4/1/1:200 // access SAP to QinQ
…
-------------------------------------------------
# Service config: one of epipes into b-VPLS protected tunnel
# as per R7.0 R4
-------------------------------------------------
A:7750_ALU>config service service vpls 3 customer 1 b-vpls create
   sap eth-tunnel-3 create
   …
service epipe 2000
    pbb-tunnel 100 backbone-dest-mac to-AS20 isid 2000
    sap 3/1/1:400 create
```

Example:

```
port 1/1/1
     ethernet
          encap-type dot1q
port 2/2/2
     ethernet
          encap-type dot1q
config eth-tunnel 1
     path 1
          member 1/1/1
          control-tag 100
          precedence primary
          eth-cfm
               mep 51 domain 1 association 1 direction
                  down
               ccm-enable
               low-priority-defect allDef
               mac-address 00:AE:AE:AE:AE:AE
               control-mep
               no shutdown
          no shutdown
```

```
        path 2
            member 2/2/2
            control-tag 200
            eth-cfm
                mep
                    mep 52 domain 1 association 2
                        direction down
                    ccm-enable
                    low-priority-defect allDef
                    mac-address 00:BE:BE:BE:BE:BE
                    control-mep
                    no shutdown
            no shutdown

    config service vpls 1 b-vpls
        sap eth-tunnel-1
    config service epipe 1000
        pbb-tunnel 1 backbone-dest-mac remote-beb
        sap 3/1/1:400.10
```

4.3.2 MC-LAG Multihoming for Native PBB

This section describes a configuration example for BEB C configuration given the following assumptions:

- BEB C and BEB D are MC-LAG peers
- B-VPLS 100 on BEB C and BEB D
- VPLS 1000 on BEB C and BEB D
- MC-LAG 1 on BEB C and BEB D

CLI Syntax:

```
service pbb
    source-bmac ab-ac-ad-ef-00-00
port 1/1/1
    ethernet
        encap-type qinq
lag 1
    port 1/1/1 priority 20
    lacp active administrative-key 32768
redundancy
    multi-chassis
        peer 1.1.1.3 create
            source-address 1.1.1.1
            mc-lag
                lag 1 lacp-key 1 system-id
                    00:00:00:01:01:01
```

```
                        system-priority 100
                        source-bmac-lsb use-lacp-key
            service vpls 100 bvpls
                 sap 2/2/2:100 // bvid 100
                 mac-notification
                      no shutdown

            service vpls 101 bvpls
                 sap 2/2/2:101 // bvid 101
                 mac-notification
                      no shutdown
            // no per BVPLS source-bmac configuration, the chassis
             one (ab-ac-ad-ef-00-00) is used

            service vpls 1000 ivpls
                 backbone-vpls 100
                 sap lag-1:1000 //automatically associates the SAP
                   with ab-ac-ad-ef-00-01 (first 36 bits from BVPLS
                   100 sbmac+16bit source-bmac-lsb)

            service vpls 1001 ivpls
                 backbone-vpls 101
                 sap lag-1:1001 //automatically associates the SAP
                   with ab-ac-ad-ef-00-01(first 36 bits from BVPLS
                   101 sbmac+16bit source-bmac-lsb)
```

4.3.3 Access Multi-Homing over MPLS for PBB Epipes

This section gives an example configuration for BEB1 from Figure 134.

```
*A:BEB1>config>service# info
----------------------------------------------
        pbb
            source-bmac 00:00:00:00:11:11
            mac-name "remote-BEB" 00:44:44:44:44:44
        exit
        sdp 1 mpls create
            far-end 1.1.1.4
            ldp
            keep-alive
                shutdown
            exit
            source-bmac-lsb 33:33 control-pw-vc-id 100
            no shutdown
        exit
        vpls 10 customer 1 b-vpls create
            service-mtu 1532
            stp
                shutdown
            exit
```

```
            spb 1024 fid 1 create
                no shutdown
            exit
            sap 1/1/1:10 create
                spb create
                    no shutdown
                exit
            exit
            sap 1/1/5:10 create
                spb create
                    no shutdown
                exit
            exit
            no shutdown
        exit
        epipe 100 customer 1 create
            pbb
                tunnel 10 backbone-dest-mac "remote-BEB" isid 100
            exit
            spoke-sdp 1:100 create
                use-sdp-bmac
                no shutdown
            exit
            no shutdown
        exit
        epipe 101 customer 1 create
            pbb
                tunnel 10 backbone-dest-mac "remote-BEB" isid 101
            exit
            spoke-sdp 1:101 create
                use-sdp-bmac
                no shutdown
            exit
            no shutdown
        exit
-------------------------------------------
*A:BEB1>config>service#
```

The SDP control pseudowire information can be seen using this command:

```
*A:BEB1# show service sdp 1 detail

===============================================================================
Service Destination Point (Sdp Id : 1) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1  -1.1.1.4
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id                : 1                     SDP Source         : manual
...
Src B-MAC LSB         : 33-33                 Ctrl PW VC ID      : 100
Ctrl PW Active        : Yes

...
===============================================================================
*A:BEB1#
```

The configuration of a pseudowire to support remote active/standby PBB Epipe operation can be seen using this command:

```
*A:BEB1# show service id 100 sdp 1:100 detail

===============================================================================
Service Destination Point (Sdp Id : 1:100) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:100  -(1.1.1.4)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id              : 1:100                     Type              : Spoke
...
Use SDP B-MAC       : True
...
===============================================================================
*A:BEB1#8.C
```

4.4 PBB Configuration Command Reference

This chapter describes the PBB configuration command reference.

4.4.1 Command Hierarchies

4.4.1.1 Global Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls]
          [etree] [create]
          — [no] spb instance [fid value] [create]
              — level level-number
                  — bridge-priority value
                  — ect-algorithm name fid-range fid-range
                  — forwarding-tree-topology[st|spf]
              — lsp-lifetime seconds
              — no lsp-lifetime
              — overload [timeout seconds]
              — no overload
              — overload-on-boot [timeout seconds]
              — no overload-on-boot
              — [no] shutdown
              — timers
                  — lsp-wait lsp-wait [lsp-initial-wait lsp-initial-wait] [lsp-second-
                    wait lsp-second-wait]
                  — no lsp-wait
                  — spf-wait spf-wait [spf-initial-wait spf-initial-wait] [spf-second-
                    wait spf-second-wait]
                  — no spf-wait
          — spbm-control-vpls service-id fid fid
          — no spbm-control-vpls
          — mrp
              — [no] attribute-table-high-wmark high-water-mark
              — [no] attribute-table-low-wmark low-water-mark
              — [no] attribute-table-size max-attributes
              — flood-time flood-time
              — no flood-time
              — [no] shutdown

config
    — service
        — pbb
            — mac-name name ieee-address
            — no mac-name
```

```
                    — source-bmac ieee-address
                    — no source-bmac

config
    — service
        — [no] vpls service-id [customer customer-id] [i-vpls] [b-vpls] [[create]
            — pbb
                — backbone-vpls service-id [isid isid]
                — no backbone-vpls
                    — igmp-snooping
                        — [no] mrouter-dest mac-name
                    — mld-snooping
                        — [no] mrouter-dest mac-name
                    — [no] sap sap-id
                        — igmp-snooping
                            — [no] mrouter-port
                        — mld-snooping
                            — [no] mrouter-port
                    — [no] sdp sdp-id:vc-id
                        — igmp-snooping
                            — [no] mrouter-port
                        — mld-snooping
                            — [no] mrouter-port
                    — [no] stp
                — [no] force-qtag-forwarding
                — mac-notification
                    — [no] count value
                    — [no] interval value
                    — renotify value
                    — no renotify
                — [no] propagate-mac-flush-from-bvpls
                — [no] send-bvpls-flush {[all-from-me] | [all-but-mine]}
                — [no] send-flush-on-bvpls-failure
                — source-bmac ieee-address
                — no source-bmac
                — [no] use-sap-bmac
```

4.4.1.2 SAP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [mvpls] [b-vpls | i-vpls]
          [create]
            — sap sap-id [split-horizon-group group-name] [create] [capture-sap]
            — no sap sap-id
                — mrp
                    — [no] join-time value
                    — [no] leave-all-time value
                    — [no] leave-time value
                    — [no] mrp-policy policy-name
```

```
                    — [no] periodic-time value
                    — [no] periodic-timer
                — [no] spb create
                    — [no] shutdown
                    — lsp-pacing-interval milli-seconds
                    — no lsp-pacing-interval
                    — retransmit-interval seconds
                    — no retransmit-interval
                    — metric value
                    — no metric
                        — hello-interval seconds
                        — no hello-interval
                        — hello-multiplier multiplier
                        — no hello-multiplier
```

4.4.1.3 Mesh SDP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id]  [vpn vpn-id] [mvpls] [b-vpls | i-vpls]
          [create]
          — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
          — no mesh-sdp sdp-id[:vc-id]
              — mrp
                  — [no] join-time value
                  — [no] leave-all-time value
                  — [no] leave-time value
                  — [no] mrp-policy policy-name
                  — [no] periodic-time value
                  — [no] periodic-timer
```

4.4.1.4 Spoke SDP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id]  [vpn vpn-id] [mvpls] [b-vpls | i-vpls]
          [create]
          — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-
              name]
          — no spoke-sdp sdp-id[:vc-id]
              — mrp
                  — [no] join-time value
                  — [no] leave-all-time value
                  — [no] leave-time value
                  — [no] mrp-policy policy-name
                  — [no] periodic-time value
                  — [no] periodic-timer
```

```
— [no] spb create
    — [no] shutdown
    — lsp-pacing-interval milli-seconds
    — no lsp-pacing-interval
    — retransmit-interval seconds
    — no retransmit-interval
    — metric value
    — no metric
        — hello-interval seconds
        — no hello-interval
        — hello-multiplier multiplier
        — no hello-multiplier
```

4.4.1.5 BGP-MH for I-VPLS Commands

```
config
    — service
        — vpls service-id [customer customer-id]  [vpn vpn-id] [mvpls] [b-vpls | i-vpls] [create]
        — no vpls service-id
            — no site name
                — boot-timer seconds
                — no boot-timer
                — failed-threshold [1..1000]
                — failed-threshold all
                — [no] mesh-sdp-binding
                — monitor-oper-group name
                — no monitor-oper-group
                — sap sap-id
                — no sap
                — [no] shutdown
                — site-activation-timer seconds
                — no site-activation-timer
                — site-min-down-timer min-down-time
                — no site-min-down-timer
                — site-id value
                — no site-id
                — split-horizon-group group-name
                — no split-horizon-group
                — spoke-sdp sdp-id:vc-id
                — no spoke-sdp
```

4.4.2 Command Descriptions

4.4.2.1 VPLS Service Commands

vpls

Syntax **vpls** *service-id* **customer** *customer-id* **vpn** *vpn-id* [**m-vpls**] [**b-vpls** | **i-vpls**] [**create**]
vpls *service-id*
no vpls *service-id*

Context config>service

Description This command creates or edits a Virtual Private LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. When a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

When a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *service-id* — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every SR OS on which this service is defined.

Values 1 to 2147483648

customer *customer-id* — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn *vpn-id* — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default null (0)

m-vpls — Specifies a management VPLS.

b-vpls | i-vpls — Creates a backbone-vpls or ISID-vpls for use with PBB.

site

Syntax **site** *name* [**create**]
no site *name*

Context config>service>vpls

Description This command configures a VPLS site.

The no form of the command removes the name from the configuration.

Parameters *name* — Specifies a site name up to 32 characters in length.

create — This keyword is mandatory while creating a VPLS service.

boot-timer

Syntax **boot-timer** *seconds*
no boot-timer

Context config>service>vpls>site

Description This command configures for how long the service manger waits after a node reboot before running the DF election algorithm. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged.

The **no** form of the command reverts the default.

Default 10

Parameters *seconds* — Specifies the site boot-timer in seconds.

Values 0 to 100

failed-threshold

Syntax **failed-threshold** [1..1000]
failed-threshold all

Context config>service>vpls>site

Description This command defines the number of objects should be down for the site to be declared down. Both administrative and operational status must be evaluated and if at least one is down, the related object is declared down.

Default failed-threshold all

Parameters *1 . 1000* — Specifies the threshold for the site to be declared down.

mesh-sdp-binding

Syntax [**no**] **mesh-sdp-binding**

Context config>service>vpls>site

Description This command enables applications to all mesh SDPs.

The **no** form of reverts the default.

Default no mesh-sdp-binding

monitor-oper-group

Syntax **monitor-oper-group** *group-name*
no monitor-oper-group

Context config>service>vpls>site

Description This command specifies the operational group to be monitored by the object under which it is configured. The **oper-group** *name* must be already configured under the **config>service** context before its name is referenced in this command.

The **no** form of the command removes the association.

sap

Syntax **sap** *sap-id*
no sap

Context config>service>vpls>site

Description This command configures a SAP for the site.

The **no** form of the command removes the SAP ID from the configuration.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

site-activation-timer

Syntax **site-activation-timer** *seconds*
no site-activation-timer

Context config>service>vpls>site

Description This command configures the time-period the system keeps the local sites in standby status, waiting for BGP updates from remote PEs before running the DF (designated-forwarder) election algorithm to decide whether the site should be unblocked. This timer if terminated if an update is received for which the remote PE has transitioned from DF to non-DF.

The no form of the command removes the value from the configuration.

Default 2

Parameters *seconds* — Specifies the site activation timer in seconds.

Values 0 to 100

site-min-down-timer

Syntax **site-min-down-timer** *min-down-time*
no site-min-down-timer

Context config>service>vpls>site

Description This command configures the BGP multi-homing site minimum down time. When set to a non-zero value, if the site goes operationally down it will remain operationally down for at least the length of time configured for the **site-min-down-timer**, regardless of whether other state changes would have caused it to go operationally up. This timer is restarted every time that the site transitions from up to down. Setting this parameter to zero allows the minimum down timer to be disabled for this service.

The above operation is optimized in the following circumstances:

- If the site goes down on the designated forwarder but there are no BGP multi-homing peers with the same site in an operationally up state, then the **site-min-down-timer** is not started and is not used.
- If the site goes down on the designated forwarder but there are no active BGP multi-homing peers, then the **site-min-down-timer** is not started and is not used.

- If the s**ite-min-down-timer** is active and a BGP multi-homing update is received from the designated forwarder indicating its site has gone down, the **site-min-down-timer** is immediately terminated and this PE becomes the designated forwarder if the BGP multi-homing algorithm determines it should be the designated forwarder.

The **no** form of this command reverts to the default value.

Default Taken from the value of **site-min-down-timer** configured for Multi-Chassis BGP multi-homing under the **config>redundancy>bgp-multi-homing** context.

Parameters *min-down-time* — Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down.

Values 0 to 100 seconds

site-id

Syntax **site-id** *value*
no site-id

Context config>service>vpls>site

Description This command configures the identifier for the site in this service.

Parameters *value* — Specifies the site identifier.

Values 1 to 65535

split-horizon-group

Syntax **split-horizon-group** *group-name*
no split-horizon-group

Context config>service>vpls>site

Description This command configures the value of split horizon group associated with this site.

The **no** form of the command reverts the default.

Default no split-horizon-group

Parameters *group-name* — Specifies a split horizon group name.

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id*
no spoke-sdp

Context config>service>vpls>site

Description This command binds a service to an existing Service Distribution Point (SDP).

The **no** form of the command removes the parameter from the configuration.

service-name

Syntax **service-name** *service-name*
no service-name

Context config>service>vpls

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a specified service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service when it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0 to 9).

eth-tunnel

Syntax **eth-tunnel** *tunnel-id*

Context config>service>vpls

Description This command associates a BVPLS SAP with the global Ethernet tunnel object specified by tunnel-id. Only one-to-one mapping between SAP and Ethernet tunnel is supported in the initial implementation. The global eth-tunnel tunnel-id with at least a member port must be configured in advance for the command to be successful. A SAP will be instantiated using the active path components (member port and control-tag) for VPLS forwarding. The last member port in the Ethernet tunnel cannot be deleted if there is a SAP configured on that eth-tunnel. This command is only available in the BVPLS context.

The **no** form of this command removes the sap from the Ethernet tunnel object.

Default no sap is specified

Parameters *tunnel-id* — Specifies the value of the Ethernet tunnel identifier to be used for the SAP.

Values 1 to 64

mesh-sdp

Syntax **mesh-sdp sdp-id**[**:vc-id**] [**vc-type** {**ether** | **vlan**}] [**root-leaf-tag** | **leaf-ac**]
no mesh-sdp sdp-id[**:vc-id**]

Context config>service>vpls

Description This command binds a VPLS service to an existing Service Distribution Point (SDP). Mesh SDPs bound to a service are logically treated like a single bridge “port” for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke-SDPs and SAPs) and not transmitted on any mesh SDPs.

This command creates a binding between a service and an SDP. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate the SDP with a valid service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. After it is removed, no packets are forwarded to the far-end router.

Default none

spoke-sdp

Syntax **spoke-sdp sdp-id**[**:vc-id**] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group group-name**] **endpoint** [**no-endpoint**] [**root-leaf-tag** | **leaf-ac**]
no spoke-sdp sdp-id[**:vc-id**]

Context config>service>vpls

Description This command binds a service to an existing Service Distribution Point (SDP). A spoke-SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke-SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the sdp-id does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. After it is removed, no packets are forwarded to the far-end router.

Default none

spb

Syntax [**no**] **spb** *instance* [**fid** *value*] [**create**]

Context config>service>vpls b-vpls
config>service>vpls b-vpls>sap>spb
config>service>vpls b-vpls>spoke-sdp>spb

Description This command enables Shortest Path Bridging (SPB) on a B-VPLS instance. SPB uses IS-IS that supports multiple instances, therefore an instance must be specified. The declaration of SPB in this context is the control configuration for the SPB. This is an SPB management interface and it manages the configuration for IS-IS. Various parameters that define this SPB instance are configured under this SPB instance. Several of the parameters are shared with other B-VPLS service instances using SPB.

SPB enables an instance of IS-IS protocol with the no shutdown command. Alternatively, the IS-IS protocol instance under SPB is disabled with the shutdown command in the **config>service>vpls b-vpls>spb** context.

A Forwarding Identifier (FID) is optionally specified which is an abstraction of the B-VID used for forwarding in SPB. When no FID is configured the control VPLS is advertised with FID value 1. When a FID value is specified, the control VPLS is advertised and associated with the FID value specified. The default algorithm for any FID declared or implicit is low-path-id. When a FID is specified, the ect-algorithm can be specified for the FID and changed only when there are no VPLS, SAPs or SDP bindings associated with the FID. The FID for a control instance cannot be changed after it is created. To change a FID the SPB component would have to be shutdown, deleted and recreated with a new FID.

Note: SPB operates with disable-learning, disable aging and discard-unknown. The state of these commands is ignored when SPB is configured.

Default no spb

Parameters *instance* — Specifies the instance ID for an SPB IS-IS instance.

Values 1024 to 2047 (4 available)

Default 1024

value — Specifies the FID value.

Values 1 to 4095

Default 1

spb

Syntax [**no**] **spb** [**create**]

Context config>service>vpls b-vpls>sap>spb
config>service>vpls b-vpls>spoke-sdp>spb

Description This command enables Shortest Path Bridging (SPB) on SAP or Spoke SDP. The B-VPLS may be a control B-VPLS or user B-VPLS. Since SPB uses IS-IS that supports multiple instances, SPB inherits the instance from the control B-VPLS.

SPB at this context level is enabled immediately. SPB enables an instance of IS-IS protocol with the no shutdown command. Alternatively, the IS-IS protocol instance under SPB is disabled with the shutdown command in the **config>service>vpls b-vpls>spb** context.

Default no spb

spbm-control-vpls

Syntax **spbm-control-vpls** *service-id* **fid** *fid*
no spbm-control-vpls

Context config>service>vpls service-id b-vpls

Description This command associates a user B-VPLS with a particular control B-VPLS and a FID. The ECT algorithm and the behavior of unicast and multicast come from the association to the FID.

A Forwarding Identifier (FID) is specified which is an abstraction of the B-VID used for forwarding in SPB. The ect-algorithm is associated with the FID and can be changed only when there are no VPLS, SAPs or SDP bindings associated with the FID. The FID must be independent from the FID assigned to other services.

Default none

Parameters *service-id* — The B-VPLS service identifier.

Values 1 to 2147483647 | *svc-name*: 64 characters max

fid — The forwarding identifier.

Values 1 to 4095

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls b-vpls>spb
config>service>vpls b-vpls>sap>spb
config>service>vpls b-vpls>spoke-sdp>spb

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within.

The **no** form of this command administratively enables an entity.

SPB Interface — In the config>service>vpls b-vpls>spb> context, the command disables the IS-IS interface. By default, the IS-IS interface is disabled (shutdown).

lsp-lifetime

Syntax **lsp-lifetime** *seconds*
no lsp-lifetime

Context config>service>vpls b-vpls>spb

Description This command sets the time, in seconds, SPB wants the LSPs it originates to be considered valid by other routers in the domain. This is a control B-VPLS command.

Each LSP received is maintained in an LSP database until the lsp-lifetime expires unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from this formula: lsp-lifetime/2.

The **no** form of the command reverts to the default value.

Default 1200 — LSPs originated by SPB should be valid for 1200 seconds (20 minutes).

Parameters *seconds* — The time, in seconds, that SPB wants the LSPs it originates to be considered valid by other routers in the domain.

Values 350 to 65535

timers

Syntax **timers**

Context config>service>vpls b-vpls>spb

Description This command enables the context to configure SPB timers.

lsp-wait

Syntax **lsp-wait** *lsp-wait* [**lsp-initial-wait** *lsp-initial-wait*] [**lsp-second-wait** *lsp-second-wait*]

Context config>service>vpls b-vpls>spb>timers

Description This command is used to customize LSP generation throttling. Timers that determine when to generate the first, second and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

Note: The IS-IS timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Parameters *lsp-wait* — Specifies the maximum interval in milliseconds between two consecutive occurrences of an LSP being generated.

Values 10 to 120000

Default 5000

lsp-initial-wait — Specifies the initial LSP generation delay in milliseconds. Values < 100 ms are internally rounded down to 0, so that there is no added initial LSP generation delay.

Values 10 to 100000

Default 10

lsp-second-wait — Specifies the hold time in milliseconds between the first and second LSP generation.

Values 10 to 100000

Default 1000

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>service>vpls b-vpls>spb

Description This command administratively sets the SPB to operate in the overload state for a specific time period, in seconds, or indefinitely. During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by SPB and will not used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The overload command can be useful in circumstances where SPB is overloaded or used prior to executing a shutdown command to divert traffic around the switch.

The **no** form of the command causes the router to exit the overload state.

Default no overload

Parameters *seconds* — The time, in seconds, that this router must operate in overload state.

Values 60 to 1800

Default Infinity (overload state maintained indefinitely)

overload-on-boot

Syntax **overload-on-boot** [**timeout** *seconds*]
no overload-on-boot

Context config>service>vpls b-vpls>spb>

Description When the router is in an overload state, SPB the B-VPLS is used only if there is no other SPB B-VPLS to reach the destination. This command configures the IGP upon boot up in the overload state until one of the following events occur:

- The timeout timer expires.
- A manual override of the current overload state is entered with the **config**>**service**>**vpls** *instance*>**b-vpls**>**spb**>**no overload** command.

The **no** form of the command does not affect the overload-on-boot function.

If no timeout is specified, SPB IS-IS goes into overload indefinitely after a reboot. After the reboot, the SPB IS-IS status displays a permanent overload state:

```
L1 LSDB Overload : Manual on boot (Indefinitely in overload)
```

This state can be cleared with the **config**>**service**>**vpls** *instance* >**b-vpls**>**spb**>**no overload** command.

When specifying a timeout value, SPB IS-IS goes into overload for the configured timeout after a reboot. After the reboot, SPB IS-IS status displays the remaining time the system stays in overload:

```
L1 LSDB Overload : Manual on boot (Overload Time Left : 17)
```

The overload state can be cleared before the timeout expires with **config**>**service**>**vpls** *instance*>**b-vpls**>**spb**>**no overload** command.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Default no overload-on-boot

Parameters *seconds* — The time, in seconds, that this router must operate in overload state.

Values 60 to 1800

Default Infinity (overload state maintained indefinitely)

spf-wait

Syntax [**no**] **spf-wait** *spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

Context config>service>vpls b-vpls>spb>timers

Description This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second and subsequent SPF calculations after a topology change occurs can be controlled with this command.

Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc., until it reaches the *spf-wait* value. The SPF interval will stay at *spf-wait* value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

Default no spf-wait

Parameters *spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

Values 10 to 120000

Default 10000

spf-initial-wait — Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 to 100000

Default 1000

spf-second-wait — Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 10 to 100000

Default 1000

level

Syntax **level** *level-number*

Context config>service>vpls b-vpls>spb

Description This command creates the context to configure SPB Level 1 or Level 2 area attributes. This is IS-IS levels. Only Level 1 can be configured.

A Level 1 adjacency can be established only with other Level 1 B-VPLS. A Level 2 adjacency can be established only with other Level 2 B-VPLS. Currently there is no support for level 1 and level 2 in the same instance of SPB.

Default level 1

Parameters *level-number* — The SPB level number.

Values 1, 2

bridge-priority

Syntax **bridge-priority** *value*

Context config>service>vpls b-vpls>spb>level level-number

Description This command configures the four bit bridge priority for Shortest Path Bridging. This value is added to the 6 byte bridge Identifier (which is the system-id) in the top four bits of a two byte field. Note the actual value will be bit shifted 12 bits left effective putting this in the high bits of the 16 bits added to system ID.

The bridge priority is important in choosing the Root Bridge for the single tree algorithm (lowest value = best). Bridge priority also factors into the tie breaker for SPF algorithms as described in the SPB standard. The bridge-identifier (system-id) of the control B-VPLS determines the tiebreaker when the bridge-priorities are equal.

Default 8

Parameters *value* — The bridge-priority value.

Values 0 to 15

ect-algorithm

Syntax **ect-algorithm** *name* **fid-range** *fid-range*

Context config>service>vpls b-vpls>spb>level level-number

Description This command configures the ect-algorithm associated with a FID. Names are:

- low-path-id
- high-path-id

The algorithm for low-path-id chooses the path with the lowest metric and uses the sum of each Bridge-ID to break-ties (in this case preferring the lowest bridge identifiers).

The algorithm for high-path-id choose the path with the lowest metric and the sum of each Bridge-ID (after each one is modified by the algorithm mask) to break-ties (in this case preferring the highest bridge identifiers).

A Forwarding Identifier (FID) is an abstraction of the IEEE 802.1 SPB Base VID and represents the VLAN (B-VPLS) in IS-IS LSPs. B-VPLS services with the same FID share B-MACs and I-SIDs. (the SAP encapsulation VLAN tag may be set to the same value as the FID or to any other valid VLAN tag). One or more FIDs can be associated with an ECT-algorithm by using the FID range. User B-VPLS services may share the same FID as the control B-VPLS or use independent FIDs where each FID has an assigned ect-algorithm. B-VPLS services with i-vpls services must have an independent FID. B-VPLS services with only PBB Epipes may share FIDs with other B-VPLS services including the control B-VPLS service.

The ect-algorithm is associated with the FID and can only be changed only when there are no VPLS, SAPs or SDP bindings associated with the FID. The FID must be independent from the FID assigned to other services.

Default low-path-id

Parameters *name* — low-path-id, high-path-id.

fid-range — Range of Forwarding Identifier values.

Values 1 to 4095

forwarding-tree-topology

Syntax **forwarding-tree-topology unicast** [**st** | **spf**]

Context config>service>vpls b-vpls>spb>level level-number

Description This command sets the unicast forwarding to follow the shortest path tree defined by the ECT algorithm shortest path forwarding (spf) or to follow a single tree. (st). Shortest path trees make use of more link resources.

Multicast traffic is defaulted to follow the single tree topology. A single tree unicast would make Multicast and unicast follow the same path.

Default spf

lsp-pacing-interval

Syntax **lsp-pacing-interval** *milli-seconds*
no lsp-pacing-interval

Context config>service>vpls>sap>spb
config>service>vpls>spoke-sdp>spb

Description This command configures the interval during which LSPs are sent from the interface.

To avoid overwhelming neighbors that have less CPU processing power with LSPs, the pacing interval can be configured to limit how many LSPs are sent during an interval. LSPs may be sent in bursts during the interval up to the configured limit. If a value of 0 is configured, no LSPs are sent from the interface.

The **no** form of the command reverts to the default value.

Note: The IS-IS timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default 100 — LSPs are sent in 100 millisecond intervals.

Parameters *milli-seconds* — The interval in milliseconds during which IS-IS LSPs are sent from the interface, expressed as a decimal integer.
0 to 65535

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>service>vpls b-vpls>sap>spb>
config>service>vpls b-vpls>spoke-sdp>spb>

Description This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface. This command is valid only for interfaces on control B-VPLS.

The no form of the command reverts to the default value.

Default 100

Parameters *seconds* — The interval in seconds that SPB IS-IS LSPs can be sent on the interface.

Values 1 to 65535

metric

Syntax **metric** *value*
No metric

Context config>service>vpls b-vpls>sap>spb>level
config>service>vpls b-vpls>spoke-sdp>spb>level

Description This configures metric for this SPB interface SAP/spoke-sdp. This command is valid only for interfaces on control B-VPLS.

Parameters *value* — The configuration metric value.

Values 1 to 16777215

Default 1000

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>service>vpls b-vpls>sap>spb>level
config>service>vpls b-vpls>spoke-sdp>spb>level

Description This command configures the interval in seconds between hello messages issued on this interface at this level. This command is valid only for interfaces on control B-VPLS.

The no form of the command to reverts to the default value.

Default 3 — Hello interval default for the designated inter-system.

9 — Hello interval default for non-designated inter-systems.

Parameters *seconds* — The hello interval in seconds expressed as a decimal integer.

Values 1 to 20000

hello-multiplier

Syntax **hello-multiplier** *multiplier*
no hello-multiplier

Context config>service>vpls b-vpls>sap>spb>level

config>service>vpls b-vpls>spoke-sdp>spb>level

Description This command configures the number of missing hello PDUs from a neighbor SPB declares the adjacency down. This command is valid only for interfaces on control B-VPLS.

The no form of the command reverts to the default value.

Default 3 — SPB can miss up to 3 hello messages before declaring the adjacency down.

Parameters *multiplier* — The multiplier for the hello interval expressed as a decimal integer.

Values 2 to 100

mrp

Syntax **mrp**

Context config>service>vpls
config>service>vpls>mesh-sdp
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command configures Multiple Registration Protocol (MRP) parameters. MRP is valid only under B-VPLS.

attribute-table-size

Syntax [**no**] **attribute-table-size** *value*

Context config>service>vpls>mrp

Description This command controls the number of attributes accepted on a per B-VPLS basis. When the limit is reached, no new attributes will be registered.

If a new lower limit (smaller than the current number of attributes) from a local or dynamic I-VPLS is being provisioned, a CLI warning will be issued stating that the system is currently beyond the new limit. The value will be accepted, but any creation of new attributes will be blocked under the attribute count drops below the new limit; the software will then start enforcing the new limit.

Default maximum number of attributes

Parameters *value* — The maximum number of attributes accepted per B-VPLS.

Values 1 to 2048 (for 7450 ESS-7, 7450 ESS-12, 7750 SR-7, or 7750 SR-12)

attribute-table-high-wmark

Syntax [**no**] **attribute-table-high-wmark** *high-water-mark*

Context config>service>vpls>mrp

Description This command specifies the percentage filling level of the MMRP attribute table where logs and traps are sent.

Default attribute-table-high-wmark 95

Parameters *high-water-mark* — The maximum filling level of the MMRP attribute table, as a percentage.

Values 1 to 100

attribute-table-low-wmark

Syntax [**no**] **attribute-table-low-wmark** *low-water-mark*

Context config>service>vpls>mrp

Description This command specifies the MMRP attribute table low watermark as a percentage. When the percentage filling level of the MMRP attribute table drops below the configured value, the corresponding trap is cleared and/or a log entry is added.

Default attribute-table-low-wmark 90

Parameters *low-water-mark* — The minimum filling level of the MMRP attribute table, as a percentage.

Values 1 to 100

flood-time

Syntax **flood-time** *flood-time*
no flood-time

Context config>service>vpls>mrp

Description This command configures the amount of time, in seconds, after a status change in the VPLS service during which traffic is flooded. When that time expires, traffic will be delivered according to the MMRP registrations that exist in the VPLS. When “no flood-time” is executed, flooding behavior is disabled.

Default no flood-time

Parameters *flood-time* — Specifies the MRP flood time, in seconds.

Values 3 to 600

mrp

Syntax **mrp**

Context config>service

Description This command configures a Multi-service Route Processor (MRP).

mrp-policy

Syntax **mrp-policy** *policy-name* [**create**]
[**no**] **mrp-policy** *policy-name*

Context config>service>mrp

Description This command enables the context for a MRP policy. The mrp-policy specifies either a forward or a drop action for the Group BMAC attributes associated with the ISIDs specified in the match criteria. The mrp-policy can be applied to multiple BVPLS services as long as the scope of the policy is template.

Any changes made to the existing policy, using any of the sub-commands, will be applied immediately to all services where this policy is applied. For this reason, when many changes are required on a mrp-policy, Nokia recommends that the policy be copied to a work area. That work-in-progress policy can be modified until complete and then written over the original mrp-policy. Use the **config mrp-policy copy** command to maintain policies in this manner.

The **no** form of the command deletes the mrp-policy. An MRP policy cannot be deleted until it is removed from all the SAPs or SDPs where it is applied.

Default no mrp-policy is defined

Parameters *policy-name* — Specifies the redirect policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

create — This keyword is required when first creating the configuration context. When the context is created, it is possible to navigate into the context without the **create** keyword.

scope

Syntax **scope** {**exclusive** | **template**}
no scope

Context config>service>mrp>mrp-policy

Description This command configures the filter policy scope as exclusive or template. If the scope of the policy is template and is applied to one or more services, the scope cannot be changed.

The **no** form of the command sets the scope of the policy to the default of template.

Default template

Parameters **exclusive** — When the scope of a policy is defined as exclusive, the policy can only be applied to a single entity (SAP or SDP). Attempting to assign the policy to a second entity will result in an error message. If the policy is removed from the entity, it will become available for assignment to another entity.

template — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs or network ports.

default-action

Syntax **default-action** {**block** | **allow**}

Context config>service>mrp>mrp-policy

Description This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs do not match the specified criteria in all of the entries of the mrp-policy.

When multiple default-action commands are entered, the last command will overwrite the previous command.

Default default-action-allow

Parameters **block** — Specifies that all MMRP attributes will not be declared or registered unless there is a specific mrp-policy entry which causes them to be allowed on this SAP/SDP.

allow — Specifies that all MMRP attributes will be declared and registered unless there is a specific mrp-policy entry which causes them to be blocked on this SAP/SDP.

entry

Syntax [**no**] **entry** *entry-id*

Context config>service>mrp>mrp-policy

Description This command creates or edits an mrp-policy entry. Multiple entries can be created using unique entry-id numbers within the policy. The implementation exits the policy on the first match found and executes the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit. An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.

The no form of the command removes the specified entry from the mrp-policy. Entries removed from the mrp-policy are immediately removed from all services where the policy is applied.

The no form of the command removes the specified entry-id.

Default none

Parameters *entry-id* — An entry-id uniquely identifies a match criteria and the corresponding action. It is recommended that multiple entries be given entry-ids in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

Values 1 to 65535

create — Keyword; required when first creating the configuration context. When the context is created, one can navigate into the context without the create keyword.

description

Syntax **description** *description-string*
no description

Context config>service>mrp>mrp-policy>entry
config>service>mrp>mrp-policy

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

match

Syntax [**no**] **match**

Context config>service>mrp>mrp-policy>entry

Description This command creates the context for entering/editing match criteria for the mrp-policy entry. When the match criteria have been satisfied the action associated with the match criteria is executed. In the current implementation just one match criteria (ISID based) is possible in the entry associated with the mrp-policy. Only one match statement can be entered per entry.

The **no** form of the command removes the match criteria for the entry-id.

isid

Syntax [**no**] i**sid** *value* [**to** *higher-value*]
no isid

Context config>service>mrp>mrp-policy>entry>match

Description This command configures an ISID value or a range of ISID values to be matched by the mrp-policy parent when looking at the related MMRP attributes (Group BMACs). The pbb-etype value for the related SAP (inherited from the Ethernet port configuration) or for the related SDP binding (inherited from SDP configuration) will be used to identify the ISID tag.

Multiple isid statements are allowed under a match node. The following rules govern the usage of multiple isid statements:

- overlapping values are allowed:
  - isid from 1 to 10
  - isid from 5 to 15
  - isid 16
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “isid from 1 to 16” statement.
- there is no consistency check with the content of isid statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation t o execute the associated action for that entry and then to exit the mrp-policy.
- If there are no isid statements under a match criteria but the mac-filter type is isid the following behaviors apply for different actions:
  - For end-station – it treats any ISID value as no match and goes to next entry or default action which must be “block” in this case
  - For allow – it treats any ISID value as a match and allows it
  - For block – it treats any ISID value as a match and blocks it

The **no** form of the command can be used in two ways:

no isid - removes all the previous statements under one match node

no isid *value* **to** *higher-value* - removes a specific ISID value or range. Must match a previously used positive statement: for example if the command “isid 16 to 100” was used using “no isid 16 to 50” will not work but “no isid 16 to 100 will be successful.

Default no isid

Parameters *value* — Specifies an ISID to be used for matching in 24 bits. When used with **to** *higher-value*, *value* specifies the lowest ISID value in the range.

Values 0 to 16777215

higher-value — Specifies the highest ISID value in the range.

Values 0 to 16777215

action

Syntax **action** {**block** | **allow** | **end-station**}
no action

Context config>service>mrp>mrp-policy>entry

Description This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs match the specified ISID criteria in the related entry.

The action keyword must be entered for the entry to be active. Any filter entry without the action keyword will be considered incomplete and will be inactive. If neither keyword is specified (no action is used), this is considered a No-Op policy entry used to explicitly set an entry inactive without modifying match criteria or removing the entry itself. Multiple action statements entered will overwrite previous actions parameters when defined. To remove a parameter, use the no form of the action command with the specified parameter.

The **no** form of the command removes the specified action statement. The entry is considered incomplete and hence rendered inactive without the action keyword.

Default no action

Parameters **block** — Specifies that the matching MMRP attributes will not be declared or registered on this SAP/SDP.

allow — Specifies that the matching MMRP attributes will be declared and registered on this SAP/SDP.

end-station — Specifies that an end-station emulation is present on this SAP/SDP for the MMRP attributes related with matching ISIDs. Equivalent action with the block keyword on that SAP/SDP– the attributes associated with the matching ISIDs do not get declared or registered on the SAP/SDP. The matching attributes on the other hand are mapped as static MMRP entries on the SAP/SDP which implicitly instantiates in the data plane as a MFIB entry associated with that SAP/SDP for the related Group BMAC. For the other SAPs/SDPs in the BVPLS with MRP enabled (no shutdown) this means permanent declaration of the matching attributes, same as in the case when the IVPLS instances associated with these ISIDs were locally configured.

If an mrp-policy has end-station action in one entry, the only default action allowed in the policy is block. Also no other actions are allowed to be configured in other entry configured under the policy.

This policy will apply even if the MRP is shutdown on the local SAP/SDP or for the whole BVPLS to allow for manual creation of MMRP entries in the data plane. Specifically the following rules apply:

- If service vpls mrp shutdown then MMRP on all SAP/SDPs is shutdown - MRP PDUs pass-through transparently
- If service vpls mrp no shutdown and endstation statement (even with no ISID values in the related match statement) is used in a mrp-policy applied to SAP/SDP - no declaration is sent on SAP/SDP. The provisioned ISIDs in the match statement are registered on that SAP/SDP and are propagated on all the other MRP enabled endpoints.

copy

Syntax **copy mrp-policy** *source-name* **to** *dest-name*

Context config>service>mrp

Description This command copies existing mrp-policy list entries for a specific policy name to another policy name. The copy command is a configuration level maintenance tool used to create new mrp-policy using existing mrp-policy.

An error will occur if the destination policy name exists.

Parameters **mrp-policy** — Indicates that source-name and dest-name are MRP policy names.

source-name — Identifies the source mrp-policy from which the copy command will attempt to copy. The mrp-policy with this name must exist for the command to be successful.

dest-name — Identifies the destination mrp-policy to which the copy command will attempt to copy. If the mrp-policy with dest-name exist within the system an error message is generated.

renum

Syntax **renum** *old-entry-id* **to** *new-entry-id*

Context config>service>mrp>mrp-policy

Description This command renumbers existing MRP policy entries to properly sequence policy entries. This may be required in some cases since the implementation exits when the first match is found and executes the actions according to the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters *old-entry-id* — Specifies the entry number of an existing entry.

Values 1 to 65535

new-entry-id — Specifies the new entry number to be assigned to the old entry. If the new entry exists, an error message is generated.

Values 1 to 65535

join-time

Syntax [**no**] **join-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the interval between transmit opportunities that are applied to the Applicant state machine. An instance of this Join Period Timer is required on a per-Port, per-MRP Participant basis. For additional information, refer to IEEE 802.1ak-2007 section 10.7.4.1.

Default 2

Parameters *value* — The interval between transmit opportunities, in tenths of a second.

Values 1 to 10

leave-time

Syntax [**no**] **leave-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the period of time that the Registrar state machine will wait in the leave state before transitioning to the MT state when it is removed. An instance of the timer is required for each state machine that is in the leave state. The Leave Period Timer is set to the value leave-time when it is started.

A registration is normally in an “in” state where there is an MFIB entry and traffic is being forwarded. When a “leave all” is performed (periodically around every 10-15 seconds per SAP/SDP binding - see leave-all-time-below), a node sends a message to its peer indicating a leave all is occurring and puts all of its registrations in leave state.

The peer refreshes its registrations based on the leave all PDU it receives and sends a PDU back to the originating node with the state of all its declarations.

Refer to IEEE 802.1ak-2007 section 10.7.4.2.

Default 30

Parameters *value* — The period of time that the Registrar state machine waits in the leave state before transitioning to the MT state, in tenths of a second.

Values 30 to 60

leave-all-time

Syntax [**no**] **leave-all-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the frequency with which the LeaveAll state machine generates LeaveAll PDUs. The timer is required on a per-Port, per-MRP Participant basis. The Leave All Period Timer is set to a random value, T, in the range LeaveAllTime<T<1.5*leave-all-time when it is started. Refer to IEEE 802.1ak-2007 section 10.7.4.3.

Default 100

Parameters *value* — The frequency with which the LeaveAll state machine generates LeaveAll PDUs, in tenths of a second.

Values 60 to 300

periodic-time

Syntax [**no**] **periodic-time** *value*

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command controls the frequency the Periodic Transmission state machine generates periodic events if the Periodic Transmission Timer is enabled. The timer is required on a per-Port basis. The Periodic Transmission Timer is set to one second when it is started.

Default 10

Parameters *value* — The frequency with which the Periodic Transmission state machine generates periodic events, in tenths of a second.

Values 10 to 100

periodic-timer

Syntax [**no**] **periodic-timer**

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command enables or disables the Periodic Transmission Timer.

Default disabled

send-flush-on-failure

Syntax [**no**] **send-flush-on-failure**

Context config>service>vpls

Description This command enables sending out flush-all-from-me messages to all LDP peers included in affected VPLS, in the event of physical port failures or “operationally down” events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, in the case of link failure, is resolved by RSTP running between a PE router and CE devices. If the endpoint is configured within the VPLS and send-flush-on-failure is enabled, flush-all-from-me messages will be sent out only when all spoke-SDPs associated with the endpoint go down.

This feature cannot be enabled on management VPLS.

Default no send-flush-on-failure

pbb

Syntax **pbb**

Context config>service
config>service>vpl
config>service>epipe

Description This command configures global PBB parameters.

mac-name

Syntax **mac-name** *name ieee-address*
no mac-name *name*

Context config>service>pbb

Description This command configures the MAC name for the MAC address. It associates an ASCII name with an IEEE MAC to improve the PBB Epipe configuration. It can also change the dest-BMAC in one place instead of 1000s of Epipe.

Parameters *name* — Specifies the MAC name up to 32 characters in length.

ieee-address — The MAC address assigned to the MAC name. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

source-bmac

Syntax **source-bmac** *ieee-address*
no source-bmac

Context config>service>pbb

Description This command configures the source B-VPLS MAC address to use with PBB and provisions a chassis level source BMAC.

Parameters *ieee-address* — The MAC address assigned to the BMAC. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

backbone-vpls

Syntax **backbone-vpls** *vpls-id*[:*isid*]
no backbone-vpls

Context config>service>vpls>pbb

Description This command associated the I-VPLS with the B-VPLS service. The ISID value is used to mux/demux packets for the VPLS flowing through the B-VPLS.

Parameters *vpls-id* — This value represents the VPLS ID value associated with the B-VPLS.

isid — Defines ISID associated with the I-VPLS.

Default The default is the service-id.

Values 0 to 16777215

source-bmac

Syntax **source-bmac** *ieee-address*

Context config>service>vpls>pbb

Description This command configures the base source BMAC for the B-VPLS. The first 32 bits must be the same with what is configured in the MC-LAG peer. If not configured here, it will inherit the chassis level BMAC configured under the new PBB object added in the previous section. If the **use-sap-bmac** command is on, the value of the last 16 bits (lsb) of the source BMAC must be part of the **reserved-source-bmac-lsb** configured at chassis level, under service PBB component. If that is not the case, the command will fail.

use-sap-bmac

Syntax [**no**] **use-sap-bmac**

Context config>service>vpls>pbb

Description This command enables on a per BVPLS basis the use of source BMACs allocated to multi-homed SAPs (assigned to an MC-LAG) in the related IVPLS or Epipe service. The command will fail if the value of the source-bmac assigned to the BVPLS is the hardware (chassis) BMAC. In other words, the **source-bmac** must be a configured one.

Default no use-sap-bmac

mac-notification

Syntax **mac-notification**

Context config>service>vpls

Description This command controls the settings for the MAC notification message.

The mac-notification message must be generated under the following events:

1. When enabled in the BVPLS using no shutdown, a MAC notification will be sent for every active MC-LAG link. The following 3 cases assume no shutdown in the BVPLS.
2. Whenever a related MC-LAG link becomes active (related MC-LAG link = has at least 1 SAP associated with the BVPLS) if the MC-LAG peering is initialized and the PE peers are synchronized.
3. 1st SAP on an active MC-LAG is associated (via IVPLS/Epipe) with the BVPLS
4. The link between IVPLS/Epipe and BVPLS is configured and there are I-SAPs configured on an active MC-LAG link.

The MAC notification is not sent for the following events:

1. Change of source-bmac or source-bmac-lsb
2. On changes of use-sap-bmac parameter
3. If MC-LAG peering is not (initialized and in sync).

interval

Syntax [**no**] **interval** *value*

Context config>service>vpls>mac-notification

Description This command controls the frequency of subsequent MAC notification messages.

Default Inherits the chassis level configuration from config>service>mac-notification

Parameters *value* — Specifies the frequency of subsequent MAC notification messages, in tenths of a second.

Values 1 to 100

renotify

Syntax **renotify** *value*
no renotify

Context config>service>vpls>pbb>mac-notification

Description This command controls the periodic interval at which sets of MAC notification messages are sent. At each expiration of the renotify timer, a new burst of notification messages is sent, specifically <count> frames at <interval> deci-seconds.

Default no renotify

Parameters *value* — Specifies the time interval between re-notification, in seconds.

Values 240 to 840

count

Syntax [**no**] **count** *value*

Context config>service>vpls>pbb>mac-notification

Description This command configures how often MAC notification messages are sent.

Parameters *value* — Specifies, in seconds, how often MAC notification messages are sent.

Values 1 to 10

Default Inherits the chassis level configuration from config>service>mac-notification

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls

Description This command disables the sending of the notification message in the BVPLS domain.

Default shutdown

backbone-vpls

Syntax **backbone-vpls** *service-id* [**isid** *isid*]
no backbone-vpls

Context config>service>vpls>pbb

Description This command configures B-VPLS service associated with the I-VPLS.

Parameters *service-id* — Specifies the service ID.

Values 1 to 2147483648

isid — Specifies the ISID.

Values 0 to 16777215

igmp-snooping

Syntax **igmp-snooping**

Context config>service>vpls>pbb>bvpls
config>service>vpls>pbb>bvpls>sap
config>service>vpls>pbb>bvpls>sdp

Description This command configures IGMP snooping attributes for I-VPLS.

mld-snooping

Syntax **mld-snooping**

Context config>service>vpls>pbb>bvpls
config>service>vpls>pbb>bvpls>sap
config>service>vpls>pbb>bvpls>sdp

Description This command configures MLD snooping attributes for I-VPLS.

mrouter-dest

Syntax [**no**] **mrouter-dest** *mac-name*

Context config>service>vpls>pbb>bvpls>igmp-snooping
config>service>vpls>pbb>bvpls>mld-snooping

Description This command configures the destination BMAC address name to be used in the related backbone VPLS to reach a specific IGMP or MLD snooping MRouter. The name is associated at system level with the MAC address, using the command mac-name.

Parameters *mac-name* — Specifies the MAC name.

Values 32 chars max

sap

Syntax [**no**] **sap** *sap-id*

Context config>service>vpls
config>service>vpls>pbb>backbone-vpls

Description This command configures attributes of a SAP on the B-VPLS service.

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>pbb>bvpls>sap>igmp-snooping
config>service>vpls>pbb>bvpls>sdp>igmp-snooping
config>service>vpls>pbb>bvpls>sap>mld-snooping
config>service>vpls>pbb>bvpls>sdp>mld-snooping

Description This command specifies whether a multicast router is attached behind this SAP or spoke-SDP.

Configuring a SAP or spoke-SDP as an mrouter-port will have a double effect. Firstly, all multicast traffic received on another SAP or spoke-SDP will be copied to this SAP or spoke-SDP. Secondly, IGMP or MLD reports generated by the system as a result of someone joining or leaving a multicast group, will be sent to this SAP or SDP.

If two multicast routers exist in the local area network, one of them will become the active querier. The other multicast router (non-querier) stops sending IGMP or MLD queries, but it should still receive reports to keep its multicast trees up to date. To support this, the mrouter-port should be enabled on all SAPs or spoke-SDPs connecting to a multicast router.

The IGMP version to be used for the reports (v1, v2 or v3) or MLD version (v1 or v2) can only be determined after an initial query has been received. Until such time no reports are sent on the SAP, even if mrouter-port is enabled.

If the **send-queries** command is enabled on this SAP or spoke-SDP, the **mrouter-port** parameter can not be set.

Default no mrouter-port

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context config>service>vpls>pbb>backbone-vpls

Description This command configures attributes of a SDP binding on the B-VPLS service.

Parameters *sdp-id* — Specifies the SDP ID.

Values 1 to 17407

vc-id — Specifies the VC ID.

Values 1 to 4294967295

stp

Syntax [**no**] **stp**

Context config>service>Vpls>pbb>backbone-vpls

Description This command enables or disable STP through B-VPLS service.

force-qtag-forwarding

Syntax [**no**] **force-qtag-forwarding**

Context config>service>vpls>pbb

Description This command forces the addition of a IEEE 802.1q tag after the Customer MAC (CMAC) addresses when the PBB header is built, as it egresses a related BVPLS.

It is used to preserve the dot1q and DE bits from the customer domain when the service delimiting qtags are stripped when the packet is ingressing a PBB Epipe or an IVPLS. The VLAN value of the service delimiting QTAG if one exists is used for the corresponding inserted dot1q field. If a service delimiting QTAG does not exist, then the value of zero is used for all the inserted QTAG bits.

The **no** form of this command sets default behavior.

Default disabled

propagate-mac-flush-from-bvpls

Syntax [**no**] **propagate-mac-flush-from-bvpls**

Context config>service>vpls>pbb

Description This command enables the propagation in the local PBB of any regular LDP MAC Flush received in the related B-VPLS. If an LDP MAC flush-all-but-mine is received in the B-VPLS context, the command controls also whether a flush is performed for all the customer MACs in the associated FDB. The command does not have any effect on a PBB MAC Flush (LDP MAC flush with PBB TLV) received in the related B-VPLS context.

The **no** form of this command disables the propagation of LDP MAC Flush i from the related B-VPLS.

Default no propagate-mac-flush-from-bvpls

send-flush-on-bvpls-failure

Syntax [**no**] **send-flush-on-bvpls-failure**

Context config>service>vpls>pbb

Description This command enables the generation in the local I-VPLS of an LDP MAC flush-all-from-me following a failure of SAP/the whole endpoint/spoke-SDP in the related B-VPLS. The failure of mesh-SDP in B-VPLS does not generate the I-VPLS MAC flush.

The **no** form of this command disables the generation of LDP MAC flush in I-VPLS on failure of SAP/endpoint/spoke-SDP in the related B-VPLS.

Default no send-flush-on-bvpls-failure

mrp-policy

Syntax [**no**] **mrp-policy**

Context config>service>vpls>sap>mrp
config>service>vpls>spoke-sdp>mrp
config>service>vpls>mesh-sdp>mrp

Description This command instructs MMRP to use the mrp-policy defined in the command to control which group BMAC attributes will be declares and registered on the egress SAP/Mesh-SDP/ Spoke-SDP. The Group BMACs will be derived from the ISIDs using the procedure used in the PBB solution. The Group MAC = standard OUI with the last 24 bits being the ISID value. If the policy-name refers to a non-existing mrp-policy the command should return error. Changes to a mrp-policy are allowed and applied to the SAP/SDPs under which the policy is referenced.

Default no mrp-policy

send-bvpls-flush

Syntax [**no**] **send-bvpls-flush** {[**all-from-me**] | [**all-but-mine**]}

Context config>service>vpls

Description This command configures the BVPLS flush. If B-SDPs are used and MAC notification mechanism is turned on in the related BVPLS (MPLS use case), it makes sense to turn off the T-LDP MAC Flush.

Parameters **all-from-me** — Flushes on a negative event, such as pseudowire failure

all-but-mine — Flushes on a positive event, such as pseudowire activation

mac-notification

Syntax **mac-notification**

Context config>service>pbb

Description This command controls the settings for the MAC notification messages.

interval

Syntax [**no**] **interval** *value*

Context config>service>pbb>mac-notification

Description This command controls the frequency of subsequent MAC notification messages.

Default 100 ms

Parameters *value* — Specifies the frequency of subsequent MAC notification messages, in tenths of a second

Values 1 to 100

count

Syntax [**no**] **count** *value*

Context config>service>pbb>mac-notification

Description This command configures how often MAC notification messages are sent.

Parameters *value* — Specifies, in seconds, how often MAC notification messages are sent

Values 1 to 10

Default 3

epipe

Syntax **epipe** *service-id* **customer** *customer-id* [**vpn** *vpn-id*] [**vc-switching**] [**create**]
epipe *service-id*
no epipe *service-id*

Context config>service

Description This command configures an Epipe service instance. This command is used to configure a point-to-point epipe service. An Epipe connects two endpoints defined as Service Access Points (SAPs). Both SAPs may be defined in one or, for the 7750 SR, they may be defined in separate devices connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far end SAP is generalized into a Service Distribution Point (SDP). This SDP describes a destination and the encapsulation method used to reach it.

No MAC learning or filtering is provided on an Epipe.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. After a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

After a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

By default, no Epipe services exist until they are explicitly created with this command.

The **no** form of this command deletes the Epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values 1 to 2147483648

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number. This parameter applies only to the 7750 SR.

Values 1 to 2147483647

Default null (0)

vc-switching — Specifies if the pseudowire switching signaling is used for the spoke-SDPs configured in this service. This parameter applies only to the 7750 SR.

create — Keyword used to create the service instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

tunnel

Syntax **tunnel** *service-id* **backbone-dest-mac** {*mac-name* | *ieee-mac*} **isid** *isid*
no tunnel

Context config>service>epipe>pbb

Description This command configures a Provider Backbone Bridging (PBB) tunnel with Backbone VPLS (B-VPLS) service information.

Parameters *service-id* — Specifies the B-VPLS service for the PBB tunnel associated with this service

Values 1 to 2147483648

backbone-dest-mac {*mac-name* | *ieee-mac*} — Specifies the backbone destination MAC-address for PBB packets

isid — Specifies a 24 bit service instance identifier for the PBB tunnel associated with this service. As part of the PBB frames, it is used at the destination PE as a demultiplexor field.

Values 0 to 16777215

4.5 PBB Show, Clear, and Debug Command Reference

This section describes the PBB show, clear, and debug command reference.

4.5.1 Command Hierarchies

4.5.1.1 Show Commands

show
- **eth-cfm**
  - **association** [*ma-index*] [**detail**]
  - **cfm-stack-table**
  - **cfm-stack-table port** [{**all-ports** | **all-sdps** | **all-virtuals**}] [**level** 0 to 7] [**direction up** | **down**]
  - **cfm-stack-table** *port-id* [**vlan** *qtag*[*.qtag*]] [**level** 0 to 7] [**direction up** | **down**]
  - **cfm-stack-table sdp** *sdp-id*[**:***vc-id*] [**level** 0 to 7] [**direction up**| **down**]
  - **cfm-stack-table virtual** *service-id* [**level** 0 to 7]
  - **cfm-stack-table facility** [{**all-ports** | **all-lags** | **all-lag-ports** | **all-tunnel-meps** | **all-router-interfaces**}] [**level** 0 to 7] [**direction up** | **down**]
  - **cfm-stack-table facility collect-lmm-stats**
  - **cfm-stack-table facility lag** *id* [**tunnel** 1 to 4094] [**level** 0 to 7] [**direction up** | **down**]
  - **cfm-stack-table facility port** *id* [**level** 0 to 7] [**direction up** | **down**]
  - **cfm-stack-table facility router-interface** *ip-int-name* [**level** 0 to 7] [**direction up** | **down**]
  - **domain** [*md-index*] [**association** *ma-index* | **all-associations** [**detail**]]
  - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
- **service**
  - **id** *service-id*
    - **i-vpls**
    - **mrp-policy mac** [*ieee-address*]
    - **mrp**
    - **spb**
      - **adjacency** [**detail**]
      - **base**
      - **database**
      - **fate-sharing**
      - **fid** [*fid*] **fate-sharing**
      - **fid** [*fid*] **user-service**
      - **fid** [*fid*] **fdb**
      - **fid** [*fid*] **mfib** [**group-mac** *ieee-address*]
      - **fid** [*fid*] **mfib** [**isid** *isid*]
      - **hostname**
      - **interface**
      - **mfib** [**detail**]

```
                    — routes
                    — spf
                    — spf-log
                    — status
            — mrp-policy [mrp-policy]
            — mrp-policy mrp-policy [association]
            — mrp-policy mrp-policy [entry entry-id]
            — pbb
                — base
                — mac-name [detail]
```

4.5.1.2 Clear Commands

```
clear
    — service
        — statistics
            — id service-id
                — counters
                — mesh-sdp sdp-id[:vc-id] {all | counters | stp | mrp}
                — mrp
                — spoke-sdp sdp-id[:vc-id] {all | counters | stp | mrp}
                — stp
                — spb
                    — adjacency [detail]
                    — database
                    — spf-log
                    — status
            — sap sap-id {all | counters | stp | l2pt | mrp}
```

4.5.1.3 Debug Commands

```
debug
    — service
        — id service-id
            — [no] mrp
                — all-events
                — [no] applicant-sm
                — [no] leave-all-sm
                — [no] mmrp-mac ieee-address
                — [no] mrpdu
                — [no] periodic-sm
                — [no] registrant-sm
                — [no] sap sap-id
                — [no] sdp sdp-id:vc-id
```

4.5.2 Command Descriptions

4.5.2.1 PBB Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

eth-cfm

Syntax **eth-cfm**

Context show

Description This command displays 802.1ag CFM information.

association

Syntax **association** [*ma-index*] [**detail**]

Context show>eth-cfm

Description Shows association information.

Parameters *ma-index* — Specifies the MA index value

Values 1 to 4294967295

detail — Displays all association detail

Output The following output is an example of ETH-CFM association information.

Sample Output

```
ALU-IPD# show eth-cfm association
===============================================================================
CFM Association Table
===============================================================================
Md-index   Ma-index   Name                      Int  Hold  Bridge-id  MEPS TxSid
-------------------------------------------------------------------------------
10         1          port1/1/1                 10   n/a   none       1    no
12         1          ipinterface192.168.2.0    1    n/a   none       2    yes
12         4000       vpls-4000-l2              1    n/a   4000       2    yes
12         4001       vpls-4001-l2              1    n/a   4001       2    yes
12         5001       vprn-5001-10.101.28.1     1    n/a   5001       2    no
13         1000       vpls-1000-l3              1    n/a   1000       3    yes
13         1500       epipe-1500-l3             1    n/a   1500       2    yes
13         2000       vpls-2000-l3              1    n/a   2000       5    yes
13         2002       vpls-2002-l3              1    n/a   2002       2    yes
```

```
13          3000        vpls-3000-l3              1    n/a   3000        4    yes
13          4000        vpls-4000-l3              1    n/a   4000        2    yes
13          4001        vpls-4001-l3              1    n/a   4001        2    yes
14          100         vpls-100-l4               1    n/a   100         4    yes
14          1000        vpls-1000-l4              10   n/a   1000        1    no
14          2000        vpls-2000-l4              10   n/a   2000        0    yes
14          4000        vpls-4000-l4              10   n/a   4000        0    yes
14          4001        vpls-4001-l4              1    n/a   4001        1    yes
15          1000        vpls-1000-l5              10   n/a   1000        0    no
15          2000        vpls-2000-l5              10   n/a   2000        0    yes
15          4000        vpls-4000-l5              10   n/a   4000        0    yes
===============================================================================
ALU-IPD#
```

cfm-stack-table

Syntax **cfm-stack-table**
cfm-stack-table [{**all-ports** | **all-sdps** | **all-virtuals**}] [**level** 0 to 7] [**direction up** | **down**]
cfm-stack-table port *port-id* [**vlan** *qtag*[*.qtag*]] [**level** 0 to 7] [**direction up** | **down**]
cfm-stack-table sdp *sdp-id*[**:***vc-id*] [**level** 0 to 7] [**direction up** | **down**]
cfm-stack-table virtual *service-id* [**level** 0 to 7]
cfm-stack-table facility [{**all-ports** | **all-lags** | **all-lag-ports** | **all-tunnel-meps** | **all-router-interfaces**}] [**level** 0 to 7] [**direction up** | **down**]
cfm-stack-table facility collect-lmm-stats
cfm-stack-table facility lag *lag-id* [**tunnel** 1 to 4094] [**level** 0 to 7] [**direction up** | **down**]
cfm-stack-table facility port *port-id* [**level** 0 to 7] [**direction up** | **down**]
cfm-stack-table facility router-interface *ip-int-name* [**level** 0 to 7] [**direction up** | **down**]

Context show>eth-cfm

Description This command displays stack table information. This stack table is used to display the various management points MEPs and MIPs that are configured on the system. These can be Service based or facility based. The various option allow the operator to be specific. If no parameters are include then the entire stack table will be displayed.

Parameters **all-ports** — Displays stack table information for all ports

all-sdps — Displays stack table information for all SDPs

all-virtuals — Displays stack table information for all virtual circuits

all-lags — Displays stack table information for all LAGs

all-lag-parts — Displays stack table information for all LAG parts

all-tunnel-meps — Displays stack table information for all tunnel MEPs

all-router-interfaces — Displays stack table information for all router interfaces

level 0 to 7 — Displays the MD level of the maintenance point

Values 0 to 7

direction up | down — Displays the direction in which the MP faces on the bridge port

port-id — Displays the bridge port or aggregated port on which MEPs or MHFs are configured

lag-id — Displays stack table information for the specified LAG

qtag — Specifies a VLAN qtag

service-id — Displays CFM stack table information for the specified SDP

tunnel 1 to 4094 — Displays information for the specified tunnel

Values 1 to 4094

facility — Displays the CFM stack table information for facility MEPs. The base command will display all the facility MEPs. Options may be included in order to further parse the table for specific facility MEP information.

sdp-id[:*vc-id*] — Displays CFM stack table information for the specified SDP

ip-int-name — Displays stack table information for the specified IP interface

collect-lmm-stats — Displays collected LMM statistics for the facility

Output The following output is an example of ETH-CFM stack table information.

Sample Output

```
show eth-cfm cfm-stack-table
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
A = AisRx, L = CSF LOS Rx, F = CSF AIS/FDI rx, r = CSF RDI rx
G = receiving grace PDU (MCC-ED or VSM) from at least one peer
===============================================================================
CFM SAP Stack Table
===============================================================================
Sap                Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
1/2/1:51.28          2 D         12       5001   28 d8:1c:01:02:00:01 --C---- -
1/2/1:1000.1000      3 U         13       1000   28 00:00:00:00:00:28 ---E--- -
1/2/1:1001.1001      1 B          0          0  MIP d8:1c:01:02:00:01 ------- -
1/2/1:1500.1500      3 U         13       1500   28 00:00:00:00:00:28 ------- -
1/2/1:2000.2000      3 U         13       2000  128 d8:1c:01:02:00:01 ------- -
1/2/1:2000.2000      4 B         14       2000  MIP 00:00:00:00:01:28 ------- -
1/2/1:3000.3000      4 B          0          0  MIP d8:1c:01:02:00:01 ------- -
1/2/1:4000.*         3 U         13       4000   28 00:00:00:00:00:28 ------- -
1/2/1:4001.*         3 U         13       4001   28 00:00:00:00:00:28 ------- -
1/2/1:4001.*         4 D         14       4001   28 00:00:00:00:00:28 ------- -
===============================================================================

===============================================================================
CFM Ethernet Tunnel Stack Table
===============================================================================
Eth-tunnel         Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

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

```
CFM Ethernet Ring Stack Table
===============================================================================
Eth-ring           Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
CFM Facility Port Stack Table
===============================================================================
Port     Tunnel    Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
1/1/1     0          0 D         10          1   28 d8:1c:01:01:00:01 ------- -
===============================================================================

===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
CFM Facility Tunnel Stack Table
===============================================================================
Port/Lag Tunnel    Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
CFM Facility Interface Stack Table
===============================================================================
Interface          Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
v28-v29-1/1/6        2 D         12          1   28 00:00:00:00:00:28 ------- -
===============================================================================

===============================================================================
CFM SAP Primary VLAN Stack Table
===============================================================================
Sap
  Primary VlanId   Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
1/1/10:2002.*
    2002             3 U         13       2002   28 00:00:00:00:00:28 ------- -
1/2/1:4000.*
    4000             4 B         14       4000  MIP d8:1c:01:02:00:01 ------- -
    4000             5 B         15       4000  MIP d8:1c:01:02:00:01 ------- -
===============================================================================

===============================================================================
CFM SDP Stack Table
===============================================================================
Sdp                Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
2829:4001            2 D         12       4001   28 00:00:00:00:00:28 ------- -
===============================================================================
```

```
===============================================================================
CFM SDP Primary VLAN Stack Table
===============================================================================
Sdp
  Primary VlanId   Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
2829:4000
    4000             2 D         12       4000   28 00:00:00:00:00:28 ------- -
    4000             4 B         14       4000  MIP d8:1c:ff:00:00:00 ------- -
    4000             5 B         15       4000  MIP d8:1c:ff:00:00:00 ------- -
===============================================================================

===============================================================================
CFM Virtual Stack Table
===============================================================================
Service            Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
100                  4 U         14        100   28 00:00:00:00:00:28 ------- -
1000                 4 U         14       1000   28 d8:1c:ff:00:00:00 ---E--- -
2000                 3 U         13       2000   28 00:00:00:00:00:28 ------- -
3000                 3 U         13       3000   28 00:00:00:00:00:28 R-C---- -
===============================================================================
```

domain

Syntax **domain** [*md-index*] [**association** *ma-index* | **all-associations** [**detail**]]

Context show>eth-cfm>domain

Description This command displays domain information.

Parameters *md-index* — Specifies the maintenance domain (MD) index value

Values 1 to 4294967295

ma-index — Specifies the MA index value

Values 1 to 4294967295

all-associations — Displays information all maintenance associations

detail — Displays detailed information.

Output The following output is an example of ETH-CFM domain information.

Sample Output

```
*A:node-1# show eth-cfm domain
===============================================================================
CFM Domain Table
===============================================================================
Md-index   Level Name                                        Format
-------------------------------------------------------------------------------
10         0     InfrastructureL0                           CharString
12         2                                                none
```

```
13          3                                                   none
14          4                                                   none
15          5                                                   none
===============================================================================
```

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]

Context show>eth-cfm>domain

Description This command displays Maintenance Endpoint (MEP) information.

Parameters *mep-id* — Specifies the maintenance association end point identifier

Values 1 to 8191

md-index — Specifies the maintenance domain (MD) index value

Values 1 to 4294967295

ma-index — Specifies the MA index value

Values 1 to 4294967295

loopback — Displays loopback information for the specified MEP

linktrace — Displays linktrace information for specified MEP

Output The following output is an example of ETH-CFM MEP information.

Sample Output

```
# show eth-cfm mep 28 domain 13 association 1000
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index          : 13                        Direction         : Up
Ma-index          : 1000                      Admin             : Enabled
MepId             : 28                        CCM-Enable        : Enabled
IfIndex           : 37781504                  PrimaryVid        : 65537000
Description       : (Not Specified)
FngAlarmTime      : 0                         FngResetTime      : 1000
FngState          : fngReset                  ControlMep        : False
LowestDefectPri   : xcon                      HighestDefect     : none
Defect Flags      : bDefErrorCCM
Mac Address       : 00:00:00:00:00:28         Collect LMM Stats : disabled
LMM FC Stats      : None
LMM FC In Prof    : None
TxAis             : noTransmit                TxGrace           : noTransmit
Facility Fault    : disabled
CcmLtmPriority    : 7                         CcmPaddingSize    : 0 octets
CcmTx             : 189745                    CcmSequenceErr    : 0
CcmTxIfStatus     : Up                        CcmTxPortStatus   : Up
CcmTxRdi          : False                     CcmTxCcmStatus    : transmit
CcmIgnoreTLVs     : (Not Specified)
Fault Propagation: disabled                   FacilityFault     : n/a
```

```
MA-CcmInterval   : 1                        MA-CcmHoldTime    : 0ms
MA-Primary-Vid   : Disabled
Eth-1Dm Threshold: 3(sec)                   MD-Level          : 3
Eth-1Dm Last Dest: 00:00:00:00:00:00
Eth-Dmm Last Dest: 00:00:00:00:00:00
Eth-Ais          : Disabled
Eth-Ais Tx defCCM: allDef
Eth-Tst          : Enabled                  Eth-Tst Pattern   : allZerosNoCrc
Eth-Tst dataLeng*: 64                       Eth-Tst Priority  : 7
Eth-Tst Dest Mac : 00:00:00:00:00:00        Eth-Tst Dest MEP  : 0
Eth-Tst Threshold: 1(bitError)
Eth-Tst Last Dest: 00:00:00:00:00:00
Eth-CSF          : Disabled
Eth-Cfm Grace Tx : Enabled                  Eth-Cfm Grace Rx  : Enabled
Eth-Cfm ED Tx    : Disabled                 Eth-Cfm ED Rx     : Enabled
Eth-Cfm ED Rx Max: 0
Eth-Cfm ED Tx Pri: CcmLtmPri (7)
Redundancy:
    MC-LAG State : n/a
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
===============================================================================
* indicates that the corresponding row element may have been truncated.

show eth-cfm mep 28 domain 13 association 1000 all-remote-mepids
===============================================================================
Eth-CFM Remote-Mep Table
===============================================================================
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
-------------------------------------------------------------------------------
29         True  False Up       Up     00:00:00:00:00:29 01/04/2017 06:03:05
31         True  False Up       Up     00:00:00:00:00:31 01/04/2017 06:03:09
===============================================================================
Entries marked with a 'T' under the 'AD' column have been auto-discovered.

show eth-cfm mep 28 domain 13 association 1000 all-remote-mepids detail
===============================================================================
Eth-CFM Remote-MEP Information
===============================================================================
Remote MEP ID    : 29                   CC Rx State      : True
Auto Discovered  : False                RDI              : False
Port Status TLV  : Up                   I/F Status TLV   : Up
MAC Address      : 00:00:00:00:00:29    CCM Last Change  : 01/04/2017 06:03:05
Chass. ID SubType: chassisComponent
Chassis ID       : 63:73:65:73:2D:76:32:39
                   "cses-v29"
Remote MEP ID    : 31                   CC Rx State      : True
Auto Discovered  : False                RDI              : False
Port Status TLV  : Up                   I/F Status TLV   : Up
MAC Address      : 00:00:00:00:00:31    CCM Last Change  : 01/04/2017 06:03:09
Chass. ID SubType: chassisComponent
Chassis ID       : 63:73:65:73:2D:56:33:31
                   "cses-V31"
===============================================================================

show eth-cfm mep 28 domain 13 association 1000 remote-mepid 29
===============================================================================
```

```
Eth-CFM Remote-Mep Table
==============================================================================
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
------------------------------------------------------------------------------
29         True  False Up       Up     00:00:00:00:00:29 01/04/2017 06:03:05
==============================================================================
Entries marked with a 'T' under the 'AD' column have been auto-discovered.

show eth-cfm mep 28 domain 13 association 1000 remote-mepid 29 detail
===============================================================================
Eth-CFM Remote-MEP Information
===============================================================================
Remote MEP ID    : 29                   CC Rx State      : True
Auto Discovered  : False                RDI              : False
Port Status TLV  : Up                   I/F Status TLV   : Up
MAC Address      : 00:00:00:00:00:29    CCM Last Change  : 01/04/2017 06:03:05
Chass. ID SubType: chassisComponent
Chassis ID       : 63:73:65:73:2D:76:32:39
                   "cses-v29"
===============================================================================
```

i-vpls

Syntax **i-vpls**

Context show>service>id

Description Displays I-VPLS services associated with the B-VPLS service. This command only applies when the service is a B-VPLS.

Output The following output is an example of service I-VPLS information.

Sample Output

```
*A:SetupCLI# show service id 2002 i-vpls
============================================================
Related iVpls services for bVpls service 2002
============================================================
iVpls SvcId         Oper ISID           Admin               Oper
------------------------------------------------------------------------------
2001                122                 Up                  Down
------------------------------------------------------------------------------
Number of Entries : 1
------------------------------------------------------------------------------
*A:alcag1-R6#
*A:term17>show>service>id# i-vpls
===============================================================================
Related iVpls services for bVpls service 2000
===============================================================================
iVpls SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
2100                2100                Up                  Up
2110                123                 Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 2
```

```
-------------------------------------------------------------------------------
*A:SetupCLI#
```

base

Syntax **base**

Context show>service>pbb

Description This command displays information about a PBB base.

Output The following output is an example of PBB base information.

Sample

```
*A:Dut-B# show service pbb base
===============================================================================
PBB MAC Information
===============================================================================
MAC-Notif Count : 3
MAC-Notif Interval : 1
Source BMAC : Default
===============================================================================
```

mac-name

Syntax **mac-name** [**detail**]

Context show>service>pbb

Description This command displays information on a specific MAC name.

Parameters **detail** — Displays detailed MAC name information.

Output The following output is an example of PBB MAC information.

Sample

```
*A:Dut-B# show service pbb mac-name
===============================================================================
MAC Name Table
===============================================================================
MAC-Name MAC-Address
-------------------------------------------------------------------------------
test 00:03:03:03:03:02
===============================================================================
*A:Dut-B# show service pbb mac-name test detail
===============================================================================
Services Using MAC name='test' addr='00:03:03:03:03:02'
===============================================================================
Svc-Id ISID
```

```
------------------------------------------------------------------------
501 501
------------------------------------------------------------------------
Number of services: 1
========================================================================
*A:Dut-B#
```

id

Syntax **id** *service-id*

Context show>service

Description This command displays information on a specific service ID.

Parameters *service-id* — Specifies the service ID.

Output The following output is an example of service ID information.

Sample

```
*A:Dut-B# show service id 1 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id : 1 Vpn Id : 0
Service Type : b-VPLS
Description : (Not Specified)
Customer Id : 1
Last Status Change: 05/17/2009 19:33:11
Last Mgmt Change : 05/17/2009 19:31:59
Admin State : Up Oper State : Up
MTU : 2000 Def. Mesh VC Id : 1
SAP Count : 1 SDP Bind Count : 0
Snd Flush on Fail : Disabled Host Conn Verify : Disabled
Propagate MacFlush: Disabled
Oper Backbone Src : 00:03:00:00:04:01 Use SAP B-MAC : enabled
i-Vpls Count : 0
Epipe Count : 900
*A:Dut-B# show service id 501 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id : 501 Vpn Id : 0
Service Type : Epipe
Description : (Not Specified)
Customer Id : 1
Last Status Change: 05/17/2009 19:41:32
Last Mgmt Change : 05/17/2009 19:40:03
Admin State : Up Oper State : Up
MTU : 1514
Vc Switching : False
SAP Count : 1 SDP Bind Count : 0
-------------------------------------------------------------------------------
PBB Tunnel Point
```

```
-------------------------------------------------------------------------------
B-vpls Backbone-dest-MAC Isid AdmMTU OperState Flood Oper-dest-MAC
-------------------------------------------------------------------------------
1 test 501 2000 Up Yes 00:03:03:03:03:02
-------------------------------------------------------------------------------
*A:Dut-B#
```

mrp

Syntax **mrp**

Context show>service>id

Description This command displays information on a a per service MRP configuration.

Output The following output is an example of service MRP information.

Sample Output

```
*A:PE-A# show service id 10 mrp
-------------------------------------------------------------------------------
MRP Information
-------------------------------------------------------------------------------
Admin State        : Up                Failed Register Cnt: 0
Max Attributes     : 2048              Attribute Count    : 10
Flood Time         : Off
-------------------------------------------------------------------------------
*A:PE-A#
```

mrp-policy

Syntax **mrp-policy** [*mrp-policy*]
mrp-policy *mrp-policy* [**association**]
mrp-policy *mrp-policy* [**entry** *entry-id*]

Context show>service

Description This command displays MRP policy information.

Parameters *mrp-policy* — Specifies the MRP policy.

Values 32 chars max

entry-id — Specifies the entry ID.

Values 1 to 65535

association — Displays associations for the specified MRP policy.

mmrp

Syntax **mmrp mac** [*ieee-address*]

Context show>service>id

Description This command displays information on MACs. If a MAC address is specified, information will be displayed relevant to the specific group. No parameter will display information on all group MACs on a server.

Parameters *ieee-address* — Specifies a MAC address as a hex string in the form of xx:xx:xx:xx:xx:xx: or xx-xx-xx-xx-xx-xx

Output The following output is an example of service MRRP MAC information.

Sample Output

```
*A:PE-A# show service id 10 mmrp mac 01:1E:83:00:00:65
-------------------------------------------------------------------------------
SAP/SDP                                 MAC Address       Registered  Declared
-------------------------------------------------------------------------------
sap:1/1/4:10                            01:1e:83:00:00:65 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:65 No          Yes
sap:2/2/5:10                            01:1e:83:00:00:65 Yes         Yes
-------------------------------------------------------------------------------
*A:PE-A#

*A:PE-A# show service id 10 mmrp mac
-------------------------------------------------------------------------------
SAP/SDP                                 MAC Address       Registered  Declared
-------------------------------------------------------------------------------
sap:1/1/4:10                            01:1e:83:00:00:65 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:66 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:67 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:68 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:69 No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6a No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6b No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6c No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6d No          Yes
sap:1/1/4:10                            01:1e:83:00:00:6e No          Yes
sap:1/2/2:10                            01:1e:83:00:00:65 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:66 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:67 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:68 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:69 No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6a No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6b No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6c No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6d No          Yes
sap:1/2/2:10                            01:1e:83:00:00:6e No          Yes
sap:2/2/5:10                            01:1e:83:00:00:65 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:66 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:67 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:68 Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:69 Yes         Yes
```

```
sap:2/2/5:10                            01:1e:83:00:00:6a Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6b Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6c Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6d Yes         Yes
sap:2/2/5:10                            01:1e:83:00:00:6e Yes         Yes
-------------------------------------------------------------------------------
*A:PE-A#
```

spb

Syntax **spb**

Context clear>service>id

Description This command clears STP related data.

adjacency

Syntax **adjacency** [**detail**]

Context show>service>id>spb

Description This command displays SPB adjacency information.

Parameters **detail** — Shows detailed information

Output The following output is an example of service SPB adjacency information.

Sample Output

```
===============================================================================
ISIS Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT Enab
-------------------------------------------------------------------------------
Dut-B                   L1    Up    19   sap:1/2/2:1.1                No
Dut-C                   L1    Up    21   sap:1/2/3:1.1                No
-------------------------------------------------------------------------------
Adjacencies : 2
===============================================================================
```

base

Syntax **base**

Context show>service>id>spb

Description This command displays SPB base information.

Output The following output is an example of service SPB base information.

Sample Output

```
*A:Dut-A# show service id 100001 spb base
===============================================================================
Service SPB Information
===============================================================================
Admin State        : Up                 Oper State         : Up
ISIS Instance      : 1024               FID                : 1
Bridge Priority    : 8                  Fwd Tree Top Ucast : spf
Fwd Tree Top Mcast : st
Bridge Id          : 80:00.00:10:00:01:00:01
Mcast Desig Bridge : 80:00.00:10:00:01:00:01

===============================================================================
ISIS Interfaces
===============================================================================
Interface                        Level CircID  Oper State   L1/L2 Metric
-------------------------------------------------------------------------------
sap:1/2/2:1.1                    L1    65536   Up           10/-
sap:1/2/3:1.1                    L1    65537   Up           10/-
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
FID ranges using ECT Algorithm
-------------------------------------------------------------------------------
1-99      low-path-id
100-100   high-path-id
101-4095  low-path-id
===============================================================================
```

database

Syntax **database**

Context show>service>id>spb

Description This command displays SPB database information.

Output The following output is an example of service SPB database information.

Sample Output

```
*A:Dut-A# show service id 100001 spb database
===============================================================================
ISIS Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------

Displaying Level 1 database
-------------------------------------------------------------------------------
Dut-A.00-00                           0xc       0xbaba   1103     L1
```

```
Dut-B.00-00                              0x13      0xe780   1117     L1
Dut-C.00-00                              0x13      0x85a    1117     L1
Dut-D.00-00                              0xe       0x174a   1119     L1
Level (1) LSP Count : 4
===============================================================================
```

fate-sharing

Syntax **fate-sharing**

Context show>service>id>spb

Description This command displays SPB fate-sharing information on User B-VPLS service, in correspond to associated Control B-VPLS service.

Output The following output is an example of service SPB fate sharing information.

Sample Output

```
*A:Dut-A# Node show service id spb fate-sharing
===============================================================================
User service fate-shared sap/sdp-bind information
===============================================================================
Control   Control Sap/             FID       User      User Sap/
SvcId     SdpBind                            SvcId     SdpBind
-------------------------------------------------------------------------------
500       1/1/20:500               502       502       1/1/20:502
===============================================================================
```

fdb

Syntax **fdb**

Context show>service>id>spb

Description This command displays SPB Forwarding database information (FDB).

Output The following output is an example of service SPB FDB information.

Sample Output

```
*A:Dut-A# show service id 100001 spb fdb
===============================================================================
User service FDB information
===============================================================================
MacAddr           UCast Source          State   MCast Source          State
-------------------------------------------------------------------------------
00:10:00:01:00:02 1/2/2:1.1             ok      1/2/2:1.1             ok
00:10:00:01:00:03 1/2/3:1.1             ok      1/2/3:1.1             ok
00:10:00:01:00:04 1/2/2:1.1             ok      1/2/2:1.1             ok
-------------------------------------------------------------------------------
```

```
Entries found: 3
===============================================================================
```

fid

Syntax **fid** [*fid*] **fate-sharing**
fid [*fid*] **user-service**
fid [*fid*] **fdb**
fid [*fid*] **mfib** [**group-mac** *ieee-address*]
fid [*fid*] **mfib** [**isid** *isid*]

Context show>service>id>spb

Description This command displays SPB control service FID information.

Parameters *fid* — A user service FID may be specified. All user service FIDs are displayed if the FID is not specified.

Values 1 to 4095

fate-sharing — Displays fate-sharing information

user-service — Specifies user VPLS information for each control VPLS per forwarding data-base identifier. A user service FID may be specified. All user service FIDs are displayed if the FID is not specified.

fdb — Displays forwarding database (FDB) information

mfib — Displays multicast forwarding information base (MFIB) information

ieee-address — Specifies the 48-bit IEEE 802.3 group MAC address

isid — Specifies the value of ISID of the group MAC address of this entry

Values 0 to 16777215

Output The following output is an example of service SPB FID fare sharing information.

Sample Output

```
*A:Dut-A# show service id 100001 spb fid fate-sharing
===============================================================================
Control service fate-shared sap/sdp-bind information
===============================================================================
Control   Control Sap/             FID       User      User Sap/
SvcId     SdpBind                            SvcId     SdpBind
-------------------------------------------------------------------------------
500       1/1/20:500               502       502       1/1/20:502
===============================================================================

*A:Dut-A# show service id 100001 spb fid fdb
===============================================================================
Control service FDB information
===============================================================================
Fid        MacAddr            UCast Source             MCast Source
```

```
                              Last Update              Last Update
-------------------------------------------------------------------------------
1          00:10:00:01:00:01  local                    local
                              04/04/2012 15:11:24      04/04/2012 15:11:24
1          00:10:00:01:00:02  1/2/2:1.1                1/2/2:1.1
                              04/04/2012 15:51:45      04/04/2012 15:51:45
1          00:10:00:01:00:03  1/2/3:1.1                1/2/3:1.1
                              04/04/2012 15:51:56      04/04/2012 15:51:56
1          00:10:00:01:00:04  1/2/2:1.1                1/2/2:1.1
                              04/04/2012 15:52:11      04/04/2012 15:52:11
-------------------------------------------------------------------------------
Entries found: 4
===============================================================================
*A:Dut-A# show service id 100001 spb fid mfib
===============================================================================
Control service MFIB information
===============================================================================
FID   MacAddr           ISID      Source                   Last Update
-------------------------------------------------------------------------------
1     01:1E:83:00:27:11 10001     1/2/2:1.1                04/04/2012 15:51:45
                                  1/2/3:1.1                04/04/2012 15:51:56
                                  local                    04/04/2012 15:42:44
100   01:1E:83:00:27:12 10002     1/2/2:1.1                04/04/2012 15:51:45
                                  1/2/3:1.1                04/04/2012 15:51:56
                                  local                    04/04/2012 15:43:09
-------------------------------------------------------------------------------
Entries found: 6
===============================================================================
```

hostname

Syntax **hostname**

Context show>service>id>spb

Description This command displays SPB system-id to hostname mapping.

Output The following output is an example of service SPB hostname information.

Sample Output

```
*A:Dut-A# show service id 100001 spb hostname
===============================================================================
Hosts
===============================================================================
System Id               Hostname
-------------------------------------------------------------------------------
0000.00AA.AAAA          cses-B02
0000.00BB.BBBB          cses-B07
===============================================================================
```

interface

Syntax **interface**

Context show>service>id>spb

Description This command displays SPB interfaces.

Output The following output is an example of service SPB interface information.

Sample Output

```
*A:Dut-A# show service id 100001 spb interface
===============================================================================
ISIS Interfaces
===============================================================================
Interface                       Level CircID  Oper State   L1/L2 Metric
-------------------------------------------------------------------------------
sap:1/1/20:500                  L1    65536   Up           10/-
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
```

mfib

Syntax **mfib** [**group-mac** *ieee-address*] [**isid** *isid*]

Context show>service>id>spb

Description This command displays multicast forwarding data-base (MFIB) information.

Parameters *ieee-address* — Specifies a MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

isid — Specifies an I-SID

Values 0 to 16777215

Output The following output is an example of service SPB MFIB information.

Sample Output

```
*A:Dut-A# show service id 100001 spb mfib
===============================================================================
User service MFIB information
===============================================================================
MacAddr           ISID     Status
-------------------------------------------------------------------------------
01:1E:83:00:27:11 10001    Ok
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
```

routes

Syntax **routes**

Context show>service>id>spb

Description This command displays SPB route information.

Output The following output is an example of service SPB route information.

Sample Output

```
*A:Dut-A# show service id 100001 spb routes
===============================================================
MAC Route Table
===============================================================
Fid  MAC                                          Ver.   Metric
      NextHop If                    SysID
---------------------------------------------------------------

Fwd Tree: unicast
---------------------------------------------------------------
1    00:10:00:01:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
1    00:10:00:01:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
1    00:10:00:01:00:04                           10     20
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
100  00:10:00:02:00:04                           10     20
      sap:1/2/3:1.1                Dut-C

Fwd Tree: multicast
---------------------------------------------------------------
1    00:10:00:01:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
1    00:10:00:01:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
1    00:10:00:01:00:04                           10     20
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:02                           10     10
      sap:1/2/2:1.1                Dut-B
100  00:10:00:02:00:03                           10     10
      sap:1/2/3:1.1                Dut-C
100  00:10:00:02:00:04                           10     20
      sap:1/2/3:1.1                Dut-C
---------------------------------------------------------------
No. of MAC Routes: 12
===============================================================

===============================================================
ISID Route Table
===============================================================
Fid  ISID                                        Ver.
```

```
      NextHop If                    SysID
-------------------------------------------------------------
1    10001                                      10
      sap:1/2/2:1.1                Dut-B
      sap:1/2/3:1.1                Dut-C
100  10002                                      10
      sap:1/2/2:1.1                Dut-B
      sap:1/2/3:1.1                Dut-C
-------------------------------------------------------------
No. of ISID Routes: 2
=============================================================
A:Dut-A#  show service id spb fate-sharing
================================================================================
User service fate-shared sap/sdp-bind information
================================================================================
Control   Control Sap/             FID       User      User Sap/
SvcId     SdpBind                            SvcId     SdpBind
--------------------------------------------------------------------------------
500       1/1/20:500               502       502       1/1/20:502
================================================================================
```

spf

Syntax **spf**

Context show>service>id>spb

Description This command displays SPF information.

Output The following output is an example of service SPB SPF information.

Sample Output

```
A:cses-B01#  show service id spb spf
================================================================================
Path Table
================================================================================
Node                            Interface                 Nexthop
--------------------------------------------------------------------------------

Fwd Tree: unicast,   ECT Alg: low-path-id
--------------------------------------------------------------------------------
cses-B07.00                     sap:1/1/20:500             cses-B07
cses-B01.00                  sap:1/1/20:500             cses-B07
cses-B07.00                     sap:1/1/20:500             cses-B07

Fwd Tree: unicast,   ECT Alg: high-path-id
--------------------------------------------------------------------------------
cses-B07.00                     sap:1/1/20:500             cses-B07
cses-B01.00                  sap:1/1/20:500             cses-B07
cses-B07.00                     sap:1/1/20:500             cses-B07

Fwd Tree: multicast,   ECT Alg: low-path-id
--------------------------------------------------------------------------------
cses-B07.00                     sap:1/1/20:500             cses-B07
```

```
cses-B01.00                  sap:1/1/20:500             cses-B07
cses-B07.00                     sap:1/1/20:500             cses-B07

Fwd Tree: multicast,   ECT Alg: high-path-id
-------------------------------------------------------------------------------
cses-B07.00                     sap:1/1/20:500             cses-B07
cses-B01.00                  sap:1/1/20:500             cses-B07
cses-B07.00                     sap:1/1/20:500             cses-B07
===============================================================================
```

spf-log

Syntax **spf-log**

Context show>service>id>spb

Description This command displays SPF Log information.

Output The following output is an example of service SPB SPF log information.

Sample Output

```
A:cses-B01#  show service id spb spf-log
===============================================================================
ISIS SPF Log
===============================================================================
When                 Duration       L1 Nodes   L2 Nodes   Event Count      Type
-------------------------------------------------------------------------------
07/23/2012 16:01:13  <0.01s         1          0          1                Reg
07/23/2012 16:01:19  <0.01s         1          0          4                Reg
07/23/2012 16:01:24  <0.01s         3          0          2                Reg
07/23/2012 16:01:29  <0.01s         4          0          1                Reg
-------------------------------------------------------------------------------
Log Entries : 4
-------------------------------------------------------------------------------
```

statistics

Syntax **statistics**

Context show>service>id>spb

Description This command displays SPB statistics.

Output The following output is an example of service SPB statistics information.

Sample Output

```
A:cses-B01#  show service id spb statistics
===============================================================================
ISIS Statistics
```

```
===============================================================================
ISIS Instance     : 1024                    SPF Runs       : 4
Purge Initiated   : 0                       LSP Regens.    : 11

CSPF Statistics
Requests          : 0                       Request Drops  : 0
Paths Found       : 0                       Paths Not Found: 0

-------------------------------------------------------------------------------
PDU Type   Received   Processed  Dropped    Sent       Retransmitted
-------------------------------------------------------------------------------
LSP        31         31         0          9          0
IIH        532        532        0          533        0
CSNP       479        479        0          479        0
PSNP       9          9          0          27         0
Unknown    0          0          0          0          0
===============================================================================
```

status

Syntax **status**

Context show>service>id>spb

Description This command displays SPB status.

Output The following output is an example of service SPB status information.

Sample Output

```
A:cses-B01#  show service id spb status
===============================================================================
ISIS Status
===============================================================================
System Id            : 0000.00AA.AAAA
Admin State          : Up
Oper State           : Up
SPB Routing          : Enabled
Last Enabled         : 07/23/2012 16:01:06
Level Capability     : L1
Authentication Check : True
Authentication Type  : None
CSNP-Authentication  : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication  : Enabled
Overload-On-Boot Tim*: 0
LSP Lifetime         : 1200
LSP Wait             : 5 sec (Max)   0 sec (Initial)   1 sec (Second)
LSP MTU Size         : 1492  (Config)  1492  (Oper)
Adjacency Check      : loose
L1 Auth Type         : none
L1 CSNP-Authenticati*: Enabled
L1 HELLO-Authenticat*: Enabled
L1 PSNP-Authenticati*: Enabled
L1 Preference        : 15
```

```
L1 Ext. Preference   : 160
L1 Wide Metrics      : Enabled
L1 LSDB Overload     : Disabled
L1 LSPs              : 4
L1 Default Metric    : 10
L1 IPv6 Def Metric   : 10
Last SPF             : 07/23/2012 16:01:29
SPF Wait             : 10 sec (Max)   1000 ms (Initial)   1000 ms (Second)
Multi-topology       : Disabled
Area Addresses       : 00
Total Exp Routes(L1) : 0
IID TLV              : Disabled
All-L1-MacAddr       : 01:80:c2:00:00:14
===============================================================================
```

4.5.2.2 PBB Clear Commands

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[*:vc-id*] {**all** | **counters** |**stp** | **mrp**}

Context clear>service>statistics>id

Description This command clears the statistics for a particular mesh SDP bind.

Parameters *sdp-id* — Specifies the SDP ID for which to display information

Default All SDPs.

Values 1 to 17407

vc-id — Displays information about the virtual circuit identifier.

Values 1 to 4294967295

all — Clears all queue statistics and STP statistics associated with the SDP

counters — Clears all queue statistics associated with the SDP

stp — Clears all STP statistics associated with the SDP

mrp — Clears all MRP statistics associated with the SDP

mrp

Syntax **mrp**

Context clear>service>statistics>id

Description This command clears all MRP statistics for the service ID.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] {**all** | **counters** | **stp**| **l2pt** | **mrp**}

Context clear>service>statistics>id

Description This command clears statistics for the spoke-SDP bound to the service.

Parameters *sdp-id* — The spoke-SDP ID for which to clear statistics

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

all — Clears all queue statistics and STP statistics associated with the SDP

counters — Clears all queue statistics associated with the SDP

stp — Clears all STP statistics associated with the SDP

l2pt — Clears all L2PT statistics associated with the SDP

mrp — Clears all MRP statistics associated with the SDP

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp**| **l2pt** | **mrp**}

Context clear>service>statistics>id

Description This command clears statistics for the SAP.

Parameters *sap-id* — The SAP ID for which to clear statistics

all — Clears all queue statistics and STP statistics associated with the SAP

counters — Clears all queue statistics associated with the SAP

stp — Clears all STP statistics associated with the SAP

l2pt — Clears all L2PT statistics associated with the SAP

mrp — Clears all MRP statistics associated with the SAP

stp

Syntax **stp**

Context clear>service>statistics>id

Description Clears all spanning tree statistics for the service ID.

4.5.2.3 PBB Debug Commands

mrp

Syntax [**no**] **mrp**

Context debug>service>id

Description This command enables and configures MRP debugging.

all-events

Syntax **all-events**

Context debug>service>id>mrp

Description This command enables MRP debugging for the applicant, leave all, periodic and registrant state machines and enables debugging of received and transmitted MRP PDUs.

applicant-sm

Syntax [**no**] **applicant-sm**

Context debug>service>id>mrp

Description This command enables debugging of the applicant state machine.

The **no** form of the command disables debugging of the applicant state machine.

leave-all-sm

Syntax [**no**] **leave-all-sm**

Context debug>service>id>mrp

Description This command enables debugging of the leave all state machine.

The **no** form of the command disables debugging of the leave all state machine.

mmrp-mac

Syntax [**no**] **mmrp-mac** *ieee-address*

Context debug>service>id>mrp

Description This command filters debug events and only shows events related to the MAC address specified.

The **no** form of the command removes the debug filter.

Parameters *ieee-address* — xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)

mrpdu

Syntax [**no**] **mrpdu**

Context debug>service>id>mrp

Description This command enables debugging of the MRP PDUs that are received or transmitted.

The **no** form of the command disables debugging of MRP PDUs.

periodic-sm

Syntax [**no**] **periodic-sm**

Context debug>service>id>mrp

Description This command enables debugging of the periodic state machine.

The **no** form of the command disables debugging of the periodic state machine.

registrant-sm

Syntax [**no**] **registrant-sm**

Context debug>service>id>mrp

Description This command enables debugging of the registrant state machine.

The **no** form of the command disables debugging of the registrant state machine.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>mrp

Description This command filters debug events and only shows events for the particular SAP.

The **no** form of the command removes the debug filter.

Parameters *sap-id* — The SAP ID.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id>mrp

Description This command filters debug events and only shows events for the particular SDP.

The **no** form of the command removes the debug filter.

Parameters *sdp-id* — Specifies the SDP ID for which to display information

Default All SDPs

Values 1 to 17407

vc-id — Displays information about the virtual circuit identifier.

Values 1 to 4294967295

5 Ethernet Virtual Private Networks (EVPNs)

5.1 Overview and EVPN Applications

EVPN is an IETF technology per RFC 7432, *BGP MPLS-Based Ethernet VPN,* that uses a new BGP address family and allows VPLS services to be operated as IP-VPNs, where the MAC addresses and the information to set up the flooding trees are distributed by BGP.

EVPN is defined to fill the gaps of other L2VPN technologies such as VPLS. The main objective of the EVPN is to build E-LAN services in a similar way to RFC 4364 IP-VPNs, while supporting MAC learning within the control plane (distributed by MP-BGP), efficient multi-destination traffic delivery, and active-active multi-homing.

EVPN can be used as the control plane for different data plane encapsulations. The Nokia implementation supports the following data planes:

- **EVPN for VXLAN overlay tunnels (EVPN-VXLAN)**

  EVPN for VXLAN overlay tunnels (EVPN-VXLAN), being the Data Center Gateway (DC GW) function the main application for this feature. In such application VXLAN is expected within the Data Center and VPLS sdp-bindings or SAPs are expected for the connectivity to the WAN. R-VPLS and VPRN connectivity to the WAN is also supported.

  The EVPN-VXLAN functionality is standardized in IETF Draft *draft-ietf-bess-evpn-overlay*.

- **EVPN for MPLS tunnels (EVPN-MPLS)**

  EVPN for MPLS tunnels (EVPN-MPLS), where PEs are connected by any type of MPLS tunnel. EVPN-MPLS is generally used as an evolution for VPLS services in the WAN, being Data Center Interconnect one of the main applications.

  The EVPN-MPLS functionality is standardized in RFC 7432.

- **EVPN for PBB over MPLS tunnels (PBB-EVPN),**

  PEs are connected by PBB over MPLS tunnels in this data plane. It is usually used for large scale E-LAN and E-Line services in the WAN.

  The PBB-EVPN functionality is standardized in RFC 7623.

The 7750 SR, 7450 ESS, or 7950 XRS EVPN VXLAN implementation is integrated in the Nuage Data Center architecture, where the router serves as the DC GW.

Refer to the *Nuage Networks Virtualized Service Platform Guide* for more information about the Nuage Networks architecture and products. The following sections describe the applications supported by EVPN in the 7750 SR, 7450 ESS, or 7950 XRS implementation.

5.1.1 EVPN for VXLAN Tunnels in a Layer 2 DC GW (EVPN-VXLAN)

Figure 138 shows the use of EVPN for VXLAN overlay tunnels on the 7750 SR, 7450 ESS, or 7950 XRS when it is used as a Layer 2 DC GW.

Figure 138 **Layer 2 DC PE with VPLS to the WAN**

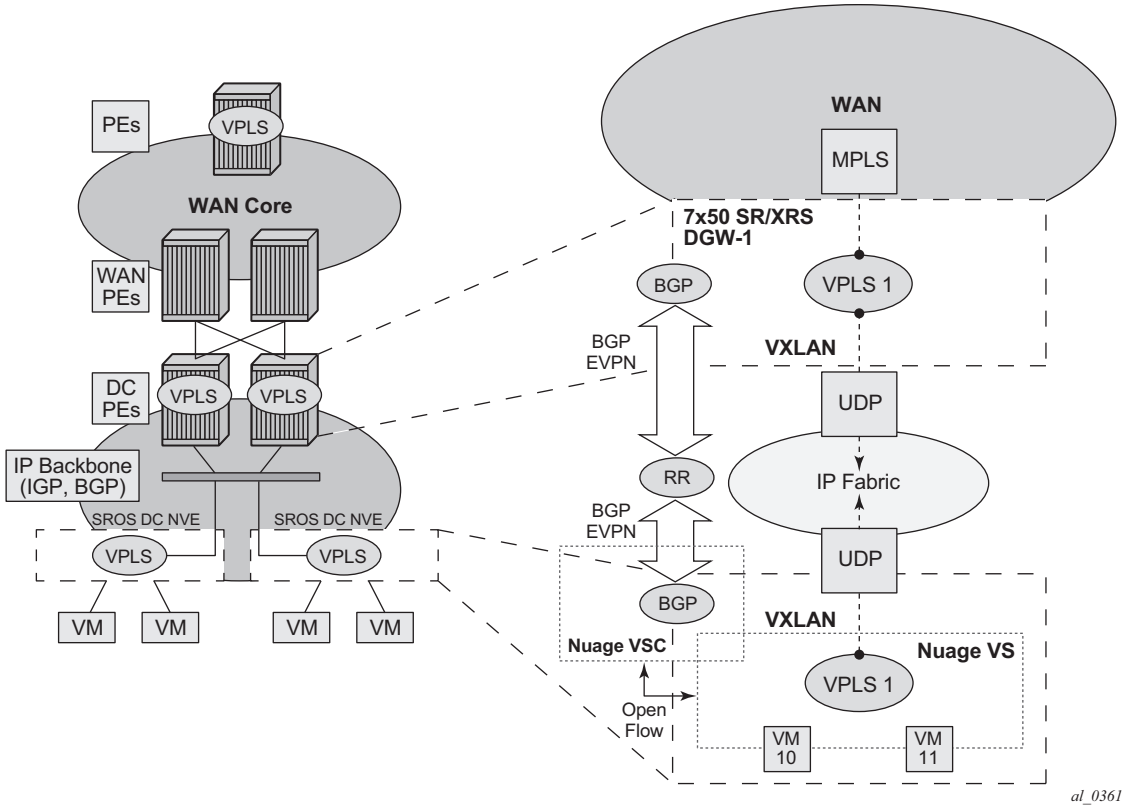

DC providers require a DC GW solution that can extend tenant subnets to the WAN. Customers can deploy the NVO3-based solutions in the DC, where EVPN is the standard control plane and VXLAN is a predominant data plane encapsulation. The Nokia DC architecture (Nuage) uses EVPN and VXLAN as the control and data plane solutions for Layer 2 connectivity within the DC and so does the SR OS.

While EVPN VXLAN will be used within the DC, most service providers use VPLS and H-VPLS as the solution to extend Layer 2 VPN connectivity. Figure 138 shows the Layer 2 DC GW function on the 7750 SR, 7450 ESS, and 7950 XRS routers, providing VXLAN connectivity to the DC and regular VPLS connectivity to the WAN.

The WAN connectivity will be based on VPLS where SAPs (null, dot1q, and qinq), spoke-SDPs (FEC type 128 and 129), and mesh-SDPs are supported.

The DC GWs can provide multi-homing resiliency through the use of BGP multi-homing.

EVPN-MPLS can also be used in the WAN. In this case, the Layer 2 DC GW function provides translation between EVPN-VXLAN and EVPN-MPLS. EVPN multi-homing can be used to provide DC GW redundancy.

5.1.2 EVPN for VXLAN Tunnels in a Layer 2 DC with Integrated Routing Bridging Connectivity on the DC GW

Figure 139 shows the use of EVPN for VXLAN overlay tunnels on the 7750 SR, 7450 ESS, or 7950 XRS when the DC provides Layer 2 connectivity and the DC GW can route the traffic to the WAN through an R-VPLS and linked VPRN.

Figure 139 **GW IRB on the DC PE for an L2 EVPN/VXLAN DC**

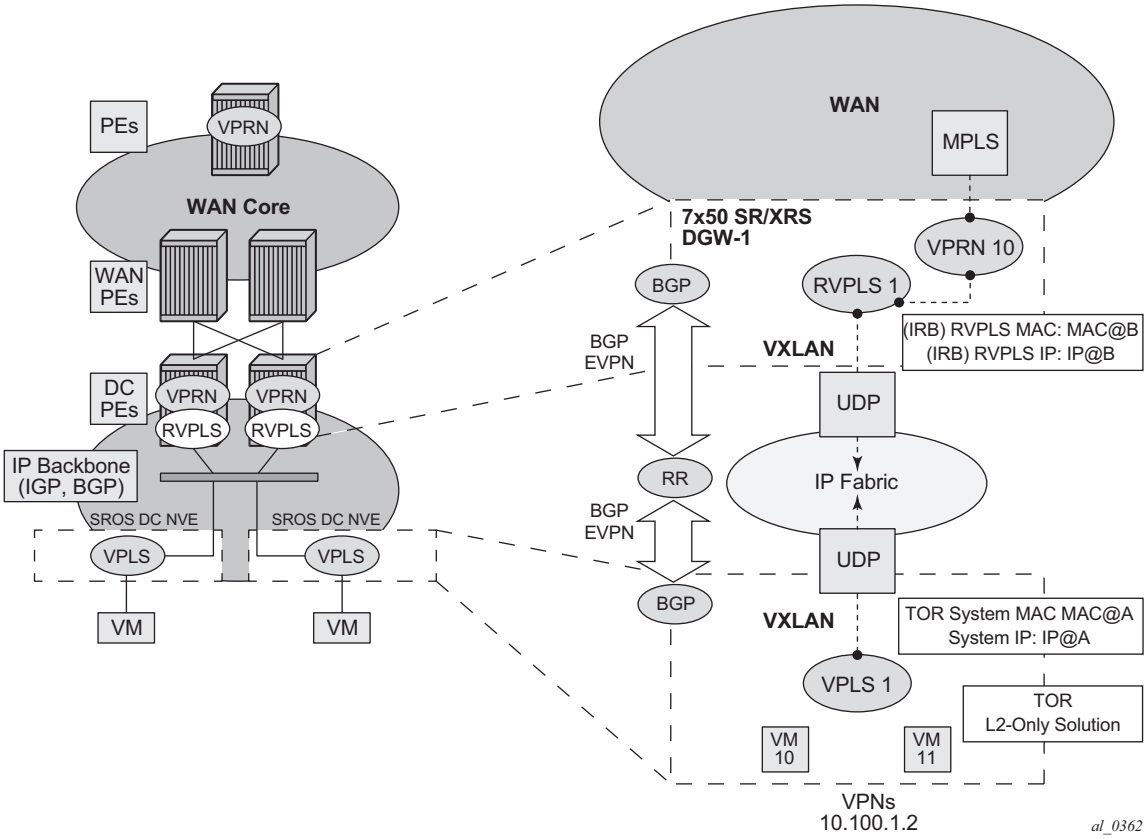

In some cases, the DC GW must provide a Layer 3 default gateway function to all the hosts in a specified tenant subnet. In this case, the VXLAN data plane will be terminated in an R-VPLS on the DC GW, and connectivity to the WAN will be accomplished through regular VPRN connectivity. The 7750 SR, 7450 ESS, and 7950 XRS support IPv4 and IPv6 interfaces as default gateways in this scenario.

5.1.3 EVPN for VXLAN Tunnels in a Layer 3 DC with Integrated Routing Bridging Connectivity among VPRNs

Figure 140 shows the use of EVPN for VXLAN tunnels on the 7750 SR, 7450 ESS, or 7950 XRS when the DC provides distributed Layer 3 connectivity to the DC tenants.

Figure 140 **GW IRB on the DC PE for an L3 EVPN/VXLAN DC**

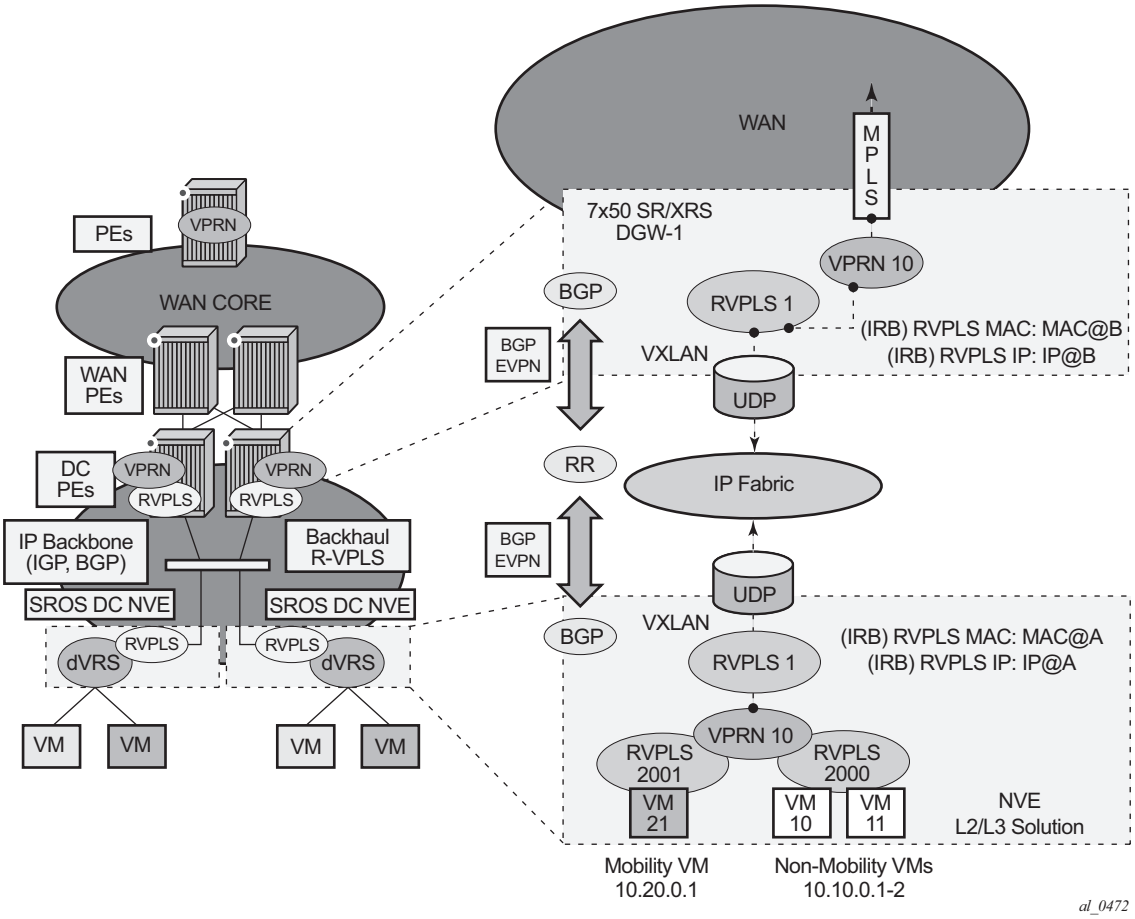

Each tenant will have several subnets for which each DC Network Virtualization Edge (NVE) provides intra-subnet forwarding. An NVE may be a Nuage VSG, VSC/VRS, or any other NVE in the market supporting the same constructs, and each subnet normally corresponds to an R-VPLS. For example, in Figure 140, subnet 10.20.0.0 corresponds to R-VPLS 2001 and subnet 10.10.0.0 corresponds to R-VPLS 2000. In this example, the NVE provides inter-subnet forwarding too, by connecting all the local subnets to a VPRN instance. When the tenant requires L3 connectivity to the IP-VPN in the WAN, a VPRN is defined in the DC GWs, which connects the tenant to the WAN. That VPRN instance will be connected to the VPRNs in the NVEs by means of an IRB (Integrated Routing and Bridging) backhaul R-VPLS. This IRB backhaul R-VPLS provides a scalable solution because it allows L3 connectivity to the WAN without the need for defining all of the subnets in the DC GW.

The 7750 SR, 7450 ESS (in mixed mode), and 7950 XRS DC GW support the IRB backhaul R-VPLS model, where the R-VPLS runs EVPN-VXLAN and the VPRN instances exchange IP prefixes (IPv4 and IPv6) through the use of EVPN. Interoperability between the EVPN and IP-VPN for IP prefixes is also fully supported.

5.1.4 EVPN for VXLAN Tunnels in a Layer 3 DC with EVPN-Tunnel Connectivity among VPRNs

Figure 141 shows the use of EVPN for VXLAN tunnels on the 7750 SR, 7450 ESS (in mixed mode), or 7950 XRS, when the DC provides distributed Layer 3 connectivity to the DC tenants and the VPRN instances are connected through EVPN tunnels.

Figure 141 **EVPN-Tunnel GW IRB on the DC PE for an L3 EVPN/VXLAN DC**

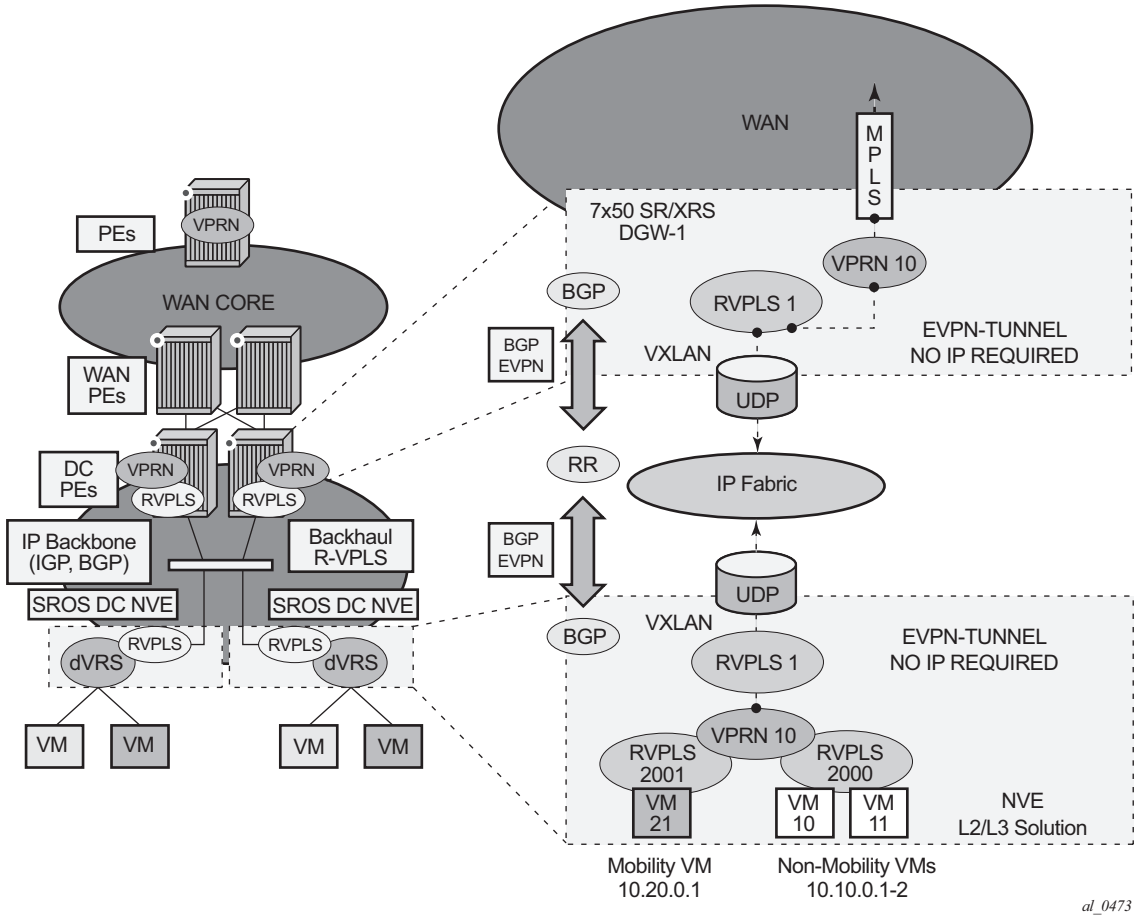

The solution described in section EVPN for VXLAN Tunnels in a Layer 3 DC with Integrated Routing Bridging Connectivity among VPRNs provides a scalable IRB backhaul R-VPLS service where all the VPRN instances for a specified tenant can be connected by using IRB interfaces. When this IRB backhaul R-VPLS is exclusively used as a backhaul and does not have any SAPs or SDP-bindings directly attached, the solution can be optimized by using EVPN tunnels.

EVPN tunnels are enabled using the **evpn-tunnel** command under the R-VPLS interface configured on the VPRN. EVPN tunnels provide the following benefits to EVPN-VXLAN IRB backhaul R-VPLS services:

- Easier provisioning of the tenant service. If an EVPN tunnel is configured in an IRB backhaul R-VPLS, there is no need to provision the IRB IPv4 addresses on the VPRN. This makes the provisioning easier to automate and saves IP addresses from the tenant space.

Note: IPv6 interfaces do not require the provisioning of an IPv6 Global Address; a Link Local Address is automatically assigned to the IRB interface.

- Higher scalability of the IRB backhaul R-VPLS. If EVPN tunnels are enabled, multicast traffic is suppressed in the EVPN-VXLAN IRB backhaul R-VPLS service (it is not required). As a result, the number of VXLAN binds in IRB backhaul R-VPLS services with EVPN-tunnels can be much higher.

This optimization is fully supported by the 7750 SR, 7450 ESS (in mixed mode), and 7950 XRS.

5.1.5 EVPN for MPLS Tunnels in E-LAN Services

Figure 142 shows the use of EVPN for MPLS tunnels on the 7750 SR, 7450 ESS, and 7950 XRS. In this case, EVPN is used as the control plane for E-LAN services in the WAN.

Figure 142 **EVPN for MPLS in VPLS Services**

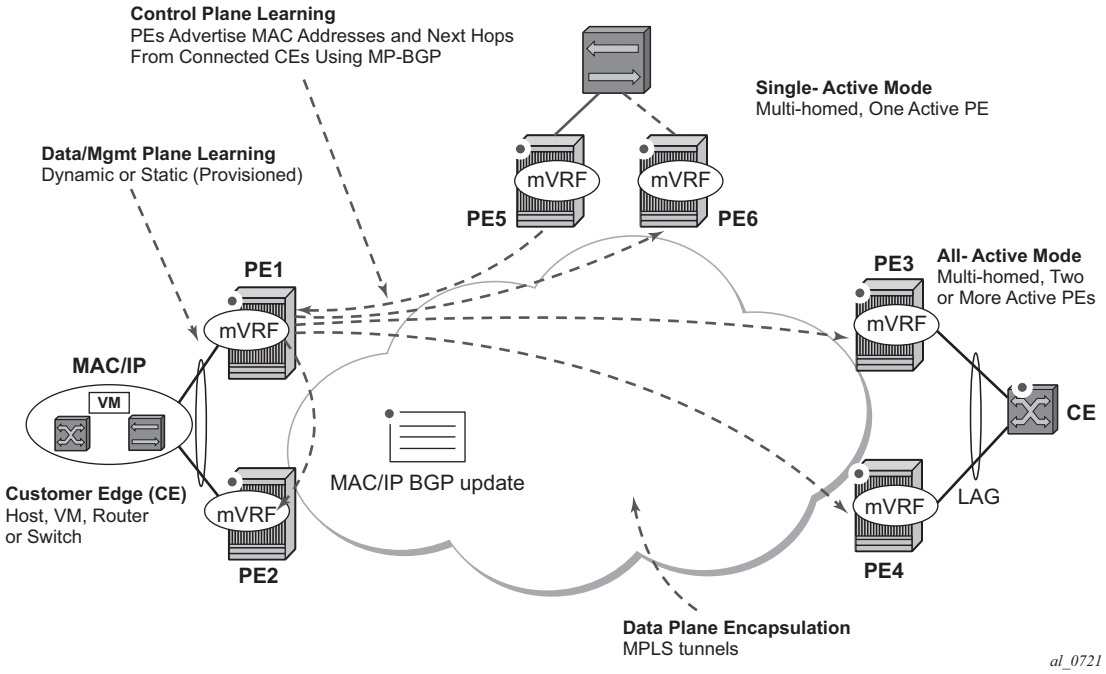

EVPN-MPLS is standardized in RFC 7432 as an L2VPN technology that can fill the gaps in VPLS for E-LAN services. A significant number of service providers offering E-LAN services today are requesting EVPN for their multi-homing capabilities, as well as the optimization EVPN provides. EVPN supports all-active multi-homing (per-flow load-balancing multi-homing) as well as single-active multi-homing (per-service load-balancing multi-homing).

EVPN is a standard-based technology that supports all-active multi-homing, and although VPLS already supports single-active multi-homing, EVPN's single-active multi-homing is perceived as a superior technology due to its mass-withdrawal capabilities to speed up convergence in scaled environments.

EVPN technology provides a number of significant benefits, including:

- superior multi-homing capabilities
- an IP-VPN-like operation and control for E-LAN services
- reduction and (in some cases) suppression of the BUM (broadcast, Unknown unicast, and Multicast) traffic in the network
- simple provision and management
- new set of tools to control the distribution of MAC addresses and ARP entries in the network

The SR OS EVPN-MPLS implementation is compliant with RFC 7432.

EVPN-MPLS can also be enabled in R-VPLS services with the same feature-set that is described for VXLAN tunnels in sections EVPN for VXLAN Tunnels in a Layer 3 DC with Integrated Routing Bridging Connectivity among VPRNs and EVPN for VXLAN Tunnels in a Layer 3 DC with EVPN-Tunnel Connectivity among VPRNs.

5.1.6 EVPN for MPLS Tunnels in E-Line Services

The MPLS network used by EVPN for E-LAN services can also be shared by E-Line services using EVPN in the control plane. EVPN for E-Line services (EVPN-VPWS) is a simplification of the RFC 7432 procedures, and is supported on the 7750 SR, 7450 ESS, and 7950 XRS in compliance with IETF Draft *draft-ietf-bess-evpn-vpws*.

5.1.7 EVPN for MPLS Tunnels in E-Tree Services

The MPLS network used by E-LAN and E-Line services can also be shared by Ethernet-Tree (E-Tree) services using the EVPN control plane. EVPN E-Tree services use the EVPN control plane extensions described in IETF Draft *draft-ietf-bess-evpn-etree* and are supported on the 7750 SR, 7450 ESS, and 7950 XRS.

5.1.8 EVPN for PBB over MPLS Tunnels (PBB-EVPN)

Figure 143 shows the use of EVPN for MPLS tunnels on the 7750 SR, 7450 ESS, and 7950 XRS. In this case, EVPN is used as the control plane for E-LAN services in the WAN.

Figure 143 **EVPN for PBB over MPLS**

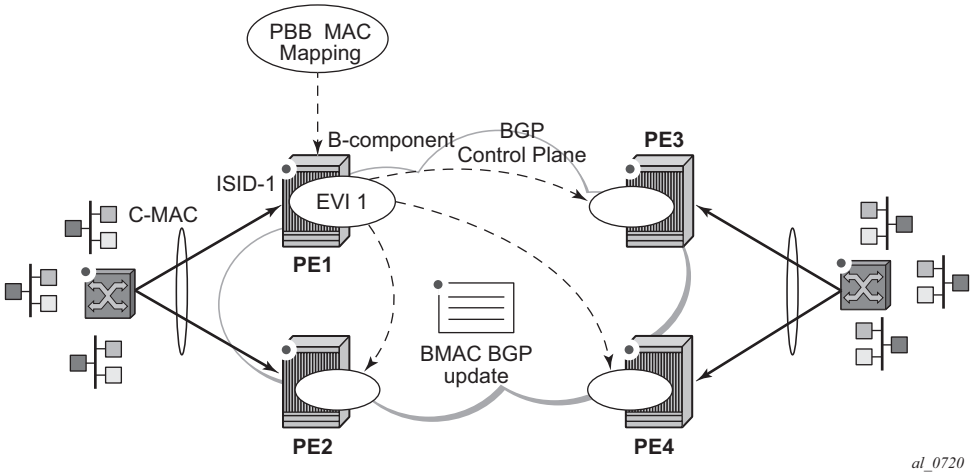

EVPN for PBB over MPLS (hereafter called PBB-EVPN) is specified in RFC 7623. It provides a simplified version of EVPN for cases where the network requires very high scalability and does not need all the advanced features supported by EVPN-MPLS (but still requires single-active and all-active multi-homing capabilities).

PBB-EVPN is a combination of 802.1ah PBB and RFC 7432 EVPN and reuses the PBB-VPLS service model, where BGP-EVPN is enabled in the B-VPLS domain. EVPN is used as the control plane in the B-VPLS domain to control the distribution of BMACs and setup per-ISID flooding trees for I-VPLS services. The learning of the CMACs, either on local SAPs/SDP-bindings or associated with remote BMACs, is still performed in the data plane. Only the learning of BMACs in the B-VPLS is performed through BGP.

The SR OS PBB-EVPN implementation supports PBB-EVPN for I-VPLS and PBB-Epipe services, including single-active and all-active multi-homing.

5.2 EVPN for VXLAN Tunnels and Cloud Technologies

This section provides information about EVPN for VXLAN tunnels and cloud technologies.

5.2.1 Introduction to VXLAN

The SR OS and Nuage solution for DC supports VXLAN (Virtual eXtensible Local Area Network) overlay tunnels as per RFC 7348.

VXLAN addresses the data plane needs for overlay networks within virtualized data centers accommodating multiple tenants. The main attributes of the VXLAN encapsulation are:

- VXLAN is an overlay network encapsulation used to carry MAC traffic between VMs over a logical Layer 3 tunnel.
- Avoids the Layer 2 MAC explosion, because VM MACs are only learned at the edge of the network. Core nodes simply route the traffic based on the destination IP (which is the system IP address of the remote PE or VTEP-VXLAN Tunnel End Point).
- Supports multi-path scalability through ECMP (to a remote VTEP address, based on source UDP port entropy) while preserving the Layer 2 connectivity between VMs. xSTP is no longer needed in the network.
- Supports multiple tenants, each with their own isolated Layer 2 domain. The tenant identifier is encoded in the VNI field (VXLAN Network Identifier) and allows up to 16M values, as opposed to the 4k values provided by the 802.1q VLAN space.

Figure 144 shows an example of the VXLAN encapsulation supported by the Nokia implementation.

Figure 144 **VXLAN Frame Format**

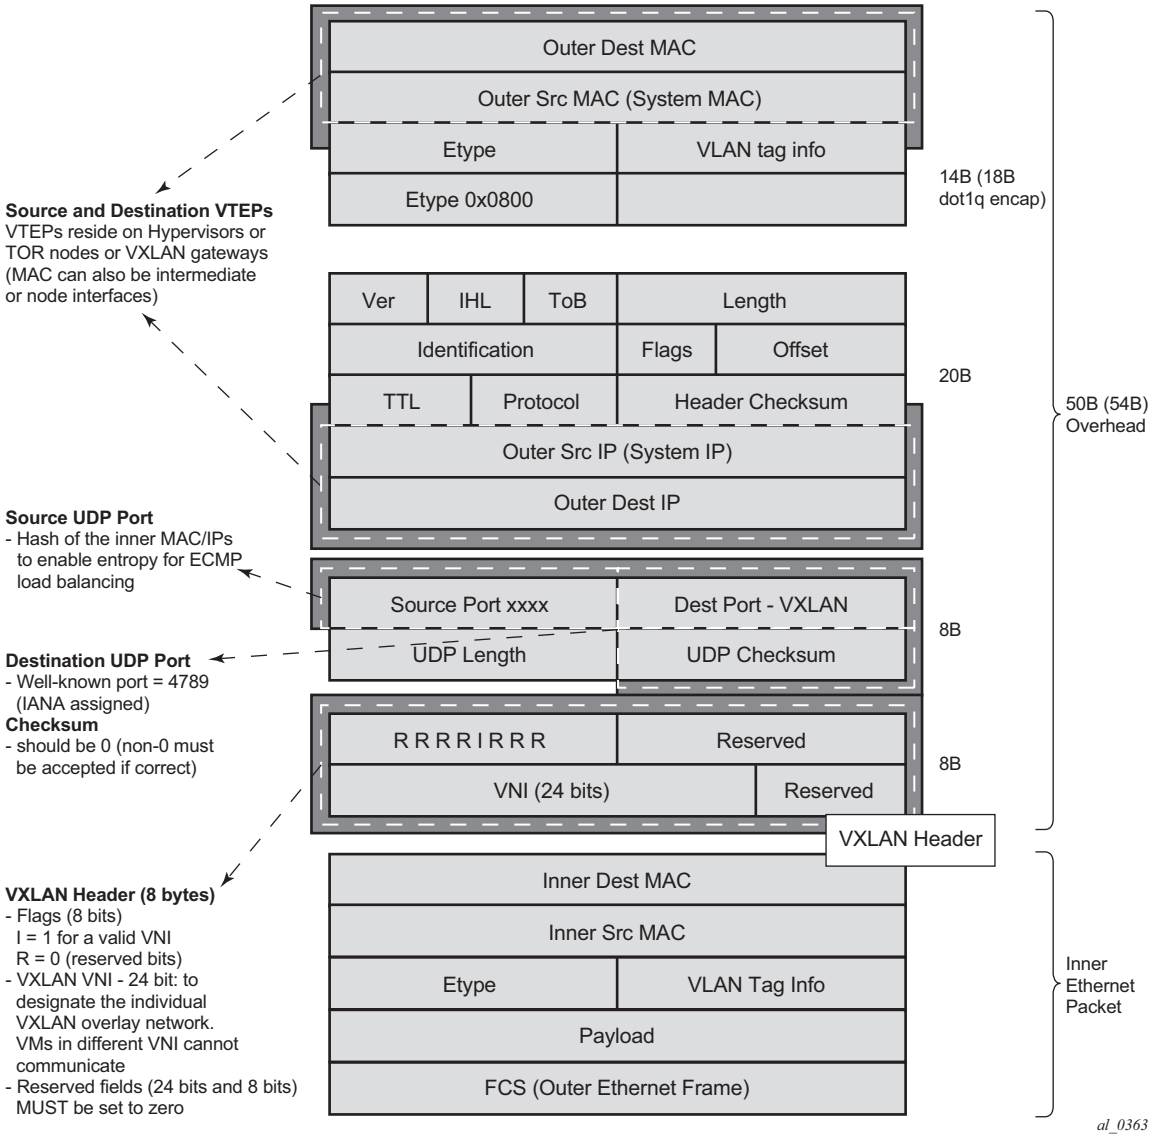

As shown in Figure 144, VXLAN encapsulates the inner Ethernet frames into VXLAN + UDP/IP packets. The main pieces of information encoded in this encapsulation are:

- VXLAN header (8 bytes)
  - Flags (8 bits) where the I flag is set to 1 to indicate that the VNI is present and valid. The rest of the flags (“Reserved” bits) are set to 0.
  - Includes the VNI field (24-bit value) or VXLAN network identifier. It identifies an isolated Layer 2 domain within the DC network.
  - The rest of the fields are reserved for future use.

- UDP header (8 bytes)
  - Where the destination port is a well-known UDP port assigned by IANA (4789).
  - The source port is derived from a hashing of the inner source and destination MAC/IP addresses that the 7750 SR, 7450 ESS, or 7950 XRS does at ingress. This will create an “entropy” value that can be used by the core DC nodes for load balancing on ECMP paths.
  - The checksum will be set to zero.
- Outer IP and Ethernet headers (34 or 38 bytes)
  - The source IP and source MAC will identify the source VTEP. That is, these fields will be populated with the PE’s system IP and chassis MAC address.

Note: The source MAC address will be changed on all the IP hops along the path, as is usual in regular IP routing.

  - The destination IP will identify the remote VTEP (remote system IP) and will be the result of the destination MAC lookup in the service Forwarding Database (FDB).

Note: All remote MACs will be learned by the EVPN BGP and associated with a remote VTEP address and VNI.

Some considerations related to the support of VXLAN on the 7750 SR, 7450 ESS, and 7950 XRS are:

- VXLAN is only supported on network or hybrid ports with null or dot1q encapsulation.
- VXLAN is supported on Ethernet/LAG and POS/APS.
- IPv4 and IPv6 unicast addresses are supported as VTEPs.
- By default, system IP addresses are supported, as VTEPs, for originating and terminating VXLAN tunnels. Non-system IPv4 and IPv6 addresses are supported by using a Forwarding Path Extension (FPE).

5.2.1.1 VXLAN ECMP and LAG

The DC GW supports ECMP load balancing to reach the destination VTEP. Also, any intermediate core node in the Data Center should be able to provide further load balancing across ECMP paths because the source UDP port of each tunneled packet is derived from a hash of the customer inner packet. The following must be considered:

- ECMP for VXLAN is supported on VPLS services, but not for BUM traffic. Unicast spraying will be based on the packet contents.
- ECMP for VXLAN on R-VPLS services is supported for VXLAN IPv6 tunnels.
- ECMP for VXLAN IPv4 tunnels on R-VPLS is only supported if the command **config**>**service**>**vpls**>**allow-ip-int-bind**>**vxlan-ipv4-tep-ecmp** is enabled on the R-VPLS.
- In the above cases where ECMP is not supported (BUM traffic in VPLS and VXLAN IPv4 on R-VPLS if not enabled), each VXLAN binding is tied to a single (different) ECMP path, so in a normal deployment with a reasonable number of remote VTEPs, there should be a fair distribution of the traffic across the paths. In other words, only per-VTEP load-balancing is supported, instead of per-flow load-balancing.
- LAG spraying based on the packet hash is supported in all the cases (VPLS unicast, VPLS BUM, and R-VPLS).

5.2.1.2 VXLAN VPLS Tag Handling

The following describes the behavior on the 7750 SR, 7450 ESS, and 7950 XRS with respect to VLAN tag handling for VXLAN VPLS services:

- Dot1q, QinQ, and null SAPs, as well as regular VLAN handling procedures at the WAN side, are supported on VXLAN VPLS services.
- No “vc-type vlan” like VXLAN VNI bindings are supported. Therefore, at the egress of the VXLAN network port, the router will not add any inner VLAN tag on top of the VXLAN encapsulation, and at the ingress network port, the router will ignore any VLAN tag received and will consider it as part of the payload.

5.2.1.3 VXLAN MTU Considerations

For VXLAN VPLS services, the network port MTU must be at least 50 Bytes (54 Bytes if dot1q) greater than the Service-MTU to allow enough room for the VXLAN encapsulation.

The Service-MTU is only enforced on SAPs, (any SAP ingress packet with MTU greater than the service-mtu will be discarded) and not on VXLAN termination (any VXLAN ingress packet will make it to the egress SAP regardless of the configured service-mtu).

Note: The router will never fragment or reassemble VXLAN packets. In addition, the router always sets the DF (Do not Fragment) flag in the VXLAN outer IP header.

5.2.1.4 VXLAN QoS

VXLAN is a network port encapsulation; therefore, the QoS settings for VXLAN are controlled from the network QoS policies.

5.2.1.4.1 Ingress

The network ingress QoS policy can be applied either to the network interface over which the VXLAN traffic arrives or under *vxlan*/*network*/*ingress* within the EVPN service.

Regardless of where the network QoS policy is applied, the ingress network QoS policy is used to classify the VXLAN packets based on the outer dot1p (if present), then the outer DSCP, to yield an FC/profile.

If the ingress network QoS policy is applied to the network interface over which the VXLAN traffic arrives then the VXLAN unicast traffic uses the network ingress queues configured on FP where the network interface resides. QoS control of BUM traffic received on the VXLAN tunnels is possible by separately redirecting these traffic types to policers within an FP ingress network queue group. This QoS control uses the per forwarding class **fp-redirect-group** parameter together with **broadcast-policer**, **unknown-policer,** and **mcast-policer** within the ingress section of a network QoS policy. This QoS control applies to all BUM traffic received for that forwarding class on the network IP interface on which the network QoS policy is applied.

The ingress network QoS policy can also be applied within the EVPN service by referencing an FP queue group instance, as follows:

```
configure
    service
        vpls <service-id>
            vxlan vni <vni-id>
                network
                    ingress
```

```
qos <network-policy-id>
    fp-redirect-group <queue-group-name>
            instance <instance-id>
```

In this case, the redirection to a specific ingress FP queue group applies as a single entity (per forwarding class) to all VXLAN traffic received only by this service. This overrides the QoS applied to the related network interfaces for traffic arriving on VXLAN tunnels in that service but does not affect traffic received on a spoke-SDP in the same service. It is possible to also redirect unicast traffic to a policer using the per forwarding class **fp-redirect-group policer** parameter, as well as the BUM traffic as above, within the ingress section of a network QoS policy. The use of **ler-use-dscp**, **ip-criteria** and **ipv6-criteria** statements are ignored if configured in the ingress section of the referenced network QoS policy. If the instance of the named queue group template referenced in the **qos** command is not configured on an FP receiving the VXLAN traffic, then the traffic uses the ingress network queues or queue group related to the network interface.

5.2.1.4.2 Egress

On egress, there is no need to specify “remarking” in the policy to mark the DSCP. This is because the VXLAN adds a new IPv4 header, and the DSCP will be always marked based on the egress network qos policy.

5.2.1.5 VXLAN Ping

A new VXLAN troubleshooting tool, VXLAN Ping, is available to verify VXLAN VTEP connectivity. The **VXLAN Ping** command is available from interactive CLI and SNMP.

This tool allows the operator to specify a wide range of variables to influence how the packet is forwarded from the VTEP source to VTEP termination. The ping function requires the operator to specify a different **test-id** (equates to originator handle) for each active and outstanding test. The required local **service** identifier from which the test is launched will determine the source IP (the system IP address) to use in the outer IP header of the packet. This IP address is encoded into the VXLAN header Source IP TLV. The service identifier will also encode the local VNI. The **outer-ip-destination** must equal the VTEP termination point on the remote node, and the **dest-vni** must be a valid VNI within the associated service on the remote node. The outer source IP address is automatically detected and inserted in the IP header of the packet. The outer source IP address uses the IPv4 system address by default.

If the VTEP is created using a non-system source IP address via the **vxlan-src-vtep** command, the outer source IP address uses the address specified by **vxlan-src-vtep**. The remainder of the variables are optional.

The VXLAN PDU will be encapsulated in the appropriate transport header and forwarded within the overlay to the appropriate VTEP termination. The VXLAN router alert (RA) bit will be set to prevent forwarding OAM PDU beyond the terminating VTEP. Since handling of the router alert bit was not defined in some early releases of VXLAN implementations, the VNI Informational bit (I-bit) is set to “0” for OAM packets. This indicates that the VNI is invalid, and the packet should not be forwarded. This safeguard can be overridden by including the **i-flag-on** option that sets the bit to “1”, valid VNI. Ensure that OAM frames meant to be contained to the VTEP are not forwarded beyond its endpoints.

The supporting VXLAN OAM ping draft includes a requirement to encode a reserved IEEE MAC address as the inner destination value. However, at the time of implementation, that IEEE MAC address had not been assigned. The inner IEEE MAC address will default to 00:00:00:00:00:00, but may be changed using the **inner-l2** option. Inner IEEE MAC addresses that are included with OAM packets will not be learned in the local Layer 2 forwarding databases.

The echo responder will terminate the VXLAN OAM frame, and will take the appropriate response action, and include relevant return codes. By default, the response is sent back using the IP network as an IPv4 UDP response. The operator can chose to override this default by changing the **reply-mode** to **overlay**. The overlay return mode will force the responder to use the VTEP connection representing the source IP and source VTEP. If a return overlay is not available, the echo response will be dropped by the responder.

Support is included for:

- IPv4 VTEP
- Optional specification of the outer UDP Source, which helps downstream network elements along the path with ECMP to hash to flow to the same path
- Optional configuration of the inner IP information, which helps the operator test different equal paths where ECMP is deployed on the source. A test will only validate a single path where ECMP functions are deployed. The inner IP information is processed by a hash function, and there is no guarantee that changing the IP information between tests will select different paths.
- Optional end system validation for a single L2 IEEE MAC address per test. This function checks the remote FDB for the configured IEEE MAC Address. Only one end system IEEE MAC Address can be configured per test.
- Reply mode UDP (default) or Overlay

- Optional additional padding can be added to each packet. There is an option that indicates how the responder should handle the pad TLV. By default, the padding will not be reflected to the source. The operator can change this behavior by including **reflect-pad** option. The **reflect-pad** option is not supported when the reply mode is set to UDP.
- Configurable send counts, intervals, times outs, and forwarding class

The VXLAN OAM PDU includes two timestamps. These timestamps are used to report forward direction delay. Unidirectional delay metrics require accurate time of day clock synchronization. Negative unidirectional delay values will be reported as “0.000”. The round trip value includes the entire round trip time including the time that the remote peer takes to process that packet. These reported values may not be representative of network delay.

The following example commands and outputs show how the VXLAN Ping function can be used to validate connectivity. The echo output includes a new header to better describe the VxLAN ping packet headers and the various levels.

```
oam vxlan-ping test-id 1 service 1 dest-vni 2 outer-ip-
destination 10.20.1.4 interval
0.1 send-count 10

TestID 1, Service 1, DestVNI 2, ReplyMode UDP, IFlag Off, PadSize 0, ReflectPad No,
SendCount 10, Interval 0.1, Timeout 5
Outer: SourceIP 10.20.1.3, SourcePort Dynamic, DestIP 10.20.1.4, TTL 10, FC be, Prof
ile
In
Inner: DestMAC 00:00:00:00:00:00, SourceIP 10.20.1.3, DestIP 127.0.0.1

! ! ! ! ! ! ! ! ! !
---- vxlan-id 2 ip-address 10.20.1.4 PING Statistics ----
10 packets transmitted, 10 packets received, 0.00% packet loss
   10 non-errored responses(!), 0 out-of-order(*), 0 malformed echo responses(.)
   0 send errors(.), 0 time outs(.)
   0 overlay segment not found, 0 overlay segment not operational
forward-delay min = 1.097ms, avg = 2.195ms, max = 2.870ms, stddev = 0.735ms
round-trip-delay min = 1.468ms, avg = 1.693ms, max = 2.268ms, stddev = 0.210ms



oam vxlan-ping test-id 2 service 1 dest-vni 2 outer-ip-destination 10.20.1.4 outer-
ip-source-udp 65000 outer-ip-ttl 64 inner-l2 d0:0d:1e:00:00:01 inner-ip-source
 192.168.1.2 inner-ip-destination 127.0.0.8 reply-mode overlay send-
count 20 interval
 1 timeout 3 padding 1000 reflect-pad fc nc profile out

TestID 2, Service 1, DestVNI 2, ReplyMode overlay, IFlag Off, PadSize 1000, ReflectP
ad
 Yes, SendCount 20, Interval 1, Timeout 3
Outer: SourceIP 10.20.1.3, SourcePort 65000, DestIP 10.20.1.4, TTL 64, FC nc, Profil
e
 out
Inner: DestMAC d0:0d:1e:00:00:01, SourceIP 192.168.1.2, DestIP 127.0.0.8
```

```
================================================================================
====================================
rc=1 Malformed Echo Request Received, rc=2 Overlay Segment Not Present, rc=3 Overlay
 Segment Not Operational, rc=4 Ok
================================================================================
====================================

1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=1 ttl=255 rtt-time=1.733ms fwd
-time=0.302ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=2 ttl=255 rtt-time=1.549ms fwd
-time=1.386ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=3 ttl=255 rtt-time=3.243ms fwd
-time=0.643ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=4 ttl=255 rtt-time=1.551ms fwd
-time=2.350ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=5 ttl=255 rtt-time=1.644ms fwd
-time=1.080ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=6 ttl=255 rtt-time=1.670ms fwd
-time=1.307ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=7 ttl=255 rtt-time=1.636ms fwd
-time=0.490ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=8 ttl=255 rtt-time=1.649ms fwd
-time=0.005ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=9 ttl=255 rtt-time=1.401ms fwd
-time=0.685ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=10 ttl=255 rtt-time=1.634ms fwd
-time=0.373ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=11 ttl=255 rtt-time=1.559ms fwd
-time=0.679ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=12 ttl=255 rtt-time=1.666ms fwd
-time=0.880ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=13 ttl=255 rtt-time=1.629ms fwd
-time=0.669ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=14 ttl=255 rtt-time=1.280ms fwd
-time=1.029ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=15 ttl=255 rtt-time=1.458ms fwd
-time=0.268ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=16 ttl=255 rtt-time=1.659ms fwd
-time=0.786ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=17 ttl=255 rtt-time=1.636ms fwd
-time=1.071ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=18 ttl=255 rtt-time=1.568ms fwd
-time=2.129ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=19 ttl=255 rtt-time=1.657ms fwd
-time=1.326ms. rc=4
1132 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=20 ttl=255 rtt-time=1.762ms fwd
-time=1.335ms. rc=4

---- vxlan-id 2 ip-address 10.20.1.4 PING Statistics ----
20 packets transmitted, 20 packets received, 0.00% packet loss
   20 valid responses, 0 out-of-order, 0 malformed echo responses
   0 send errors, 0 time outs
   0 overlay segment not found, 0 overlay segment not operational
forward-delay min = 0.005ms, avg = 0.939ms, max = 2.350ms, stddev = 0.577ms
round-trip-delay min = 1.280ms, avg = 1.679ms, max = 3.243ms, stddev = 0.375ms



oam vxlan-ping test-id 1 service 1 dest-vni 2 outer-ip-destination 10.20.1.4 send
```

```
-count 10 end-system 00:00:00:00:00:01 interval 0.1
TestID 1, Service 1, DestVNI 2, ReplyMode UDP, IFlag Off, PadSize 0, ReflectPad No,
 EndSystemMAC 00:00:00:00:00:01, SendCount 10, Interval 0.1, Timeout 5
Outer: SourceIP 10.20.1.3, SourcePort Dynamic, DestIP 10.20.1.4, TTL 10, FC be, Prof
ile
 In
Inner: DestMAC 00:00:00:00:00:00, SourceIP 10.20.1.3, DestIP 127.0.0.1

2 2 2 2 2 2 2 2 2 2
---- vxlan-id 2 ip-address 10.20.1.4 PING Statistics ----
10 packets transmitted, 10 packets received, 0.00% packet loss
   10 non-errored responses(!), 0 out-of-order(*), 0 malformed echo responses(.)
   0 send errors(.), 0 time outs(.)
   0 overlay segment not found, 0 overlay segment not operational
   0 end-system present(1), 10 end-system not present(2)
forward-delay min = 0.467ms, avg = 0.979ms, max = 1.622ms, stddev = 0.504ms
round-trip-delay min = 1.501ms, avg = 1.597ms, max = 1.781ms, stddev = 0.088ms



oam vxlan-ping test-id 1 service 1 dest-vni 2 outer-ip-destination 10.20.1.4 send
-count 10 end-system 00:00:00:00:00:01
TestID 1, Service 1, DestVNI 2, ReplyMode UDP, IFlag Off, PadSize 0, ReflectPad No,
 EndSystemMAC 00:00:00:00:00:01, SendCount 10, Interval 1, Timeout 5
Outer: SourceIP 10.20.1.3, SourcePort Dynamic, DestIP 10.20.1.4, TTL 10, FC be, Prof
ile
 In
Inner: DestMAC 00:00:00:00:00:00, SourceIP 10.20.1.3, DestIP 127.0.0.1

================================================================================
====================================
rc=1 Malformed Echo Request Received, rc=2 Overlay Segment Not Present, rc=3 Overlay
 Segment Not Operational, rc=4 Ok
mac=1 End System Present, mac=2 End System Not Present
================================================================================
====================================

92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=1 ttl=255 rtt-time=2.883ms fwd
-time=4.196ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=2 ttl=255 rtt-time=1.596ms fwd
-time=1.536ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=3 ttl=255 rtt-time=1.698ms fwd
-time=0.000ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=4 ttl=255 rtt-time=1.687ms fwd
-time=1.766ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=5 ttl=255 rtt-time=1.679ms fwd
-time=0.799ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=6 ttl=255 rtt-time=1.678ms fwd
-time=0.000ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=7 ttl=255 rtt-time=1.709ms fwd
-time=0.031ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=8 ttl=255 rtt-time=1.757ms fwd
-time=1.441ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=9 ttl=255 rtt-time=1.613ms fwd
-time=2.570ms. rc=4 mac=2
92 bytes from vxlan-id 2 10.20.1.4: vxlan_seq=10 ttl=255 rtt-time=1.631ms fwd
-time=2.130ms. rc=4 mac=2

---- vxlan-id 2 ip-address 10.20.1.4 PING Statistics ----
```

```
10 packets transmitted, 10 packets received, 0.00% packet loss
   10 valid responses, 0 out-of-order, 0 malformed echo responses
   0 send errors, 0 time outs
   0 overlay segment not found, 0 overlay segment not operational
   0 end-system present, 10 end-system not present
forward-delay min = 0.000ms, avg = 1.396ms, max = 4.196ms, stddev = 1.328ms
round-trip-delay min = 1.596ms, avg = 1.793ms, max = 2.883ms, stddev = 0.366ms
```

5.2.1.6 IGMP Snooping on VXLAN

The delivery of IP Multicast in VXLAN services can be optimized with IGMP snooping. IGMP snooping is supported in EVPN-VXLAN VPLS services. When enabled, IGMP reports will be snooped on SAPs/SDP-bindings, but also on VXLAN bindings, to create/modify entries in the MFIB for the VPLS service.

The following must be considered when configuring IGMP snooping in EVPN-VXLAN VPLS services:

- There is an additional configuration command to enable IGMP snooping on VXLAN: **config>service>vpls>igmp-snooping no shutdown** will enable the feature in the VPLS service.
- The VXLAN bindings only support basic IGMP snooping functionality. Features configurable under SAPs or SDP-bindings are not available for VXLAN. Since there is no specific IGMP snooping settings for VXLAN bindings (static mrouters or send-queries, and so on.), a specified VXLAN binding will only become a dynamic mrouter when it receives IGMP queries and will add a specified multicast group to the MFIB when it receives an IGMP report for that group.
- The corresponding **show service id igmp-snooping** and **clear service id igmp-snooping** commands are also available for VXLAN bindings. The following CLI commands show how the system displays IGMP snooping information and statistics on VXLAN bindings:

```
*A:PE1# show service id 1 igmp-snooping port-db vxlan vtep  192.0.2.72 vni 1 detail

===============================================================================
IGMP Snooping VXLAN 192.0.2.72/1 Port-DB for service 1
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 232.0.0.1
-------------------------------------------------------------------------------
Mode             : exclude              Type             : dynamic
Up Time          : 0d 19:07:05          Expires          : 137s
Compat Mode      : IGMP Version 3
V1 Host Expires  : 0s                   V2 Host Expires  : 0s
-------------------------------------------------------
Source Address  Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
No sources.
-------------------------------------------------------------------------------
IGMP Group 232.0.0.2
```

```
-------------------------------------------------------------------------------
Mode              : include               Type              : dynamic
Up Time           : 0d 19:06:39           Expires           : 0s
Compat Mode       : IGMP Version 3
V1 Host Expires   : 0s                    V2 Host Expires   : 0s
-------------------------------------------------------
Source Address  Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
10.0.0.232      0d 19:06:39  137s     dynamic  Fwd
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================

*A:PE1# show service id 1 igmp-snooping
statistics vxlan vtep  192.0.2.72 vni 1

===============================================================================
IGMP Snooping Statistics for VXLAN 192.0.2.72/1 (service 1)
===============================================================================
Message Type             Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries          0             0             556
Group Queries            0             0             0
Group-Source Queries     0             0             0
V1 Reports               0             0             0
V2 Reports               0             0             0
V3 Reports               553           0             0
V2 Leaves                0             0             0
Unknown Type             0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0
Rsvd-Scope Packets       : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
Exceeded Max Num Grp Srcs: 0
MCAC Policy Drops        : 0
===============================================================================
*A:PE1# show service id 1 mfib
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         SAP or SDP Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               *                     sap:1/1/1:1              Local    Fwd
*               232.0.0.1             sap:1/1/1:1              Local    Fwd
                                      vxlan:192.0.2.72/1       Local    Fwd
10.0.0.232      232.0.0.2             sap:1/1/1:1              Local    Fwd
```

```
                                      vxlan:192.0.2.72/1       Local     Fwd
-------------------------------------------------------------------------------
Number of entries: 3
===============================================================================
```

5.2.1.7 Static VXLAN Termination in Epipe Services

By default, the system IP address is used to terminate and generate VXLAN traffic. The following configuration example shows an Epipe service that supports static VXLAN termination:

```
config service epipe 1 customer 1 create
  sap 1/1/1:1 create
  exit
  vxlan vni 100 create
    egr-vtep 192.0.2.1
      oper-group op-grp-1
    exit
  no shutdown
```

Where:

- **vxlan vni** *vni* **create** specifies the ingress VNI the router will use to identify packets for the service. The following considerations apply.
  - In services that use EVPN, the configured VNI is only used as the ingress VNI to identify packets that belong to the service. Egress VNIs are learned from the BGP EVPN. In the case of Static VXLAN, the configured VNI is also used as egress VNI (because there is no BGP EVPN control plane).
  - The configured VNI is unique in the system, and as a result, it can only be configured in one service (VPLS or Epipe).
- **egr-vtep** *ip-address* specifies the remote VTEP the router will use when encapsulating frames into VXLAN packets. The following consideration apply.
  - When the PE receives VXLAN packets, the source VTEP is not checked against the configured egress VTEP.
  - The *IP-address* must be present in the global routing table so that the VXLAN destination is operationally up.
- **oper-group** may be added under **egr-vtep**. The expected behavior for the operational group and service status is as follows.
  - If the **egr-vtep** entry is not present in the routing table, the VXLAN destination (in the **show service id vxlan** command) and the provisioned operational group under **egr-vtep** will go into the operationally down state.
  - The service goes down if the Epipe SAP goes down, but it is not affected if the VXLAN destination goes down.

– If the service is **admin shutdown**, then in addition to the SAP, the VXLAN destination and the oper-group will also go into the operationally down state.

Note: The operational group configured under **egr-vtep** cannot be monitored on the SAP of the Epipe where it is configured.

The following features are not supported by Epipe services with VXLAN destinations.

- per-service hashing
- SDP-binds
- PBB context
- BGP-VPWS
- BGP-EVPN
- Spoke-SDP-FEC
- PW-port

5.2.1.8 Non-System IPv4 and IPv6 VXLAN Termination in VPLS, R-VPLS, and Epipe Services

By default, only VXLAN packets with the same IP destination address as the system IPv4 address of the router can be terminated and processed for a subsequent MAC lookup. A router can simultaneously terminate VXLAN tunnels destined for its system IP address and three additional non-system IPv4 or IPv6 addresses, which can be on the base router or VPRN instances. This section describes the configuration requirements for services to terminate VXLAN packets destined for a non-system loopback IPv4 or IPv6 address on the base router or VPRN.

Perform the following steps to configure a service with non-system IPv4 or IPv6 VXLAN termination:

Step 1. Create the FPE (see FPE Creation).

Step 2. Associate the FPE with VXLAN termination (see FPE Association with VXLAN Termination).

Step 3. Configure the router loopback interface (see VXLAN Router Loopback Interface).

Step 4. Configure VXLAN termination (non-system) VTEP addresses (see VXLAN Termination VTEP Addresses).

Step 5. Add the service configuration (see VXLAN Services).

FPE Creation

A Forwarding Path Extension (FPE) is required to terminate non-system IPv4 or IPv6 VXLAN tunnels.

In a non-system IPv4 VXLAN termination, the FPE function is used for additional processing required at ingress (VXLAN tunnel termination) only, and not at egress (VXLAN tunnel origination).

If the IPv6 VXLAN terminates on a VPLS or Epipe service, the FPE function is used at ingress only, and not at egress.

For R-VPLS services terminating IPv6 VXLAN tunnels and also for VPRN VTEPs, the FPE is used for the egress as well as the VXLAN termination function. In the case of R-VPLS, an internal static SDP is created to allow the required extra processing.

See section “Forwarding Path Extension” of the *7450 ESS, 7750 SR, and 7950 XRS Interface Configuration Guide* for information about FPE configuration and functions.

FPE Association with VXLAN Termination

The FPE must be associated with the VXLAN termination application. The following sample configuration shows two FPEs and their corresponding association. FPE 1 uses the base router and FPE 2 is configured for VXLAN termination on VPRN 10.

```
configure
   fwd-path-ext
       fpe 1 create
             path pxc pxc-1
             vxlan-termination
       fpe 2 create
             path pxc pxc-2
             vxlan-termination router 10
```

VXLAN Router Loopback Interface

Create the interface that will terminate and originate the VXLAN packets. The interface is created as a router interface, which is added to the Interior Gateway Protocol (IGP) and used by the BGP as the EVPN NLRI next hop.

Because the system cannot terminate the VXLAN on a local interface address, a subnet must be assigned to the loopback interface and not a host IP address that is /32 or /128. In the following example, all the addresses in subnet 11.11.11.0/24 (except 11.11.11.1, which is the interface IP) and subnet 10.1.1.0/24 (except 10.1.1.1) can be used for tunnel termination. The subnet is advertised using the IGP and is configured on either the base router or a VPRN. In the example, two subnets are assigned, in the base router and VPRN 10 respectively.

```
configure
  router
       interface "lo1"
            loopback
            address 11.11.11.1/24
       isis
            interface "lo1"
                 passive
                 no shutdown

configure
    service
        vprn 10 customer 1 create
            interface "lo1"
                loopback
                address 10.1.1.1/24
            isis
                interface "lo1"
                    passive
                    no shutdown
```

A local interface address cannot be configured as a VXLAN tunnel-termination IP address in the CLI, as shown in the following example.

```
*A:PE-3# configure service system vxlan tunnel-termination 192.0.2.3 fpe 1 create
MINOR: SVCMGR #8353 VXLAN Tunnel termination IP address cannot be configured -
 IP address in use by another application or matches a local interface IP address
```

The subnet can be up to 31 bits. For example, to use 11.11.11.1 as the VXLAN termination address, the subnet should be configured and advertised as shown in the following sample configuration.

```
 interface "lo1"
            address 11.11.11.0/31
            loopback
            no shutdown
        exit
        isis 0
            interface "lo1"
                passive
                no shutdown
            exit
            no shutdown
        exit
```

It is not a requirement for the remote PEs and NVEs to have the specific /32 or /128 IP address in their RTM to resolve the BGP EVPN NLRI next hop or forward the VXLAN packets. An RTM with a subnet that contains the remote VTEP can also perform these tasks.

Note: The system does not check for a pre-existing local base router loopback interface with a subnet corresponding to the VXLAN tunnel termination address. If a tunnel termination address is configured and the FPE is operationally up, the system starts terminating VXLAN traffic and responding ICMP messages for that address. The following conditions are ignored in this scenario:

- the presence of a loopback interface in the base router
- the presence of an interface with the address contained in the configured subnet, and no loopback

The following sample output includes an IPv6 address in the base router. It could also be configured in a VPRN instance.

```
configure
  router
      interface "lo1"
           loopback
           address 11.11.11.1/24
           ipv6
               address 200::/127
           exit
      isis
           interface "lo1"
               passive
               no shutdown
```

VXLAN Termination VTEP Addresses

The **service>system>vxlan>tunnel-termination** context allows the user to configure non-system IP addresses that can terminate the VXLAN and their corresponding FPEs.

As shown in the following example, an IP address may be associated with a new or existing FPE already terminating the VXLAN. The list of addresses that can terminate the VXLAN can include IPv4 and IPv6 addresses.

```
config service system vxlan#
       tunnel-termination 11.11.11.1 fpe 1 create
       tunnel-termination 200::1 fpe 1 create

config service vprn 10 vxlan#
       tunnel-termination 10.1.1.2 fpe 2 create
```

The **tunnel-termination** command creates internal loopback interfaces that can respond to ICMP requests. In the following sample output, an internal loopback is created when the tunnel termination address is added (for 11.11.11.1 and 200::1). The internal FPE router interfaces created by the VXLAN termination function are also shown in the output. Similar loopback and interfaces are created for tunnel termination addresses in a VPRN (not shown).

```
*A:PE1# show router interface
===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::100/64                                                PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::101/64                                                PREFERRED
_tmnx_vli_vxlan_1_131075         Up        Up/Up       Network loopback
   11.11.11.1/32                                               n/a
   200::1/128                                                  PREFERRED
   fe80::6cfb:ffff:fe00:0/64                                   PREFERRED
lo1                              Up        Up/Down     Network loopback
   11.11.11.0/31                                               n/a
system                           Up        Up/Down     Network system
   1.1.1.1/32                                                  n/a
<snip>
```

VXLAN Services

By default, the VXLAN services use the system IP address as the source VTEP of the VXLAN encapsulated frames. The **vxlan-src-vtep** command in the **service>vpls** or **service>epipe** context enables the system to use a non-system IPv4 or IPv6 address as the source VTEP for the VXLAN tunnels in that service.

A different **vxlan-src-vtep** can be used for different services, as shown in the following example where two different services use different non-system IP addresses as source VTEPs.

```
configure service vpls 1
   vxlan-src-vtep 11.11.11.1

configure service vpls 2
   vxlan-src-vtep 200::1
```

In addition, if a **vxlan-src-vtep** is configured and the service uses EVPN, the IP address is also used to set the BGP NLRI next hop in EVPN route advertisements for the service.

Note: The BGP EVPN next hop can be overridden by the use of export policies based on the following rules.

- A BGP peer policy can override a next hop pushed by the **vxlan-src-vtep** configuration.
- If the VPLS service is IPv6 (that is, the **vxlan-src-vtep** is IPv6) and a BGP peer export policy is configured with **next-hop-self**, the BGP next-hop is overridden with an IPv6 address auto-derived from the IP address of the system. The auto-derivation is based on RFC 4291. For example, ::ffff:10.20.1.3 is auto-derived from system IP 10.20.1.3.
- The policy checks the address type of the next hop provided by the **vxlan-src-vtep** command. If the command provides an IPv6 next hop, the policy is unable use an IPv4 address to override the IPv6 address provided by the **vxlan-src-vtep** command.

After the preceding steps are performed to configure a VXLAN termination, the VPLS, R-VPLS, or Epipe service can be used normally, except that the service will terminate VXLAN tunnels with a non-system IPv4 or IPv6 destination address (in the base router or a VPRN instance) instead of the system IP address only.

The FPE **vxlan-termination** function creates internal router interfaces and loopbacks that are displayed by the **show** commands. When configuring IPv6 VXLAN termination on an R-VPLS service, as well as the internal router interfaces and loopbacks, the system will create internal SDP bindings for the required egress processing. The following output shows an example of an internal FPE-type SDP binding created for IPv6 R-VPLS egress processing.

```
*A:PE1# show service sdp-using
===============================================================================
SDP Using
===============================================================================
SvcId      SdpId              Type   Far End              Opr   I.Label E.Label
                                                         State
-------------------------------------------------------------------------------
2002       17407:2002         Fpe    fpe_1.b              Up    262138  262138
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

When BGP EVPN is used, the BGP peer over which the EVPN-VXLAN updates are received can be an IPv4 or IPv6 peer, regardless of whether the next-hop is an IPv4 or IPv6 address.

The same VXLAN tunnel termination address cannot be configured on different router instances; that is, on two different VPRN instances or on a VPRN and the base router.

5.2.2 EVPN for Overlay Tunnels

This section describes the specifics of EVPN for non-MPLS Overlay tunnels.

5.2.2.1 BGP-EVPN Control Plane for VXLAN Overlay Tunnels

The IETF Draft *draft-ietf-bess-evpn-overlay* describes EVPN as the control plane for overlay-based networks. The 7750 SR, 7450 ESS, and 7950 XRS support a subset of the routes and features described in RFC 7432 that are required for the DC GW function. In particular, EVPN-specific multi-homing capabilities are not supported for VXLAN. However, multi-homing can be supported by using regular BGP multi-homing based on the L2VPN BGP address family.

Figure 145 shows the EVPN MP-BGP NLRI, required attributes and extended communities, and two route types supported for the DC GW Layer 2 applications:

- route type 3 – Inclusive Multicast Ethernet Tag route
- route type 2 – MAC/IP advertisement route

Figure 145 **EVPN-VXLAN Required Routes and Communities**

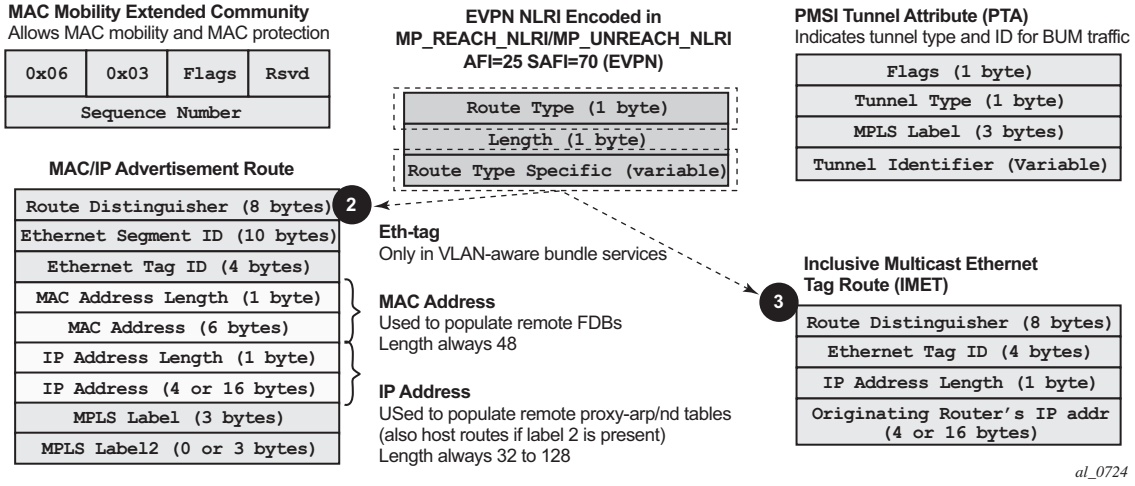

EVPN Route Type 3 – Inclusive Multicast Ethernet Tag Route

Route type 3 is used to set up the flooding tree (BUM flooding) for a specified VPLS service in the data center. The received inclusive multicast routes add entries to the VPLS flood list in the 7750 SR, 7450 ESS, and 7950 XRS. The tunnel types supported in an EVPN route type 3 when BGP-EVPN MPLS is enabled are ingress replication, P2MP MLDP, and composite tunnels.

Ingress Replication (IR) and Assisted Replication (AR) are supported for VXLAN tunnels. See Layer 2 Multicast Optimization for VXLAN (Assisted-Replication) for more information about the AR.

If **ingress-repl-inc-mcast-advertisement** is enabled, a route type 3 is generated by the router per VPLS service as soon as the service is in an operationally up state. The following fields and values are used:

- Route Distinguisher: taken from the RD of the VPLS service within the BGP context

  Note: The RD can be configured or derived from the **bgp-evpn evi** value.
- Ethernet Tag ID: 0
- IP address length: always 32
- Originating router's IP address: carries the system address (IPv4 only)

  Note: By default, the IP address of the Originating router is derived from the system IP address. However, this can be overridden by the **config>service>vpls>bgp-evpn>incl-mcast-orig-ip <ip-address>** command for the Ingress Replication (and mLDP if MPLS is used) tunnel type.
- PMSI Tunnel Attribute (PTA):
  - Tunnel type = Ingress replication (6) or Assisted Replication (10)
    - Flags—Leaf not required.
    - MPLS label—Carries the VNI configured in the VPLS service. Only one VNI can be configured per VPLS service.
    - Tunnel endpoint—Equal to the system IP address.

  As shown in Figure 146, additional flags are used in the PTA when the service is configured for AR.

Figure 146 **PMSI Attribute Flags Field for AR**

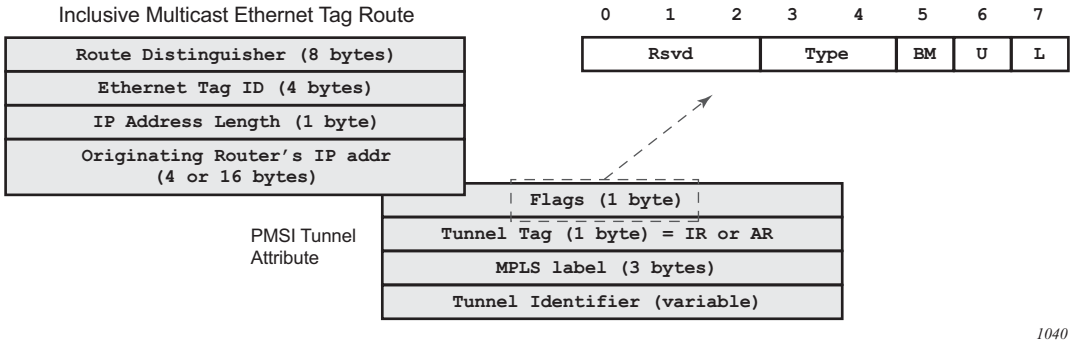

The **Flags** field is defined as a Type field (for AR) with two new flags that are defined as follows:

- T is the AR Type field (2 bits):

- 00 (decimal 0) = RNVE (non-AR support)
- 01 (decimal 1) = AR REPLICATOR
- 10 (decimal 2) = AR LEAF

- The U and BM flags defined in IETF Draft *draft-ietf-bess-evpn-optimized-ir* are not used in the SR OS.

Table 80 describes the inclusive multicast route information sent per VPLS service when the router is configured as **assisted-replication replicator** (AR-R) or **assisted-replication leaf** (AR-L). A Regular Network Virtualization Edge device (RNVE) is defined as an EVPN-VXLAN router that does not support (or is not configured for) Assisted-Replication.

Note: For AR-R, two inclusive multicast routes may be advertised if **ingress-repl-inc-mcast-advertisement** is enabled: a route with tunnel-type IR, tunnel-id = IR IP (generally system-ip) and a route with tunnel-type AR, tunnel-id = AR IP (the address configured in the **assisted-replication-ip** command).

Table 80 **AR-R AND AR-L Routes and Usage**

AR Role	Function	Inclusive Mcast Routes Advertisement
AR-R	Assists AR-LEAFs	• IR included in the Mcast route (uses IR IP) if **ingress-repl-inc-mcast-advertisement** is enabled • AR included in the Mcast route (uses AR IP, tunnel type=AR, T=1)
AR-LEAF	Sends BM only to AR-Rs	IR inclusive multicast route (IR IP, T=2) if **ingress-repl-inc-mcast-advertisement** is enabled
RNVE	Non-AR support	IR inclusive multicast route (IR IP) if **ingress-repl-inc-mcast-advertisement** is enabled

EVPN Route Type 2 – MAC/IP Advertisement Route

The 7750 SR, 7450 ESS, and 7950 XRS will generate this route type for advertising MAC addresses. The router will generate MAC advertisement routes for the following:

- Learned MACs on SAPs or sdp-bindings – if mac-advertisement is enabled
- Conditional static MACs – if mac-advertisement is enabled
- unknown-mac-routes – if unknown-mac-route is enabled, there is no bgp-mh site in the service or there is a (single) DF site

The route type 2 generated by a router uses the following fields and values:

- Route Distinguisher: taken from the RD of the VPLS service within the BGP context

Note: The RD can be configured or derived from the **bgp-evpn** evi value.

- Ethernet Segment Identifier (ESI): Value = 0:0:0:0:0:0:0:0:0:0.
- Ethernet Tag ID: 0.
- MAC address length: always 48
- MAC Address:
  - is 00:00:00:00:00:00 for the Unknown MAC route address.
  - is different from 00:…:00 for the rest of the advertised MACs.
- IP address and IP address length:
  - is the IP address associated with the MAC being advertised with a length of 32 (or 128 for IPv6).
  - if the MAC address is the Unknown MAC route, the IP address length is zero and the IP omitted.
  - in general, any MAC route without IP has IPL=0 (IP length) and the IP is omitted.
  - when received, any IPL value not equal to zero, 32, or 128 will make discard the route.
- MPLS Label 1: carries the VNI configured in the VPLS service. Only one VNI can be configured per VPLS.
- MPLS Label 2: 0
- MAC Mobility extended community: used for signaling the sequence number in case of mac moves and the sticky bit in case of advertising conditional static MACs. If a MAC route is received with a MAC mobility **ext-community**, the sequence number and the sticky bit are considered for the route selection.

When EVPN is used in an IRB backhaul R-VPLS that connects all the VPRN instances for a specified tenant and there is a need to advertise IP prefixes in EVPN, a separate route type is used: route-type 5 IP prefix route.

EVPN Route Type 5 – IP Prefix Route

Figure 147 shows the IP prefix route or route-type 5.

Figure 147 **EVPN Route-Type 5**

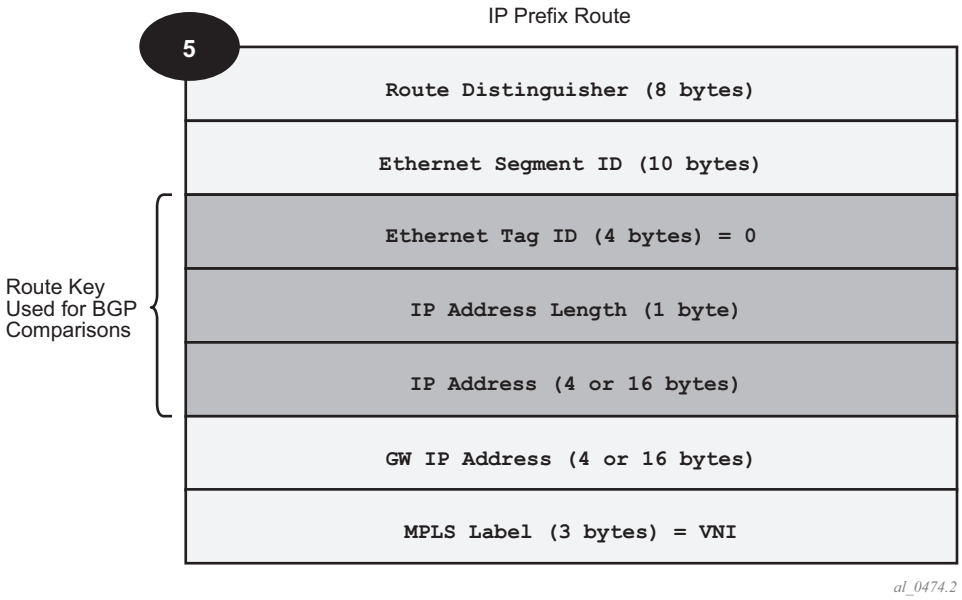

The router will generate this route type for advertising IP prefixes in EVPN. The router will generate IP Prefix advertisement routes for:

- IP prefixes existing in a VPRN linked to the IRB backhaul R-VPLS service.

The route-type 5 generated by a router uses the following fields and values:

- Route Distinguisher: taken from the RD configured in the IRB backhaul R-VPLS service within the BGP context
- Ethernet Segment Identifier (ESI): Value = 0:0:0:0:0:0:0:0:0:0
- Ethernet Tag ID: 0
- IP address length: Any value in the 0 to 128 range
- IP address: any valid IPv4 or IPv6 address
- GW IP address: can carry two different values:
  - if different from zero, the route-type 5 carries the primary IP interface address of the VPRN behind which the IP prefix is known. This is the case for the regular IRB backhaul R-VPLS model.
  - if 0.0.0.0, the route-type 5 is sent with a MAC next-hop extended community that will carry the VPRN interface MAC address. This is the case for the EVPN tunnel R-VPLS model.
- MPLS Label: carries the VNI configured in the VPLS service. Only one VNI can be configured per VPLS service.

All the routes in EVPN-VXLAN will be sent with the RFC 5512 tunnel encapsulation extended community, with the tunnel type value set to VXLAN.

5.2.2.2 EVPN for VXLAN in VPLS Services

The EVPN-VXLAN service is designed around the current VPLS objects and the additional VXLAN construct.

Figure 138 shows a DC with a Layer 2 service that carries the traffic for a tenant who wants to extend a subnet beyond the DC. The DC PE function is carried out by the 7750 SR, 7450 ESS, and 7950 XRS where a VPLS instance exists for that particular tenant. Within the DC, the tenant will have VPLS instances in all the Network Virtualization Edge (NVE) devices where they require connectivity (such VPLS instances can be instantiated in TORs, Nuage VRS, VSG, and so on). The VPLS instances in the redundant DC GW and the DC NVEs will be connected by VXLAN bindings. BGP-EVPN will provide the required control plane for such VXLAN connectivity.

The DC GW routers will be configured with a VPLS per tenant that will provide the VXLAN connectivity to the Nuage VPLS instances. On the router, each tenant VPLS instance will be configured with:

- The WAN-related parameters (saps, spoke-sdps, mesh-sdps, bgp-ad, and so on).
- The BGP-EVPN and VXLAN (VNI) parameters. The following CLI output shows an example for an EVPN-VXLAN VPLS service.

```
*A:DGW1>config>service>vpls# info
----------------------------------------------
            description "vxlan-service"
            vxlan vni 1 create
            exit
            bgp
                route-distinguisher 65001:1
                route-target export target:65000:1 import target:65000:1
            exit
            bgp-evpn
                unknown-mac-route
                mac-advertisement
                vxlan
                    no shutdown
                exit
            sap 1/1/1:1 create
            exit
            no shutdown
----------------------------------------------
```

The bgp-evpn context specifies the encapsulation type (only vxlan is supported) to be used by EVPN and other parameters like the unknown-mac-route and mac-advertisement commands. These commands are typically configured in three different ways:

- **no unknown-mac-route** and **mac-advertisement** (default option) — The router will advertise new learned MACs (on the SAPs or sdp-bindings) or new conditional static MACs.
- **unknown-mac-route** and **no mac-advertisement —** The router will only advertise an unknown-mac-route as long as the service is operationally up (if no BGP-MH site is configured in the service) or the router is the DF (if BGP-MH is configured in the service).
- **unknown-mac-route** and **mac-advertisement —** The router will advertise new learned MACs, conditional static MACs, and the unknown-mac-route. The unknown-mac-route will only be advertised under the preceding described conditions.

Other parameters related to EVPN or VXLAN are:

- Mac duplication parameters
- vxlan vni: Defines the VNI that the router will use in the EVPN routes generated for the VPLS service.

After the VPLS is configured and operationally up, the router will send/receive inclusive multicast Ethernet Tag routes, and a full-mesh of VXLAN connections will be automatically created. These VXLAN “auto-bindings” can be characterized as follows:

- The VXLAN auto-bindings model is based on an IP-VPN-like design, where no SDPs or SDP-binding objects are created by or visible to the user. The VXLAN auto-binds are composed of remote VTEPs and egress VNIs, and can be displayed with the following command:

```
A:DCGW# show service id 1 vxlan
===============================================================================
VPLS VXLAN service Network Specifics
===============================================================================
Ing Net QoS Policy : none                          Vxlan VNI Id     : 1
Ingress FP QGrp    : (none)                        Ing FP QGrp Inst : (none)
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address            Egress VNI  Num. MACs     Mcast   Oper         L2
                                                            State       PBR
-------------------------------------------------------------------------------
192.0.2.71                      1        1        BUM      Up          No
192.0.2.72                      1        0        BUM      Up          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
```

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

- The VXLAN bindings observe the VPLS split-horizon rule. This is performed automatically without the need for any split-horizon configuration.
- BGP Next-Hop Tracking for EVPN is fully supported. If the BGP next-hop for a specified received BGP EVPN route disappears from the routing table, the BGP route will not be marked as “used” and the respective entry in *show service id vxlan* will be removed.

After the flooding domain is setup, the routers and DC NVEs start advertising MAC addresses, and the routers can learn MACs and install them in the FDB. Some considerations are the following:

- All the MAC addresses associated with remote VTEP/VNIs are always learned in the control plane by EVPN. Data plane learning on VXLAN auto-bindings is not supported.
- When **unknown-mac-route** is configured, it will be generated when no (BGP-MH) site is configured, or a site is configured AND the site is DF in the PE.

Note: The **unknown-mac-route** will not be installed in the FDB (therefore, will not show up in the show service id x fdb detail command).

- While the router can be configured with only one VNI (and signals a single VNI per VPLS), it can accept any VNI in the received EVPN routes as long as the route-target is properly imported. The VTEPs and VNIs will show up in the FDB associated with MAC addresses:

```
A:PE65# show service id 1000 fdb detail
===============================================================================
Forwarding Database, Service 1000
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1000      00:00:00:00:00:01 vxlan:                   Evpn     10/05/13 23:25:57
                            192.0.2.63:1063
1000      00:00:00:00:00:65 sap:1/1/1:1000           L/30     10/05/13 23:25:57
1000      00:ca:ca:ca:ca:00 vxlan:                   EvpnS    10/04/13 17:35:43
                            192.0.2.63:1063
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

5.2.2.2.1 Resiliency and BGP Multi-Homing

The DC overlay infrastructure relies on IP tunneling, that is, VXLAN; therefore, the underlay IP layer resolves failure in the DC core. The IGP should be optimized to get the fastest convergence.

From a service perspective, resilient connectivity to the WAN may be provided by BGP-Multi-homing.

5.2.2.2.2 Use of bgp-evpn, bgp-ad, and Sites in the Same VPLS Service

All bgp-evpn (control plane for a VXLAN DC), bgp-ad (control plane for MPLS-based spoke-sdps connected to the WAN), and ONE site for BGP multi-homing (control plane for the multi-homed connection to the WAN) can be configured in one service in a specified system. If that is the case, the following considerations apply:

- The configured BGP route-distinguisher and route-target are used by BGP for the two families, that is, evpn and l2vpn. If different import/export route targets are to be used per family, vsi-import/export policies must be used.
- The pw-template-binding command under BGP, does not have any effect on evpn or bgp-mh. It is only used for the instantiation of the bgp-ad spoke-sdps.
- If the same import/export route-targets are used in the two redundant DC GWs, VXLAN binding as well as a fec129 spoke-sdp binding will be established between the two DGWs, creating a loop. To avoid creating a loop, the router will allow the establishment of an EVPN VXLAN binding and an sdp-binding to the same far-end, but the sdp-binding will be kept operationally down. Only the VXLAN binding will be operationally up.

5.2.2.2.3 Use of the unknown-mac-route

This section describes the behavior of the EVPN-VXLAN service in the router when the unknown-mac-route and BGP-MH are configured at the same time.

The use of EVPN, as the control plane of NVO networks in the DC, provides a significant number of benefits as described in IETF Draft *draft-ietf-bess-evpn-overlay*.

However, there is a potential issue that must be addressed when a VPLS DCI is used for an NVO3-based DC: all the MAC addresses learned from the WAN side of the VPLS must be advertised by BGP EVPN updates. Even if optimized BGP techniques like RT-constraint are used, the number of MAC addresses to advertise or withdraw (in case of failure) from the DC GWs can be difficult to control and overwhelming for the DC network, especially when the NVEs reside in the hypervisors.

The 7750 SR, 7450 ESS, and 7950 XRS solution to this issue is based on the use of an unknown-mac-route address that is advertised by the DC PEs. By using this unknown-mac-route advertisement, the DC tenant may decide to optionally turn off the advertisement of WAN MAC addresses in the DC GW, therefore, reducing the control plane overhead and the size of the FDB tables in the NVEs.

The use of the unknown-mac-route is optional and helps to reduce the amount of unknown-unicast traffic within the data center. All the receiving NVEs supporting this concept will send any unknown-unicast packet to the owner of the unknown-mac-route, as opposed to flooding the unknown-unicast traffic to all other NVEs that are part of the same VPLS.

Note: Although the router can be configured to generate and advertise the unknown-mac-route, the router will never honor the unknown-mac-route and will flood to the TLS-flood list when an unknown-unicast packet arrives at an ingress SAP or SDP-binding.

The use of the unknown-mac-route assumes the following:

- A fully virtualized DC where all the MACs are control-plane learned, and learned previous to any communication (no legacy TORs or VLAN connected servers).
- The only exception is MACs learned over the SAPs/SDP-bindings that are part of the BGP-MH WAN site-id. Only one site-id is supported in this case.
- No other SAPs/SDP-bindings out of the WAN site-id are supported, unless ONLY static MACs are used on those SAPs/SDP-bindings.

Therefore, when unknown-mac-route is configured, it will only be generated when one of the following applies:

- No site is configured and the service is operationally up.
- A BGP-MH site is configured AND the DC GW is Designated Forwarder (DF) for the site. In case of BGP-MH failover, the unknown-mac-route will be withdrawn by the former DF and advertised by the new DF.

5.2.2.3 EVPN for VXLAN in R-VPLS Services

Figure 139 shows a DC with a Layer 2 service that carries the traffic for a tenant who extends a subnet within the DC, while the DC GW is the default gateway for all the hosts in the subnet. The DC GW function is carried out by the 7750 SR, 7450 ESS, and 7950 XRS where an R-VPLS instance exists for that particular tenant. Within the DC, the tenant will have VPLS instances in all the NVE devices where they require connectivity (such VPLS instances can be instantiated in TORs, Nuage VRS, VSG, and so on). The WAN connectivity will be based on existing IP-VPN features.

In this model, the DC GW routers will be configured with a R-VPLS (bound to the VPRN that provides the WAN connectivity) per tenant that will provide the VXLAN connectivity to the Nuage VPLS instances. This model provides inter-subnet forwarding for L2-only TORs and other L2 DC NVEs.

On the router:

- The VPRN will be configured with an interface bound to the backhaul R-VPLS. That interface will be a regular IP interface (IP address configured or possibly a Link Local Address if IPv6 is added).
- The VPRN can support other numbered interfaces to the WAN or even to the DC.
- The R-VPLS will be configured with the BGP, BGP-EVPN and VXLAN (VNI) parameters.

On the Nuage VSGs and NVEs:

- Regular VPLS service model with BGP EVPN and VXLAN parameters.

Other considerations:

- Route-type 2 routes with MACs and IPs will be advertised. Some considerations about MAC+IP and ARP/ND entries are:
  - The 7750 SR will advertise its IRB MAC+IP in a route type 2 route and possibly the VRRP vMAC+vIP if it runs VRRP and the 7750 SR is the master.
    In both cases, the MACs will be advertised as static MACs, therefore, protected by the receiving PEs.
  - If the 7750 SR VPRN interface is configured with one or more additional secondary IP addresses, they will all be advertised in routes type 2, as static MACs.
  - The 7750 SR will process route-type 2 routes as usual, populating the FDB with the received MACs and the VPRN ARP/ND table with the MAC and IPs, respectively.

Note: ND entries received from the EVPN are installed as "Router" entries. The ARP/ND entries coming from the EVPN will be tagged as “EVPN”:

  - When a VPLS containing proxy-ARP/proxy-ND entries is bound to a VPRN (allow-ip-int-bind) all the proxy-ARP/proxy-ND entries are moved to the VPRN ARP/ND table. ARP/ND entries will be also moved to proxy-ARP/ proxy-ND entries if the VPLS is unbound.
  - EVPN will not program EVPN-received ARP/ND entries if the receiving VPRN has no IP addresses for the same subnet. The entries will be added when the IP address for the same subnet is added.
  - Static ARP/ND entries have precedence over dynamic and EVPN ARP/ND entries.

- VPRN interface binding to VPLS service will bring down the VPRN interface operational status, if the VPRN interface mac or the VRRP mac matches a static-mac or OAM mac configured in the associated VPLS service. If that is the case, a trap will be generated.
- Redundancy will be handled by VRRP. The 7750 SR master will advertise vMAC and vIP, as discussed, including the mac mobility extended community and the sticky bit.

5.2.2.3.1 EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes

Figure 140 shows a Layer 3 DC model, where a VPRN is defined in the DC GWs, connecting the tenant to the WAN. That VPRN instance will be connected to the VPRNs in the NVEs by means of an IRB backhaul R-VPLS. Since the IRB backhaul R-VPLS provides connectivity only to all the IRB interfaces and the DC GW VPRN is not directly connected to all the tenant subnets, the WAN ip-prefixes in the VPRN routing table must be advertised in EVPN. In the same way, the NVEs will send IP prefixes in EVPN that will be received by the DC GW and imported in the VPRN routing table.

Note: To generate or process IP prefixes sent or received in EVPN route type 5, the support for IP route advertisement must be enabled in BGP-EVPN. This is performed through the **bgp-evpn>ip-route-advertisement** command. This command s disabled by default and must be explicitly enabled. The command is tied to the **allow-ip-int-bind** command required for R-VPLS.

Local router interface host addresses are not advertised in EVPN by default. To advertise them, the **ip-route-advertisement incl-host** command must be enabled. For example:

```
===============================================================================
Route Table (Service: 2)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                           Active     Metric
-------------------------------------------------------------------------------
10.1.1.0/24                                     Local   Local     00h00m11s  0
       if                                               Y            0
10.1.1.100/32                                   Local   Host      00h00m11s  0
       if                                               Y            0
===============================================================================
```

For the case displayed by the output above, the behavior is the following:

- **ip-route-advertisement** only local subnet (default) - 10.1.1.0/24 is advertised
- **ip-route-advertisement incl-host** local subnet, host - 10.1.1.0/24 and 10.1.1.100/32 are advertised

Below is an example of VPRN (500) with two IRB interfaces connected to backhaul R-VPLS services 501 and 502 where EVPN-VXLAN runs:

```
vprn 500 customer 1 create
            ecmp 4
            route-distinguisher 65072:500
            auto-bind-tunnel
                resolution-filter
                resolution-filter gre ldp rsvp
            vrf-target target:65000:500
            interface "evi-502" create
                address 20.20.20.72/24
                vpls "evpn-vxlan-502"
                exit
            exit
            interface "evi-501" create
                address 10.10.10.72/24
                vpls "evpn-vxlan-501"
                exit
            exit
            no shutdown
vpls 501 customer 1 create
            allow-ip-int-bind
            vxlan vni 501 create
            exit
            bgp
                route-distinguisher 65072:501
                route-target export target:65000:501 import target:65000:501
            exit
            bgp-evpn
                ip-route-advertisement incl-host
                vxlan
                    no shutdown
                exit
            exit
            service-name "evpn-vxlan-501"
            no shutdown
        exit
```

```
vpls 502 customer 1 create
            allow-ip-int-bind
            vxlan vni 502 create
            exit
            bgp
                route-distinguisher 65072:502
                route-target export target:65000:502 import target:65000:502
            exit
            bgp-evpn
                ip-route-advertisement incl-host
                vxlan
                    no shutdown
                exit
            exit
            service-name "evpn-vxlan-502"
            no shutdown
        exit
```

When the above commands are enabled, the router will:

- Receive route-type 5 routes and import the IP prefixes and associated IP next-hops into the VPRN routing table.
  - If the route-type 5 is successfully imported by the router, the prefix included in the route-type 5 (for example, 10.0.0.0/24), will be added to the VPRN routing table with a next-hop equal to the GW IP included in the route (for example, 192.0.0.1. that refers to the IRB IP address of the remote VPRN behind which the IP prefix sits).
  - When the router receives a packet from the WAN to the 10.0.0.0/24 subnet, the IP lookup on the VPRN routing table will yield 192.0.0.1 as the next-hop. That next-hop will be resolved to a MAC in the ARP table and the MAC resolved to a VXLAN tunnel in the FDB table

Note: IRB MAC and IP addresses are advertised in the IRB backhaul R-VPLS in routes type 2.

- Generate route-type 5 routes for the IP prefixes in the associated VPRN routing table.
  - For example, if VPRN-1 is attached to EVPN R-VPLS 1 and EVPN R-VPLS 2, and R-VPLS 2 has **bgp-evpn ip-route-advertisement** configured, the 7750 SR will advertise the R-VPLS 1 interface subnet in one route-type 5.
- Routing policies can filter the imported and exported IP prefix routes accordingly.

The VPRN routing table can receive routes from all the supported protocols (BGP-VPN, OSPF, IS-IS, RIP, static routing) as well as from IP prefixes from EVPN, as shown below:

```
*A:PE72# show router 500 route-table
===============================================================================
Route Table (Service: 500)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
20.20.20.0/24                                   Local   Local     01d11h10m  0
       evi-502                                                     0
20.20.20.71/32                                  Remote  BGP EVPN  00h02m26s  169
       10.10.10.71                                                 0
156.10.10.0/24                                  Remote  Static    00h00m05s  5
       10.10.10.71                                                 1
172.16.0.1/32                                   Remote  BGP EVPN  00h02m26s  169
       10.10.10.71                                                 0
-------------------------------------------------------------------------------
No. of Routes: 4
```

The following considerations apply:

- The route Preference for EVPN IP prefixes is 169.
  - BGP IP-VPN routes have a preference of 170 by default, therefore, if the same route is received from the WAN over BGP-VPRN and from BGP-EVPN, then the EVPN route will be preferred.
- When the same route-type 5 prefix is received from different GW IPs, ECMP is supported if configured in the VPRN.
- All routes in the VPRN routing table (as long as they do not point back to the EVPN R-VPLS interface) are advertised via EVPN.

Although the description above is focused on IPv4 interfaces and prefixes, it applies to IPv6 interfaces too. The following considerations are specific to IPv6 VPRN R-VPLS interfaces:

- IPv4 and IPv6 interfaces can be defined on R-VPLS IP interfaces at the same time (dual-stack).
- The user may configure specific IPv6 Global Addresses on the VPRN R-VPLS interfaces. If a specific Global IPv6 Address is not configured on the interface, the Link Local Address interface MAC/IP will be advertised in a route type 2 as soon as IPv6 is enabled on the VPRN R-VPLS interface.
- Routes type 5 for IPv6 prefixes will be advertised using either the configured Global Address or the implicit Link Local Address (if no Global Address is configured).

  If more than one Global Address is configured, normally the first IPv6 address will be used as GW IP. The "first IPv6 address" refers to the first one on the list of IPv6 addresses shown via show router <id> interface <interface> IPv6 or via SNMP.

The rest of the addresses will be advertised only in MAC-IP routes (Route Type 2) but not used as GW IP for IPv6 prefix routes.

5.2.2.3.2 EVPN for VXLAN in EVPN Tunnel R-VPLS Services

Figure 141 shows an L3 connectivity model that optimizes the solution described in EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes. Instead of regular IRB backhaul R-VPLS services for the connectivity of all the VPRN IRB interfaces, EVPN tunnels can be configured. The main advantage of using EVPN tunnels is that they don't need the configuration of IP addresses, as regular IRB R-VPLS interfaces do.

In addition to the **ip-route-advertisement** command, this model requires the configuration of the **config>service>vprn>if>vpls <***name***> evpn-tunnel**.

Note: The **evpn-tunnel** can be enabled independently of **ip-route-advertisement**, however, no route-type 5 advertisements will be sent or processed in that case.

The example below shows a VPRN (500) with an EVPN-tunnel R-VPLS (504):

```
vprn 500 customer 1 create
    ecmp 4
    route-distinguisher 65071:500
    auto-bind-tunnel
        resolution-filter
        resolution-filter gre ldp rsvp
    vrf-target target:65000:500
    interface "evi-504" create
        vpls "evpn-vxlan-504"
            evpn-tunnel
        exit
    exit
    no shutdown
exit
vpls 504 customer 1 create
    allow-ip-int-bind
    vxlan vni 504 create
    exit
    bgp
        route-distinguisher 65071:504
        route-target export target:65000:504 import target:65000:504
    exit
    bgp-evpn
        ip-route-advertisement
        vxlan
            no shutdown
        exit
    exit
    service-name "evpn-vxlan-504"
```

```
            no shutdown
        exit
```

A specified VPRN supports regular IRB backhaul R-VPLS services as well as EVPN tunnel R-VPLS services.

Note: EVPN tunnel R-VPLS services do not support SAPs or SDP-binds.

The process followed upon receiving a route-type 5 on a regular IRB R-VPLS interface differs from the one for an EVPN-tunnel type:

- IRB backhaul R-VPLS VPRN interface:
  - When a route-type 2 that includes an IP prefix is received and it becomes active, the MAC/IP information is added to the FDB and ARP tables. This can be checked with the **show>router>arp** command and the **show>service>id>fdb detai**l command.
  - When route -type 5 is received and becomes active for the R-VPLS service, the IP prefix is added to the VPRN routing table, regardless of the existence of a route-type 2 that can resolve the GW IP address. If a packet is received from the WAN side and the IP lookup hits an entry for which the GW IP (IP next-hop) does not have an active ARP entry, the system will use ARP to get a MAC. If ARP is resolved but the MAC is unknown in the FDB table, the system will flood into the TLS multicast list. Routes type 5 can be checked in the routing table with the **show>router>route-table** command and the **show>router>fib** command.
- EVPN tunnel R-VPLS VPRN interface:
  - When route -type 2 is received and becomes active, the MAC address is added to the FDB (only).
  - When a route-type 5 is received and active, the IP prefix is added to the VPRN routing table with next-hop equal to EVPN tunnel: GW-MAC.

    For example, ET-d8:45:ff:00:01:35, where the GW-MAC is added from the GW-MAC extended community sent along with the route-type 5.

    If a packet is received from the WAN side, and the IP lookup hits an entry for which the next-hop is a EVPN tunnel: GW-MAC, the system will look up the GW-MAC in the FDB. Usually a route-type 2 with the GW-MAC is previously received so that the GW-MAC can be added to the FDB. If the GW-MAC is not present in the FDB, the packet will be dropped.
  - IP prefixes with GW-MACs as next-hops are displayed by the show router command, as shown below:

```
*A:PE71# show router 500 route-table
```

```
===============================================================================
Route Table (Service: 500)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
20.20.20.72/32                                  Remote  BGP EVPN  00h23m50s  169
       10.10.10.72                                                  0
30.30.30.0/24                                   Remote  BGP EVPN  01d11h30m  169
       evi-504 (ET-d8:45:ff:00:01:35)                               0
156.10.10.0/24                                  Remote  BGP VPN   00h20m52s  170
       192.0.0.69 (tunneled)                                        0
200.1.0.0/16                                    Remote  BGP EVPN  00h22m33s  169
       evi-504 (ET-d8:45:ff:00:01:35)                               0
-------------------------------------------------------------------------------
No. of Routes: 4
```

The GW-MAC as well as the rest of the IP prefix BGP attributes are displayed by the **show>router>bgp>routes>evpn>ip-prefix** command.

```
*A:Dut-A# show router bgp routes evpn ip-prefix prefix 3.0.1.6/32 detail
===============================================================================
BGP Router ID:10.20.1.1        AS:100         Local AS:100
===============================================================================
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
Original Attributes

Network        : N/A
Nexthop        : 10.20.1.2
From           : 10.20.1.2
Res. Nexthop   : 192.168.19.1
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None
Connector      : None
Community      : target:100:1 mac-nh:00:00:01:00:01:02
                 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 10.20.1.2
Flags          : Used  Valid  Best  IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : IP-PREFIX
ESI            : N/A                    Tag            : 1
Gateway Address: 00:00:01:00:01:02
Prefix         : 3.0.1.6/32             Route Dist.    : 10.20.1.2:1
MPLS Label    : 262140
Route Tag      : 0xb
Neighbor-AS    : N/A
```

```
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

Modified Attributes

Network        : N/A
Nexthop        : 10.20.1.2
From           : 10.20.1.2
Res. Nexthop   : 192.168.19.1
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None
Connector      : None
Community      : target:100:1 mac-nh:00:00:01:00:01:02
                 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 10.20.1.2
Flags          : Used  Valid  Best  IGP
Route Source   : Internal
AS-Path        : 111
EVPN type      : IP-PREFIX
ESI            : N/A                    Tag            : 1
Gateway Address: 00:00:01:00:01:02
Prefix         : 3.0.1.6/32             Route Dist.    : 10.20.1.2:1
MPLS Label    : 262140
Route Tag      : 0xb
Neighbor-AS    : 111
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

EVPN tunneling is also supported on IPv6 VPRN interfaces. When sending IPv6 prefixes from IPv6 interfaces, the GW-MAC in the route type 5 (IP-prefix route) is always zero. If no specific Global Address is configured on the IPv6 interface, the routes type 5 for IPv6 prefixes will always be sent using the Link Local Address as GW-IP. The following example output shows an IPv6 prefix received via BGP EVPN.

```
*A:PE71# show router 30 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 30)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
300::/64                                        Local   Local     00h01m19s  0
       int-PE-71-CE-1                                               0
500::1/128                                      Remote  BGP EVPN  00h01m20s  169
       fe80::da45:ffff:fe00:6a-"int-evi-301"                        0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
```

```
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================

*A:PE71# show router bgp routes evpn ipv6-prefix prefix 500::1/128 hunt
===============================================================================
 BGP Router ID:192.0.2.71       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : N/A
Nexthop        : 192.0.2.69
From           : 192.0.2.69
Res. Nexthop   : 192.168.19.2
Local Pref.    : 100                    Interface Name : int-71-69
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None
Connector      : None
Community      : target:64500:301 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.69
Flags          : Used  Valid  Best  IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : IP-PREFIX
ESI            : N/A                    Tag            : 301
Gateway Address: fe80::da45:ffff:fe00:*
Prefix         : 500::1/128             Route Dist.    : 192.0.2.69:301
MPLS Label     : 0
Route Tag      : 0
Neighbor-AS    : N/A
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h41m17s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

5.2.3 DC GW integration with the Nuage Virtual Services Directory (VSD)

The Nuage VSD (Virtual Services Directory) provides automation in the Nuage DC. The VSD is a programmable policy and analytics engine. It provides a flexible and hierarchical network policy framework that enables IT administrators to define and enforce resource policies.

The VSD contains a multi-tenant service directory that supports role-based administration of users, computing, and network resources. The VSD also manages network resource assignments such as IP addresses and ACLs.

To communicate with the Nuage controllers and gateways (including the 7750 SR, 7450 ESS, or 7950 XRS DC GW), VSD uses an XMPP (eXtensible Messaging and Presence Protocol) communication channel. The router can receive service parameters from the Nuage VSD through XMPP and add them to the existing VPRN/VPLS service configuration.

Note: The service must be pre-provisioned in the router using the CLI, SNMP, or other supported interfaces. The VSD will only push a limited number of parameters into the configuration. This router – VSD integration model is known as a Static-Dynamic provisioning model, because only a few parameters are dynamically pushed by VSD, as opposed to a Fully Dynamic model, where the entire service can be created dynamically by VSD.

The router – VSD integration comprises the following building blocks:

- An XMPP interface to the DC XMPP server, through which the router can discover the Data Center Nuage VSDs and select a specified VSD for each VPLS/VPRN service.
- The configuration of **vsd-domains** on those services where VSD will dynamically provision parameters. As part of the static provisioning of a service, the user will configure a domain name (that will be used between VSD and 7750 SR) using a new CLI command **vsd-domain** *name*. Any parameters sent by the VSD for an existing service will contain the **vsd-domain**. Based on that tag, the router will add the required configuration changes to the correct service.
- The dynamic provisioning of parameters in the following four use-cases:
  - L2-DOMAIN: To attach a service at the gateway to a Layer 2 (Ethernet) domain in the data center with no routing at the gateway, a VPLS service should be associated with a **vsd-domain** of type **l2-domain**. When the appropriate configuration for the domain is present/added at the VSD, the VSD will dynamically add the VXLAN VNI and BGP export and import route-targets to exchange DC EVPN routes with the VPLS service.

- L2-DOMAIN-IRB: To attach a service at the gateway to a Layer 2 (Ethernet) domain in the data center with routing at the gateway, an R-VPLS service should be associated with a **vsd-domain** of type **l2-domain-irb**. When the appropriate configuration for the domain is present/added at the VSD, the VSD will dynamically add the VXLAN VNI and BGP export and import route-targets to exchange DC EVPN routes with the R-VPLS service.
- VRF-GRE: To attach a service at the gateway to a layer 3 domain (with GRE transport) in the data center, a VPRN service should be associated with a **vsd-domain** of type **vrf-gre**. When the appropriate configuration for the domain is present/added at the VSD, the VSD will dynamically add the BGP export and import route-targets to exchange DC IP VPN routes with the VPRN service.
- VRF-VXLAN: To attach a service at the gateway to a layer 3 domain (with VXLAN transport) in the data center, an R-VPLS service (linked to an EVPN-tunnel with ip-route-advertisement enabled) should be associated with a **vsd-domain** of type **vrf-vxlan**. When the appropriate configuration for the domain is present/added at the VSD, the VSD will dynamically add the VXLAN VNI and BGP export and import route-targets to exchange DC EVPN routes with the backhaul R-VPLS connected to the data center VPRN service.

These building blocks are described in more detail in the following subsections.

5.2.3.1 XMPP Interface on the DC GW

The Extensible Messaging and Presence Protocol (XMPP) is an open technology for real-time communication using XML (Extensible Markup Language) as the base format for exchanging information. The XMPP provides a way to send small pieces of XML from one entity to another in close to real time.

In a Nuage DC, an XMPP ejabberd server will have an interface to the Nuage VSD as well as the Nuage VSC/VSG and the 7750 SR, 7450 ESS, or 7950 XRS DC GW.

Figure 148 shows the basic XMPP architecture in the data center. While a single XMPP server is represented in the diagram, XMPP allows for easy server clustering and performs message replication to the cluster. It is similar to how BGP can scale and replicate the messages through the use of route reflectors.

Also the VSD is represented as a single server, but a cluster of VSD servers (using the same data base) will be a very common configuration in a DC.

Figure 148 **Basic XMPP Architecture**

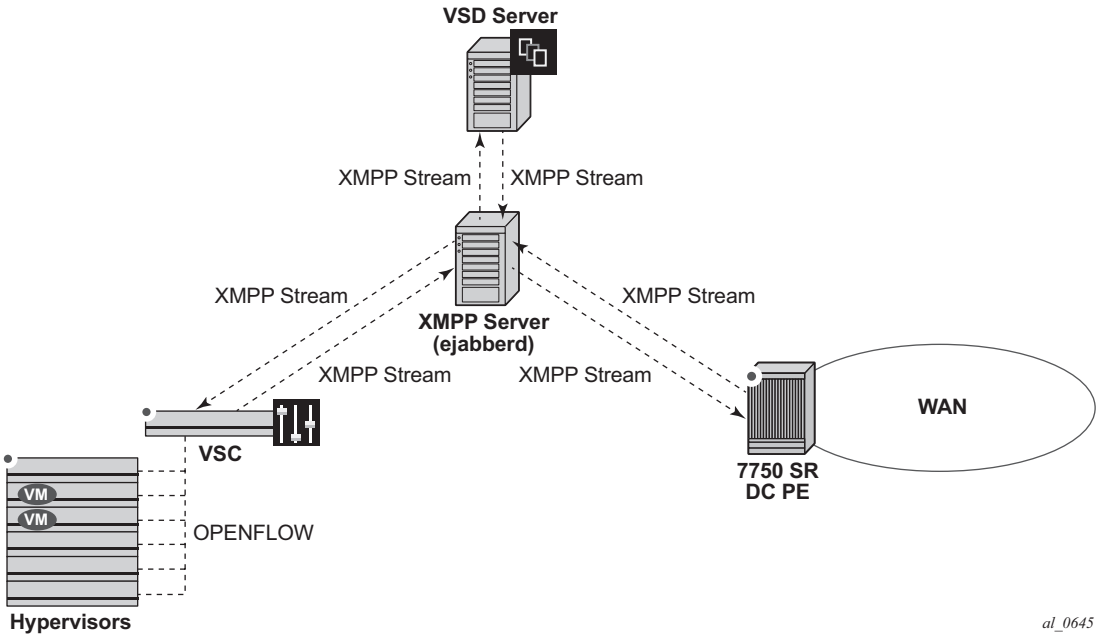

In the Nuage solution, each XMPP client, including the 7750 SR, 7450 ESS, and 7950 XRS, is referred to with a JID (JabberID) in the following format: username@xmppserver.domain. The xmppserver.domain points to the XMPP Server.

To enable the XMPP interface on the 7750 SR, 7450 ESS, or 7950 XRS, the following command must be added to indicate to which XMPP server address the DC GW has to register, as well as the router's JID:

```
*A:PE-2# configure system xmpp server
  - server <xmpp-server-name> [domain-name <fqdn>] [username <user-name>]
    [password <password>] [create] [service-name <service-name>]
  - no server <xmpp-server-name>
  - server <xmpp-server-name> [domain-name <fqdn>] [username <user-name>]
    [password <password>] [create] [router <router-instance>]
 <xmpp-server-name>   : [32 chars max]
 <fqdn>               : [256 chars max]
 <user-name>          : [32 chars max]
 <password>           : [32 chars max]
 <create>             : keyword - mandatory while creating an entry.
 <router-instance>    : <router-name>|<vprn-svc-id>
                        router-name    - "Base"|"management"
                        vprn-svc-id    - [1..2147483647]
 <service-name>       : [64 chars max]
 [no] shutdown        - Administratively enable or disable XMPP server
```

Where:

- [domain-name <fqdn>] is the domain portion of the JID.
- <user-name> and <password> is the username:password portion of the JID of the router acting as an XMPP client. Plain/MD5/anonymous authentication is supported.
- The user can choose not to configure the username portion of the JID. In that case, an in-band registration will be attempted, using the chassis MAC as username.
- The user has the option to try to establish an XMPP TCP session over a router instance by using the **router** *router-instance* command. The router name can be “Base”, “management”, or a given VPRN service identifier.
- When the xmpp server is properly configured and **no shutdown,** the 7750 SR will try to establish a TCP session with the XMPP server through the management interface first. If it fails to establish communication, the 7750 SR will use an in-band communication and will use its system IP as source IP address. **Shutdown** will not remove the dynamic configs in all the services. No server will remove all the dynamic configs in all the services.
- Only one xmpp server can be configured.

Note: The DNS must be configured on the router so that the XMPP server name can be resolved. XMPP relies on the Domain Name System (DNS) to provide the underlying structure for addressing, instead of using raw IP addresses. The DNS is configured using the following bof commands: **bof primary-dns**, **bof secondary-dns**, **bof dns-domain**.

After the XMPP server is properly configured, the router can generate or receive XMPP stanza elements, such as presence and IQ (Information/Query) messages. IQ messages are used between the VSD and the router to request and receive configuration parameters. The status of the XMPP communication channel can be checked with the following command:

```
Dut# show system xmpp server "vsd1-hy"

===============================================================================
XMPP Server Table
===============================================================================
XMPP FQDN          : vsd1-hy.alu.us
XMPP Admin User    : csproot
XMPP Oper User     : csproot
State Lst Chg Since: 0d 02:56:44        State              : Functional
Admin State        : Up                 Connection Mode    : outOfBand
Auth Type          : md5
IQ Tx.             : 47                 IQ Rx.             : 47
IQ Error           : 0                  IQ Timed Out       : 0
IQ Min. Rtt        : 0 ms               IQ Max. Rtt        : 180 ms
IQ Ack Rcvd.       : 47
Push Updates Rcvd  : 1                  VSD list Upd Rcvd  : 12
Msg Tx.            : 27                 Msg Rx.            : 27
Msg Ack. Rx.       : 27                 Msg Error          : 0
```

```
Msg Min. Rtt       : 0 ms               Msg Max. Rtt       : 180 ms
Sub Tx.            : 1                  UnSub Tx.          : 0
Msg Timed Out      : 0

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

In addition to the XMPP server, the router must be configured with a VSD **system-id** that uniquely identifies the router in the VSD:

```
*B:Dut>config>system>vsd# info
----------------------------------------------
                system-id "SR12U-46-PE"
----------------------------------------------
```

After the above configuration is complete, the router will subscribe to a VSD XMPP PubSub node to discover the available VSD servers. Then, the router will be discovered in the VSD UIs. On the router, the available VSD servers can be shown with the following command.

```
B:Dut#show system xmpp vsd

===============================================================================
Virtual Services Directory Table
===============================================================================
Id User Name                    Uptime             Status
-------------------------------------------------------------------------------
1  cna@vsd1-hy.alu-srpm.us/nua* 0d 00:44:36        Available
-------------------------------------------------------------------------------
No. of VSD's: 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
*B:Dut#show system xmpp vsd 1

===============================================================================
VSD Server Table
===============================================================================
VSD User Name      : cna@vsd1-hy.alu-srpm.us/nuage
Uptime             : 0d 00:44:39        Status             : Available
Msg Tx.            : 16                 Msg Rx.            : 10
Msg Ack. Rx.       : 4                  Msg Error          : 6
Msg TimedOut       : 0                  Msg MinRtt         : 80 ms
Msg MaxRtt         : 240 ms

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

5.2.3.2 Overview of the Static-Dynamic VSD Integration Model

In the Static-Dynamic integration model, the DC and DC GW management entities can be the same or different. The DC GW operator will provision the required VPRN and VPLS services with all the parameters needed for the connectivity to the WAN. VSD will only push the required parameters so that those WAN services can be attached to the existing DC domains.

Figure 149 shows the workflow for the attachment of the WAN services defined on the DC GW to the DC domains.

Figure 149 **WAN Services Attachment Workflow**

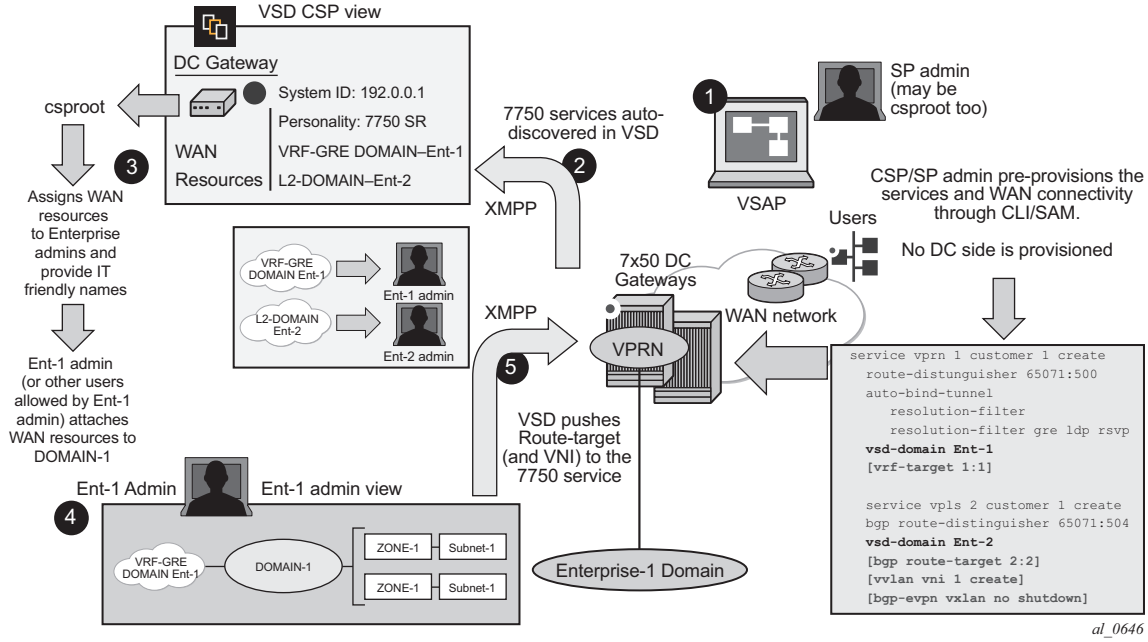

The Static-Dynamic VSD integration model can be summarized in the steps shown in Figure 149 and described in the following procedure.

1. The WAN or SP (Service Provider) administrator (which can be also the DC or Cloud Service Provider administrator) provisions the WAN services with all the parameters required for the connectivity to the WAN. This configuration is performed through the regular management interfaces, for example, CLI or SNMP. In the example above, there are two services created by the SP:
   - VPRN 1 – associated with **vsd-domain Ent-1**, which is a VRF-GRE domain.
   - VPLS 2 – associated with **vsd-domain Ent-2**, which is an L2-DOMAIN

Note: The parameters between brackets "[..]" are not configured at this step. They will be pushed by the VSD through XMPP.

2. The router communicates with the VSD through the XMPP channel and lets VSD know about its presence and available domains: Ent-1 and Ent-2. In the VSD's User Interface (UI), the router will show up as DC GW with its System ID, personality (for example, router) and the available WAN resources, that is, **vsd-domains** Ent-1 and Ent-2.
3. At VSD, the Cloud Service Provider administrator will assign the available WAN resources to Enterprises defined in VSD. In this example, VRF-GRE Ent-1 will be assigned to Enterprise-1 and L2-DOMAIN Ent-2 to Enterprise-2.
4. Each Enterprise administrator will have visibility of their own assigned WAN resource and will attach it to an existing DC Domain, assuming that both the DC domain and WAN resource are compatible. For instance, a VRF-GRE domain can only be attached to an L3 domain in the DC that uses GRE as transport.
5. When the Enterprise administrator attaches the WAN resource to the DC domain, VSD will send the required configuration parameters to the DC GW through the XMPP channel:
   - In the case of the VRF-GRE domain, VSD will only send the **vrf-target** required for the service attachment to the DC domain.
   - In the case of the L2-DOMAIN, VSD will send the **route-target** (in the **service>bgp** or **vsi-import/export** contexts) as well as the **vxlan vni** and the **bgp-evpn vxlan no shutdown** commands.

   WAN resources can also be detached from the DC domains.

5.2.3.3 VSD-Domains and Association to Static-Dynamic Services

In the Static-Dynamic integration model, VSD can only provision certain parameters in VPLS and/or VPRN services. When VSD and the DC GW exchange XMPP messages for a specified service, they use **vsd-domains** to identify those services. A **vsd-domain** is a tag that will be used by the 7750 SR, 7450 ESS, or 7950 XRS router and the VSD to correlate a configuration to a service. When redundant DC GWs are available, the **vsd-domain** for the same service can have the same or a different name in the two redundant DC GWs.

There are four different types of **vsd-domains** that can be configured in the router:

- L2-DOMAIN – it will be associated with a VPLS service in the router and, in VSD, it will be attached to an existing Nuage L2-DOMAIN. This type of domain will be used for extending Layer 2 subnets to the WAN connected to the DC GW.

- L2-DOMAIN-IRB – it will be associated with a R-VPLS service in the router and, in VSD, it will be attached to an existing Nuage L2-DOMAIN. In this case, the DC GW will be the default gateway for all the VMs and hosts in the Nuage L2-DOMAIN.
- VRF-GRE – this domain type will be associated with a VPRN service in the router that uses GRE tunnels and MP-BGP VPN-IPv4 to provide connectivity to the DC. In VSD, it will be attached to an existing Nuage L3-DOMAIN, when GRE is configured as tunnel-type for L3-DOMAINs.
- VRF-VXLAN – this domain type will be associated with a router R-VPLS service (connected to a VPRN with an evpn-tunnel VPLS interface) that uses VXLAN tunnels and EVPN to provide connectivity to the DC. In VSD, it will be attached to an existing Nuage L3-DOMAIN, when VXLAN is configured as the tunnel-type for L3-DOMAINs.

The domains will be configured in the **config>service#** context and assigned to each service.

```
# configure service vsd domain
  - domain <name> [type {l2-domain|vrf-gre|vrf-vxlan|l2-domain-irb}] [create]
  - no domain <name>
 <name>               : [32 chars max]
 <create>             : keyword
 [no] description     - Set VSD Domain Description
 [no] shutdown        - Administratively enable/disable the domain
```

5.2.3.3.1 VSD-Domain Type L2-DOMAIN

L2-DOMAIN VSD-domains will be associated with VPLS services configured without a **route-target** and **vxlan VNI**. VSD will configure the route-target and VNI in the router VPLS service. Some considerations related to L2-DOMAINs are:

- **ip-route-advertisement** and **allow-ip-int-bind** commands are not allowed in this type of domain. An example of configuration for an L2-DOMAIN association is shown below:

```
*B:Dut>config>service# info
...
      vsd
            domain nuage_501 type l2-domain create
                description "nuage_501_l2_domain"
                no shutdown
            exit
*B:Dut>config>service# vpls 501
*B:Dut>config>service>vpls# info
----------------------------------------------
            bgp
                route-distinguisher 192.0.2.2:52
                vsi-import "policy-1"
```

```
                vsi-export "policy-1"
            exit
            bgp-evpn
            exit
            sap 1/1/1:501 create
            exit
            spoke-sdp 10:501 create
                no shutdown
            exit
            vsd-domain "nuage_501"
            no shutdown
----------------------------------------------
```

- The VSD will push a dynamic **vxlan vni** and **route-target** that the router will add to the VPLS service. For the **route-targe**t, the system will check whether the VPLS service has a configured policy:
  - If there is **no vsi-import/export** policy, the received dynamic route-target will be added in the **vpls>bgp** context, and will be used for all the BGP families in the service.
  - If there is a **vsi-import/export** policy, the dynamic route-target will be added to the policy, in an auto-created community that will be shown with the following format: "**_VSD_svc-id**". That community will be added to dynamically created entries 1000 and 2000 in the first policy configured in the service **vsi-import** and **vsi-export** commands. This allows the user to allocate entries before entries 1000 and 2000 in case other modifications have to be made (user entries would have an action next-entry). An example of the auto-generated entries is shown below:

```
*A:PE# show router policy "policy-1"
    entry 900 # manual entry
        from
          as-path "null"
          family evpn
        exit
        action next-entry
          local-preference 500
        exit
    exit
    entry 1000 # automatic VSD-generated entry
        from
            community "_VSD_1"
            family evpn
        exit
        action accept
        exit
    exit
    entry 2000 # automatic VSD-generated entry
        from
            family evpn
        exit
        action accept
            community add "_VSD_1"
        exit
    exit
```

5.2.3.3.2 VSD-Domain Type L2-DOMAIN-IRB

L2-DOMAIN-IRB VSD-domains will be associated with R-VPLS services configured without a static **route-target** and **vxlan VNI**. VSD will configure the dynamic route-target and VNI in the router VPLS service. The same considerations described for L2-DOMAINs apply to L2-DOMAIN-IRB domains with one exception: **allow-ip-int-bind** is now allowed.

5.2.3.3.3 VSD-Domain Type VRF-GRE

VRF-GRE VSD-domains will be associated with VPRN services configured without a static route-target. In this case, the VSD will push a route-target that the router will add to the VPRN service. The system will check whether the VPRN service has a configured policy:

- If there is no **vrf-import** policy, the received dynamic route-target will be added in the vprn> context.
- If there is a **vrf-import** policy, the dynamic route-target will be added to the policy, in an auto-created community that will be shown with the following format: "**VSD_svc-id**" in a similar way as in L2-DOMAINs.

Note: In cases where a **vrf-import** policy is used, the user will provision the WAN **route-target** statically in a **vrf-export** policy. This **route-target** will also be used for the routes advertised to the DC.

An example of the auto-generated entry is shown below:

```
*A:PE# show router policy "policy-1"
    entry 1000 # automatic VSD-generated entry
        from
            community "_VSD_1"
            family vpn-ipv4
        exit
        action accept
        exit
    exit
```

5.2.3.3.4 VSD-Domain Type VRF-VXLAN

VRF-VXLAN VSD-domains will be associated with R-VPLS services configured without a static **route-target** and **vxlan VNI**. VSD will configure the dynamic route-target and VNI in the router VPLS service. Some considerations related to VRF-VXLAN domains are:

- **ip-route-advertisement**, **allow-ip-int-bind**, as well as the VPRN **evpn-tunnel** commands are now required for this type of VSD-domain. An example of configuration for a VRF-VXLAN association is shown below:

```
*A:Dut>config>service# info
<snip>
     vsd
            domain L3Domain-1 type vrf-vxlan create
                description "L3Domain-example"
                no shutdown
            exit
*A:Dut>config>service# vpls 20003
*A:Dut>config>service>vpls# info
----------------------------------------------
            allow-ip-int-bind
            bgp
                route-distinguisher 65000:20003
            exit
            bgp-evpn
                ip-route-advertisement
            exit
            stp
                shutdown
            exit
            service-name "vpls-20003"
            vsd-domain "L3Domain-1"
            no shutdown
----------------------------------------------
*A:sr7L2-47-PE4# configure service vprn 20002
*A:sr7L2-47-PE4>config>service>vprn# info
----------------------------------------------
            route-distinguisher 65000:20002
            auto-bind-tunnel
                resolution-filter
                resolution-filter gre ldp rsvp
            vrf-target target:10:10
            interface "toDC" create
                vpls "vpls-20003"
                    evpn-tunnel
                exit
            exit
            no shutdown
```

- The VSD will push a dynamic **vxlan VNI** and **route-target** that the router will add to the VPLS service. For the **route-target**, the system will check whether the VPLS service has a configured policy and will push the **route-target** either in the service context or the **vsi-import/export** policies, as described in the section for L2-DOMAINs.

The following commands help show the association between the 7750 SR, 7450 ESS, and 7950 XRS router services and VSD-domains, as well as statistics and configuration errors sent/received to/from VSD.

```
*A:Dut# show service service-using vsd
===============================================
Services-using VSD Domain
```

```
=========================================
Svc Id     Domain
-----------------------------------------
501        nuage_501
200001     MyL2Domain
20003      MyL3Domain
-----------------------------------------
Number of services using VSD Domain: 3
=========================================
*A:Dut# show service vsd domain "MyL3Domain"

===============================================================================
VSD Information
===============================================================================
Name               : MyL3Domain
Description        : MyL3Domain-example
Type               : vrfVxlan                      Admin State   : inService
Last Error To Vsd  : (Not Specified)
Last Error From Vsd: (Not Specified)

Statistics
-------------------------------------------------------------------------------
Last Cfg Chg Evt   : 02/06/2015 01:28:30          Cfg Chg Evts  : 671
Last Cfg Update    : 02/06/2015 02:58:41          Cfg Upd Rcvd  : 3
Last Cfg Done      : 02/06/2015 02:58:41
Cfg Success        : 667                          Cfg Failed    : 0
===============================================================================

*A:Dut# show service vsd domain "MyL3Domain" association

============================================================
Service VSD Domain
============================================================
Svc Id      Svc Type  Domain Type    Domain Admin   Origin
------------------------------------------------------------
20003       vpls      vrfVxlan       inService      manual
------------------------------------------------------------
Number of entries: 1
============================================================
```

5.2.3.4 Fully-Dynamic VSD Integration Model

In the Static-Dynamic VSD integration model, the VPLS/VPRN service, as well as most of the parameters, must be provisioned "statically" through usual procedures (CLI, SNMP, and so on). VSD will "dynamically" send the parameters that are required for the attachment of the VPLS/VPRN service to the L2/L3 domain in the Data Center. In the Fully-Dynamic VSD integration model, the entire VPLS/VPRN service configuration will be dynamically driven from VSD and no static configuration

is required. Through the existing XMPP interface, the VSD will provide the 7750 SR, 7450 ESS, and 7950 XRS routers with a handful of parameters that will be translated into a service configuration by a python-script. This python-script provides an intermediate abstraction layer between VSD and the router, translating the VSD data model into a 7750 SR, 7450 ESS, or 7950 XRS CLI data model.

In this Fully-Dynamic integration model, the DC and DC GW management entities are usually the same. The DC GW operator will provision the required VPRN and VPLS services with all the parameters needed for the connectivity to the WAN and the DC. VSD will push the required parameters so that the router can create the service completely and get it attached to an existing DC domain.

The workflow of the Fully-Dynamic integration model is shown in Figure 150.

Figure 150 **Fully-Dynamic VSD Integration Model Workflow**

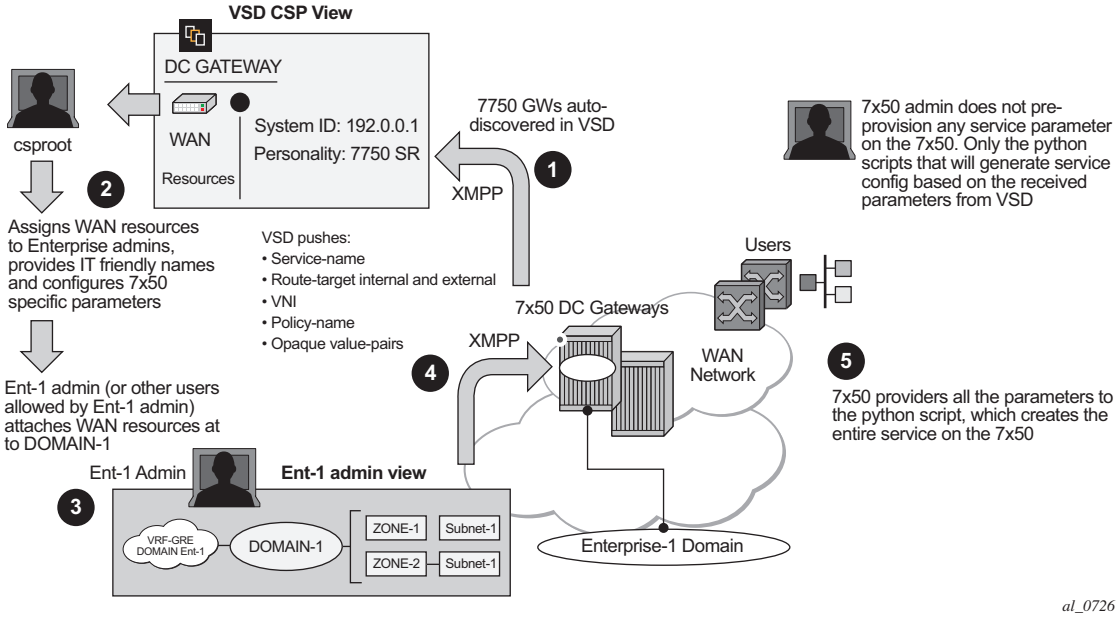

The Fully-Dynamic VSD integration model can be summarized in the steps shown in Figure 150 and described in the following procedure.

1. The 7750 SR, 7450 ESS, or 7950 XRS administrator only needs to provision parameters required for connectivity to the VSD, a s**ervice-range**, and configure the python script/policy in the system. Provisioning of service parameters is not required.

   The **service-range** defines the service identifiers to include for VSD provisioned services and, once configured, they are protected from CLI changes. The vsd-policy defines the script to be used:

```
*A:PE1>config>python# info
```

```
-----------------------------------------------
        python-script "l2-domain_services" create
            primary-url "ftp://1.1.1.1/cf2/l2domain_service.py"
            no shutdown
        exit
  python-policy "py-l2" create
            description "Python script to create L2 domains"
            vsd script "l2-domain_services"
        exit
-----------------------------------------------

*A:PE1>config>service>vsd# info
-----------------------------------------------
  service-range 64000 to 65000
-----------------------------------------------
```

When the router boots up or the gateway configuration is changed, the router sends a message to the VSD indicating its capabilities:

- System-ID
- Name and gateway type

The VSD uses this information to register the router on its list of router GWs.

When registered, the VSD and router exchange messages where the VSD communicates its list of service-names and their domain-type to the router. Based on this list, the router sends an XMPP IQ message to request the configuration of a specified service.

The 7750 SR, 7450 ESS, or 7950 XRS router will periodically audit the VSD and request a “DIFF” list of Full-Dynamic VSD domains. The VSD keeps a “DIFF” list of domains, that contains the Fully-Dynamic domain names for which the VSD has not received an IQ request from the router for a long time.

The 7750 SR, 7450 ESS, or 7950 XRS CLI user can also audit the VSD to get the DIFF list, or even the “FULL” list of all the domains in the VSD. The following command triggers this audit: **tools perform service vsd fd-domain-sync {full|diff}**.

2. Concurrently at the VSD, the Cloud Service Provider administrator will assign WAN resources to Enterprises defined in the VSD. In this example, a VRF-GRE domain will be assigned to Enterprise-1.
3. Each Enterprise administrator will have visibility of their own assigned WAN resource and will attach it to an existing DC Domain, assuming that both the DC domain and WAN resource are compatible. For instance, a VRF-GRE domain can only be attached to an L3 domain in the DC that uses GRE as transport.
4. When the Enterprise administrator attaches the service requested by the 7750 SR, 7450 ESS, or 7950 XRS router to the DC domain, the VSD will send the required configuration parameters for that service to the DC GW through the XMPP channel in an IQ Service message, including the following information:
   - Service name and service type, where the type can be:

- L2-DOMAIN
- L2-DOMAIN-IRB
- VRF-GRE
- VRF-VXLAN

– Configuration type— Static (for Static-Dynamic model) or Dynamic (for Fully-Dynamic model).
– Internal route-target (RT-i) — Used to export/import the BGP routes sent/received from/to the DC route-reflector.
– External route-target (RT-e) — Used to export/import the BGP routes sent/received from/to the WAN route-reflector. The value can be the same as the RT-i.
– VNI (VXLAN Network Identifier) — Used to configure the EVPN-VXLAN VPLS service on the router (if the domain type is L2-DOMAIN, L2-DOMAIN-IRB, or VRF-VXLAN).
– Metadata — A collection of 'opaque' <key=value> pairs including the rest of the service parameters required for the service configuration at the router.

Note: The keys or values do not need to follow any specific format. The python script interprets and translates them into the router data model.

– Python-policy— Used by the router to find the Python script that will translate the VSD parameters into configuration.

5. When the 7750 SR, 7450 ESS, or 7950 XRS router receives the IQ Service message, it builds a string with all the parameters and passes it to the Python module. The Python module is responsible for creating and activating the service, and, therefore, provides connectivity between the tenant domain and the WAN.

Note: The python-script cannot access all the CLI commands and nodes in the system. A white-list of nodes and commands is built and Python will only have access to those nodes/commands. The list can be displayed using the following command: **tools dump service vsd-services command-list**.

In addition to the *white-list*, the user can further restrict the allowed CLI nodes to the VSD by using a separate CLI user for the XMPP interface, and associating that user to a profile where the commands are limited. The CLI user for the XMPP interface is configurable:

```
config>system>security>cli-script>authorization>
   vsd
[no] cli-user <username>
```

When the system executes a python-script for VSD commands, the *vsd cli-user* profile is checked to allow the resulting commands. A single CLI user is supported for VSD, therefore, the operator can configure a single 'profile' to restrict (even further than the *whitelist*) the CLI commands that can be executed by the VSD Python scripts.

No *cli-user* means that the system will not perform any authorization checks and the VSD scripts can access any commands that are supported in the *white-list*.

5.2.3.4.1 Python Script Implementation Details

A python-script provides an intermediate abstraction layer between VSD and the router, translating the VSD data model into the 7750 SR, 7450 ESS, or 7950 XRS router CLI data model. VSD will use metadata key=value parameters to provision service specific attributes on the 7750. The XMPP messages get to the 7750 and are passed transparently to the Python module so that the CLI is generated and the service created.

Note: The CLI generated by the python-script is not saved in the configuration file; it can be displayed by running the **info include-dynamic** command on the service contexts. The configuration generated by the python-script is protected and cannot be changed by a CLI user.

The following example shows how the python-script works for a VSD generated configuration:

- The following configuration is added to the 7750 SR. In this case, *py-l2* is the python-policy received from VSD that will call the *l2domain_service.py* python script:

```
*A:PE1>config>python# info
----------------------------------------------
        python-script "l2-domain_services" create
            primary-url "ftp://1.1.1.1/cf2/l2domain_service.py"
            no shutdown
        exit
            python-policy "py-l2" create
            description "Python script to create L2 domains"
            vsd script "l2-domain_services"
        exit
----------------------------------------------
*A:PE1>config>service>vsd# info
----------------------------------------------
  service-range 64000 to 65000
----------------------------------------------
```

- VSD will send metadata containing the service parameters. This opaque parameter string will be passed to the python script and is composed of tag=value pairs, with the following format:
- In addition, other information provided by the VSD, (domain, vni, rt-i, rt-e, and service type) is bundled with the metadata string and passed to the python script. For example:

```
{'rt': 'target:64000:64000', 'rte': 'target:64000:64000', 'domain': 'L2-DOMAIN-
1', 'servicetype': 'L2DOMAIN', 'vni': '64000', 'metadata': 'rd=1:1, sap=1/1/
10:3000'}
```

The user should consider the following:

- The python script is solely responsible for generating the configuration; no configuration aspects of the Static-Dynamic model are used. The python script must be written in a manner similar to those used by RADIUS Dynamic Business Services. Currently, RADIUS Dynamic Business Services and the Fully-Dynamic VSD model are mutually exclusive, one or the other can operate on the same system, but not both at the same time.
- The following scripts must be defined in order to set up, modify, revert, and tear down the configuration for a service: *setup_script(), modify_script(), revert_script(),* and *teardown_script()*. These names must always be the same in all scripts. The *revert_script()* is only required if the *modify_script()* is defined, the latter being optional.
- When the configuration for a new domain name is received from the VSD, the metadata and the VSD parameters are concatenated into a single dictionary and *setup_script()* is called. Within the python script:
  - The VSD UI parameters are referenced as vsdParams['rt'], vsdParams ['domain'], and so on.
  - The metadata parameters are referenced as vsdParams['metadata'].
- When the startup script is executed, the **config>service>vsd>domain** is created outside the script context before running the actual script. The teardown script will remove the **vsd domain**.
  - If a specified *setup_script()* fails, the *teardown_script()* is invoked.
- When subsequent configuration messages are received from the VSD, the new parameter list is generated again from the VSD message and compared to the last parameter list that was successfully executed.
  - If the two strings are identical, no action is taken.
  - If there is a difference between the strings, the *modify_script()* function is called.
  - If the *modify_script()* fails, the *revert_script()* is invoked. The *teardown_script()* is invoked if the *revert_script()* fails or does not exist.

- The **python-policy** is always present in the attributes received from VSD; if the VSD user does not include a policy name, VSD will include 'default' as the python-policy. Hence, care must be taken to ensure that the 'default' policy is always configured in the 7750.
- If the scripts are incorrect, teardown and modify procedures could leave orphaned objects. An admin mode (**enable-vsd-config**) is available to enable an administrator to clean up these objects; it is strictly meant for cleaning orphaned objects only.

Note: The CLI configured services cannot be modified when the user is in **enable-vsd-config** mode.

- Unless the CLI user enters the **enable-vsd-config** mode, changes of the dynamic created services are not allowed in the CLI. However, changes through SNMP are allowed.
- The command **tools perform service vsd evaluate-script** is introduced to allow the user to test their setup and to modify and tear down python scripts in a lab environment without the need to be connected to a VSD. The successful execution of the command for **action setup** will create a **vsd domain** and the corresponding configuration, just as the system would do when the parameters are received from VSD.

The following example shows the use of the **tools perform service vsd evaluate-script** command:

```
*A:PE1# tools perform service vsd evaluate-script domain-name "L2-DOMAIN-5" type l2-
domain action setup policy "py-l2" vni 64000 rt-i target:64000:64000 rt-
e target:64000:64000 metadata "rd=1:1, sap=1/1/10:3000"

Success
```

The following example output shows a python-script that can set up or tear down L2-DOMAINs.

```
*A:PE1# show python python-script "l2-domain_services" source-in-use

===============================================================================
Python script "l2-domain_services"
===============================================================================
Admin state    : inService
Oper state     : inService
Primary URL    : ftp://1.1.1.1/timos86/cses-V71/cf2/l2domain_service.py
Secondary URL  : (Not Specified)
Tertiary URL   : (Not Specified)
Active URL     : primary

-------------------------------------------------------------------------------
```

```
Source (dumped from memory)
-------------------------------------------------------------------------------
      1 from alc import dyn
      2
      3 def setup_script(vsdParams):
      4
      5     print ("These are the VSD params: " + str(vsdParams))
      6     servicetype = vsdParams.get('servicetype')
      7     vni = vsdParams.get('vni')
      8     metadata = vsdParams['metadata']
      9     print ("VSD metadata: " + str(metadata))
     10     metadata = dict(e.split('=') for e in metadata.split(','))
     11     print ("Modified metadata: " + str(metadata))
     12     vplsSvc_id = dyn.select_free_id("service-id")
     13     print ("this is the free svc id picked up by the system: " + vplsSvc_id)
     14
     15     if servicetype == "L2DOMAIN":
     16         rd = metadata['rd']
     17         sap_id = metadata[' sap']
     18         print ('servicetype, VPLS id, rd, sap:', servicetype, vplsSvc_id, rd
, sap_id)
     19         dyn.add_cli("""
     20           configure service
     21               vpls %(vplsSvc_id)s customer 1 create
     22                   description vpls%(vplsSvc_id)s
     23                   bgp
     24                       route-distinguisher %(rd)s
     25                       route-target %(rt)s
     26                   exit
     27                   vxlan vni %(vni)s create
     28                   exit
     29                   bgp-evpn
     30                       evi %(vplsSvc_id)s
     31                       vxlan
     32                           no shutdown
     33                       exit
     34                   exit
     35                   service-name evi%(vplsSvc_id)s
     36                   sap %(sap_id)s create
     37                   exit
     38                   no shutdown
     39                   exit
     40               exit
     41           exit
     42         """ % {'vplsSvc_id' : vplsSvc_id, 'vni' : vsdParams['vni'], 'rt' : v
sdParams['rt'], 'rd' : metadata['rd'], 'sap_id' : sap_id})
     43         # L2DOMAIN returns setupParams: vplsSvc_id, servicetype, vni, sap
     44         return {'vplsSvc_id' : vplsSvc_id, 'servicetype' : servicetype, 'vni
' : vni, 'sap_id' : sap_id}
     45
     46
     47 #-------------------------------------------------------------------------
     48
     49 def teardown_script(setupParams):
     50     print ("These are the teardown_script setupParams: " + str(setupParams))
     51     servicetype = setupParams.get('servicetype')
     52     if servicetype == "L2DOMAIN":
     53       dyn.add_cli("""
     54         configure service
```

```
     55             vpls %(vplsSvc_id)s
     56                 no description
     57                 bgp-evpn
     58                     vxlan
     59                         shutdown
     60                     exit
     61                     no evi
     62                     exit
     63                 no vxlan vni %(vni)s
     64                 bgp
     65                    no route-distinguisher
     66                    no route-target
     67                 exit
     68                 no bgp
     69                 no bgp-evpn
     70                 sap %(sap_id)s
     71                    shutdown
     72                    exit
     73                 no sap %(sap_id)s
     74                 shutdown
     75                 exit
     76                 no vpls %(vplsSvc_id)s
     77             exit
     78         exit
     79       """ % {'vplsSvc_id' : setupParams['vplsSvc_id'], 'vni' : setupParams['
vni'], 'sap_id' : setupParams['sap_id']})
     80       return setupParams
     81
     82
     83 d = { "script" : (setup_script, None, None, teardown_script) }
     84
     85 dyn.action(d)
================================================================================
```

For instance, assuming that the VSD sends the following:

```
{'rt': 'target:64000:64000', 'rte': 'target:64000:64000', 'domain': 'L2-DOMAIN-
5', 'servicetype': 'L2DOMAIN', 'vni': '64000', 'metadata': 'rd=1:1, sap=1/1/
10:3000 '}
```

The system will execute the setup script as follows:

```
123 2015/06/16 23:35:40.22 UTC MINOR: DEBUG #2001 Base dyn-script req=setup
"dyn-script req=setup: L2-DOMAIN-5
  state=init->waiting-for-setup
"

124 2015/06/16 23:35:40.22 UTC MINOR: DEBUG #2001 Base dyn-script req=setup
"dyn-script req=setup: L2-DOMAIN-5
  state=waiting-for-setup->generating-setup
"

125 2015/06/16 23:35:40.22 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: l2-domain_services
These are the VSD params: {'rt': 'target:64000:64000', 'rte': 'target:64000:6400
0', 'domain': '', 'servicetype': 'L2DOMAIN', 'vni': '64000', 'metadata': 'rd=1:1
, sap=1/1/10:3000 '}
```

```
VSD metadata: rd=1:1, sap=1/1/10:3000
Modified metadata: {'rd': '1:1', ' sap': '1/1/10:3000 '}
this is the free svc id picked up by the system: 64000
('servicetype, VPLS id, rd, sap:', 'L2DOMAIN', '64000', '1:1', '1/1/10:3000 ')
"
Success

126 2015/06/16 23:35:40.22 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: l2-domain_services
"

127 2015/06/16 23:35:40.22 UTC MINOR: DEBUG #2001 Base dyn-script req=setup
"dyn-script req=setup: L2-DOMAIN-5
  state=generating-setup->executing-setup
"

128 2015/06/16 23:35:40.22 UTC MINOR: DEBUG #2001 Base dyn-script cli 1/1
"dyn-script cli 1/1: script:L2-DOMAIN-5(cli 698 dict 0->126)

          configure service
              vpls 64000 customer 1 create
                  description vpls64000
                  bgp
                      route-distinguisher 1:1
                      route-target target:64000:64000
                  exit
                  vxlan vni 64000 create
                  exit
                  bgp-evpn
                      evi 64000
                      vxlan
                          no shutdown
                      exit
                  exit
                  service-name evi64000
                  sap 1/1/10:3000  create
                  exit
                  no shutdown
                  exit
              exit
          exit
        "

143 2015/06/16 23:35:40.23 UTC MINOR: DEBUG #2001 Base dyn-script req=setup
"dyn-script req=setup: L2-DOMAIN-5
  state=executing-setup->established"
```

5.2.3.4.2 Provisioning Filters using the VSD Fully Dynamic Model

IP, IPv6, and MAC filters can be configured from the VSD within the context of the Fully Dynamic XMPP provisioning model for VPRN and VPLS services.

The VSD filters or filter entries are intended for use in two DC environments:

- **Dedicated DC GW Model**

The DC GW services and filter policies in this DC environment are completely owned and self-managed by the Nuage VSD. In this model, the filter cannot be changed and/or deleted by any management or policy interface other than the VSD; changes are not saved in the configuration file.

To enable the VSD to configure a filter, the python setup script must contain a **config filter ip/ipv6/mac-filter _tmnx_vsd_<filter_id> create** statement. The VSD exclusively manages the removal and change of such filters.

The following excerpt shows a sample setup script to create a filter.

```
def setup_script(vsdParams):
<snip>
       filter_id = metadata[' filter']
<snip>
       dyn.add_cli("""
              config filter ip-filter _tmnx_vsd_%(filter_id)s create
       entry 10 create
              match protocol tcp
                    dst-port eq 80
              exit
             action
               forward
             exit
<snip>
```

- **PE/BNG + DC GW Combination Model**

  The filter and point of embedding insertion in this DC environment is owned by a WAN controller. In this model, the entries in the embedded filter are populated by the VSD.

  The WAN controller creates a filter and the embedding point (through a management interface other than the VSD) by using the **config filter ip/ipv6/mac-filter <id> embed-filter vsd _tmnx_vsd_<filter-id>** command. When this command is run, a filter with the name **_tmnx_vsd_<filter-id>** will be auto-generated; the python scripts can use that name to create entries driven by the VSD.

5.2.4 Layer 2 Multicast Optimization for VXLAN (Assisted-Replication)

The Assisted-Replication feature for IPv4 VXLAN tunnels (both Leaf and Replicator functions) is supported in compliance with the non-selective mode described in IETF Draft *draft-ietf-bess-evpn-optimized-ir*.

The Assisted-Replication feature is a Layer 2 multicast optimization feature that helps software-based PE and NVEs with low-performance replication capabilities to deliver broadcast and multicast Layer 2 traffic to remote VTEPs in the VPLS service.

The EVPN and proxy-ARP/ND capabilities can reduce the amount of broadcast and unknown unicast in the VPLS service; ingress replication is sufficient for most use cases in this scenario. However, when multicast applications require a significant amount of replication at the ingress node, software-based nodes struggle due to their limited replication performance. By enabling the Assisted-Replication Leaf function on the software-based SR-series router, all the broadcast and multicast packets are sent to a 7x50 router configured as a Replicator, which replicates the traffic to all the VTEPs in the VPLS service on behalf of the Leaf. This guarantees that the broadcast or multicast traffic is delivered to all the VPLS participants without any packet loss caused by performance issues.

The Leaf or Replicator function is enabled per VPLS service by the **config>service>vpls>vxlan>assisted-replication {replicator|leaf}** command. In addition, the Replicator requires the configuration of an Assisted-Replication IP (AR-IP) address. The AR-IP loopback address indicates whether the received VXLAN packets have to be replicated to the remote VTEPs. The AR-IP address is configured using the **config>service>system>vxlan>assisted-replication-ip <ip-address>** command.

Based on the **assisted-replication {replicator|leaf}** configuration, the SR-series router can behave as a Replicator (AR-R), Leaf (AR-L), or Regular Network Virtualization Edge (RNVE) router. An RNVE router does not support the Assisted-Replication feature. Because it is configured with no assisted replication, the RNVE router ignores the AR-R and AR-L information and replicates to its flooding list where VTEPs are added based on the regular ingress replication routes.

5.2.4.1 Replicator (AR-R) Procedures

An AR-R configuration is shown in the following example.

```
*A:PE-2>config>service>system>vxlan# info
----------------------------------------------
            assisted-replication-ip 2.2.2.2
----------------------------------------------
*A:PE-2>config>service>vpls# info
----------------------------------------------
            vxlan vni 4000 create
                assisted-replication replicator
            exit
            bgp
            exit
            bgp-evpn
                evi 4000
                vxlan
                    no shutdown
                exit
<snip>
            no shutdown
```

```
-------------------------------------------
```

In this example configuration, the BGP advertises a new inclusive multicast route with tunnel-type = AR, type (T) = AR-R, and tunnel-id = originating-ip = next-hop = assisted-replication-ip (IP address 2.2.2.2 in the preceding example). In addition to the AR route, the AR-R sends a regular IR route if **ingress-repl-inc-mcast-advertisement** is enabled.

Note: You should disable the **ingress-repl-inc-mcast-advertisement** command if the AR-R does not have any SAP or SDP-bindings and is used solely for Assisted-Replication functions.

The AR-R builds a flooding list composed of ACs (SAPs and SDP-bindings) and VXLAN tunnels to remote nodes in the VPLS. All objects in the flooding list are broadcast/multicast (BM) and unknown unicast (U) capable. The following sample output of the **show service id vxlan** command shows that the VXLAN destinations in the flooding list are tagged as “BUM”.

```
*A:PE-2# show service id 4000 vxlan
===============================================================================
Vxlan Src Vtep IP: N/A
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 4000
Creation Origin: manual
Assisted-Replication: replicator
RestProtSrcMacAct: none
===============================================================================
VPLS VXLAN service Network Specifics
===============================================================================
Ing Net QoS Policy : none                       Vxlan VNI Id     : 4000
Ingress FP QGrp    : (none)                     Ing FP QGrp Inst : (none)
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI  Num. MACs   Mcast Oper  L2
                                                                    State PBR
-------------------------------------------------------------------------------
192.0.2.3                             4000        0           BUM   Up    No
192.0.2.5                             4000        0           BUM   Up    No
192.0.2.6                             4000        0           BUM   Up    No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 3
-------------------------------------------------------------------------------
===============================================================================
```

When the AR-R receives a BUM packet on an AC, the AR-R forwards the packet to its flooding list (including the local ACs and remote VTEPs).

When the AR-R receives a BM packet on a VXLAN tunnel, it checks the IP DA of the underlay IP header and performs the following BM packet processing.

Note: The AR-R function is only relevant to BM packets; it does not apply to unknown unicast packets. If the AR-R receives unknown unicast packets, it sends them to the flooding list, skipping the VXLAN tunnels.

- If the destination IP matches its AR-IP, the AR-R forwards the BM packet to its flooding list (ACs and VXLAN tunnels). The AR-R performs source suppression to ensure that the traffic is not sent back to the originating Leaf.
- If the destination IP matches its regular VXLAN termination IP (IR-IP), the AR-R skips all the VXLAN tunnels from the flooding list and only replicates to the local ACs. This is the default Ingress Replication (IR) behavior.

5.2.4.2 Leaf (AR-L) procedures

An AR-L is configured as shown in the following example.

```
A:PE-3>config>service>vpls# info
----------------------------------------------
            vxlan vni 4000 create
                assisted-replication leaf replicator-activation-time 30
            bgp
            exit
            bgp-evpn
                evi 4000
                vxlan
                    no shutdown
                exit
                mpls
                    shutdown
                exit
            exit
            stp
                shutdown
            exit
            sap 1/1/1:4000 create
                no shutdown
            exit
            no shutdown
----------------------------------------------
```

In this example configuration, the BGP advertises a new inclusive multicast route with a tunnel-type = IR, type (T) = AR-L and tunnel-id = originating-ip = next-hop = IR-IP (IP address terminating VXLAN normally, either system-ip or vxlan-src-vtep address).

The AR-L builds a single flooding list per service but controlled by the BM and U flags. These flags are displayed in the following **show service id vxlan** command sample output.

```
A:PE-3# show service id 4000 vxlan
===============================================================================
Vxlan Src Vtep IP: N/A
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 4000
Creation Origin: manual
Assisted-Replication: leaf      Replicator-Activation-Time: 30
RestProtSrcMacAct: none
===============================================================================
VPLS VXLAN service Network Specifics
===============================================================================
Ing Net QoS Policy : none                          Vxlan VNI Id     : 4000
Ingress FP QGrp    : (none)                        Ing FP QGrp Inst : (none)
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI  Num. MACs   Mcast Oper  L2
                                                                    State PBR
-------------------------------------------------------------------------------
2.2.2.2                               4000        0           BM    Up    No
4.4.4.4                               4000        0           -     Up    No
192.0.2.2                             4000        0           U     Up    No
192.0.2.5                             4000        0           U     Up    No
192.0.2.6                             4000        0           U     Up    No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 5
-------------------------------------------------------------------------------
===============================================================================
```

The AR-L creates the following VXLAN destinations when it receives and selects a Replicator-AR route or the Regular-IR routes.

- A VXLAN destination to each remote PE that sent an IR route. These bindings have the U flag set.
- A VXLAN destination to the selected AR-R. These bindings have only the BM flag set; the U flag is not set.
- The non-selected AR-Rs create a binding with flag “-” (in the CPM) that is displayed by the **show service id vxlan** command. Although the VXLAN destinations to non-selected AR-Rs do not carry any traffic, the destinations count against the total limit and must be considered when accounting for consumed VXLAN destinations in the router.

The BM traffic is only sent to the selected AR-R, whereas the U (unknown unicast) traffic is sent to all the destinations with the U flag.

The AR-L performs per-service load-balancing of the BM traffic when two or more AR-Rs exist in the same service. The AR Leaf creates a list of candidate PEs for each AR-R (ordered by IP and VNI; candidate 0 being the lowest IP and VNI). The replicator is selected out of a modulo function of the service-id and the number of replicators, as shown in the following sample output.

```
A:PE-3# show service id 4000 vxlan assisted-replication replicator
```

```
===============================================================================
Vxlan AR Replicator Candidates
===============================================================================
VTEP Address           Egress VNI     In Use  In Candidate List Pending Time
-------------------------------------------------------------------------------
2.2.2.2                4000           yes     yes               0
4.4.4.4                4000           no      yes               0
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

A change in the number of Replicator-AR routes (for example, if a route is withdrawn or a new route appears) affects the result of the hashing, which may cause a different AR-R to be selected.

Note: An AR-L waits for the configured replicator-activation-time before sending the BM packets to the AR-R. In the interim, the AR-L uses regular ingress replication procedures. This activation time allows the AR-R to program the Leaf VTEP. If the timer is zero, the AR-R may receive packets from a not-yet-programmed source VTEP, in which case it will discard the packets.

The following list summarizes other aspects of the AR-L behavior.

- When a Leaf receives a BM packet on an AC, it sends the packet to its flood list that includes access SAP or SDP-bindings and VXLAN destinations with BM or BUM flags. If a single AR-R is selected, only a VXLAN destination will include the BM flags.
- Control plane-generated BM packets, such as ARP/ND (when proxy-ARP/ND is enabled) or Eth-CFM, follow the behavior of regular data plane BM packets.
- When a Leaf receives an unknown unicast packet on an AC, it sends the packet to the flood-list, skipping the AR destination because the U flag is set to 0. To avoid packet re-ordering, the unknown unicast packets do not go through the AR-R.
- When a Leaf receives a BUM packet on an overlay tunnel, it forwards the packet to the flood list, skipping the VXLAN tunnels (that is, the packet is sent to the local ACs and never to a VXLAN tunnel). This is the default IR behavior.
- When the last Replicator-AR route is withdrawn, the AR-L removes the AR destination from the flood list and falls back to ingress replication.

Figure 151 shows the expected replication behavior for BM traffic when received at the access on an AR-R, AR-L, or RNVE router. Unknown unicast follows regular ingress replication behavior regardless of the role of the ingress node for the specific service.

Figure 151 **AR BM Replication Behavior for a BM Packet**

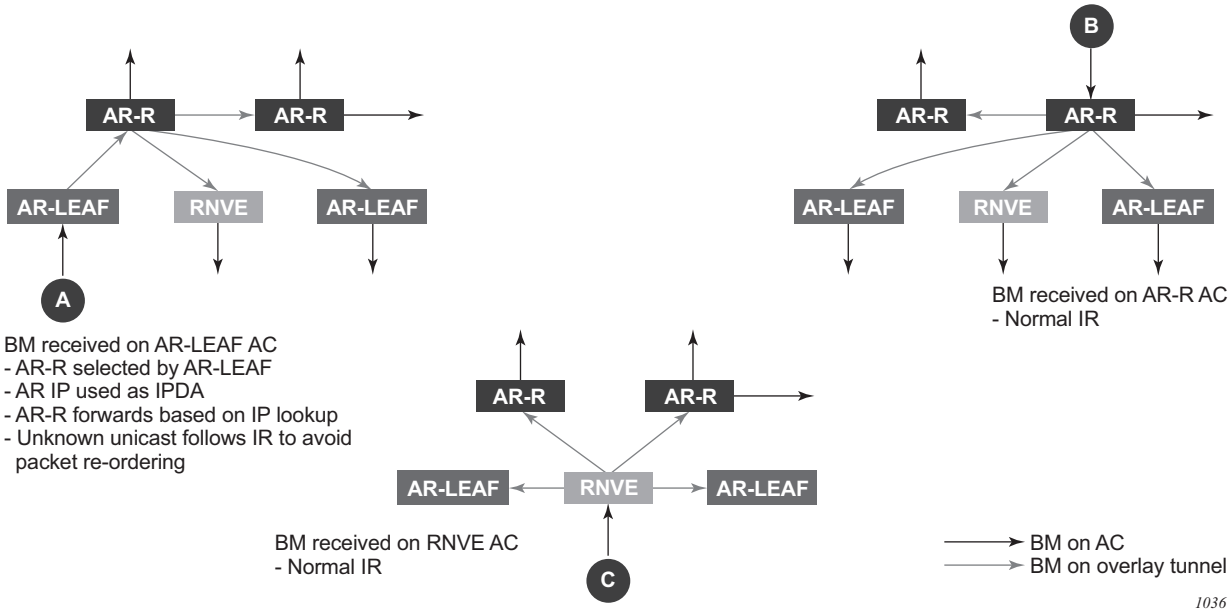

5.2.4.3 Assisted-Replication Interaction with Other VPLS Features

The Assisted-Replication feature has the following limitations.

- The following features are not supported on the same service where the Assisted-Replication feature is enabled.
  - Aggregate QoS per VNI
  - VXLAN IPv6 transport
  - IGMP-snooping
- Assisted-Replication Leaf and Replicator functions are mutually exclusive within the same VPLS service.
- The Assisted-Replication feature is supported with IPv4 non-system-ip VXLAN termination. However, the configured assisted-replication-ip (AR-IP) must be different from the tunnel termination IP address.
- The AR-IP address must be a /32 loopback interface on the base router.
- The Assisted-Replication feature is only supported in EVPN-VXLAN services (VPLS with bgp-evpn vxlan enabled). Although services with a combination of EVPN-MPLS and EVPN-VXLAN are supported, the Assisted-Replication configuration is only relevant to the VXLAN.

5.2.5 DC GW Policy Based Forwarding/Routing to an EVPN ESI (Ethernet Segment Identifier)

The Nuage VSP (Virtual Services Platform) supports a service chaining function that ensures traffic traverses a number of services (also known as Service Functions) between application hosts (FW, LB, NAT, IPS/IDS, and so on.) if the operator meeds to do so. In the DC, tenants want the ability to specify these functions and their sequence, so that services can be added or removed without requiring changes to the underlying application.

This service chaining function is built based on a series of policy based routing/ forwarding redirecting rules that are automatically coordinated and abstracted by the Nuage VSD (Virtual Services Directory). From a networking perspective, the packets are 'hop-by-hop' redirected based on the location of the corresponding SF (Service Function) in the DC fabric. The location of the SF is specified by its VTEP and VNI and is advertised by BGP-EVPN along with an Ethernet Segment Identifier (ESI) that is uniquely associated with the SF.

Refer to the Nuage VSP documentation for more information about the Nuage Service Chaining solution.

The 7750 SR, 7450 ESS, or 7950 XRS can be integrated as the first hop in the chain in a Nuage DC. This service chaining integration is intended to be used as described in the following three use-cases.

5.2.5.1 Policy Based Forwarding in VPLS Services for Nuage Service Chaining Integration in L2-Domains

Figure 152 shows the 7750 SR, 7450 ESS, and 7950 XRS Service Chaining integration with the Nuage VSP on VPLS services. In this example, the DC GW, PE1, is connected to an L2-DOMAIN that exists in the DC and must redirect the traffic to the Service Function SF-1. The regular Layer 2 forwarding procedures would have taken the packets to PE2, as opposed to SF-1.

Figure 152 **PBF to ESI Function**

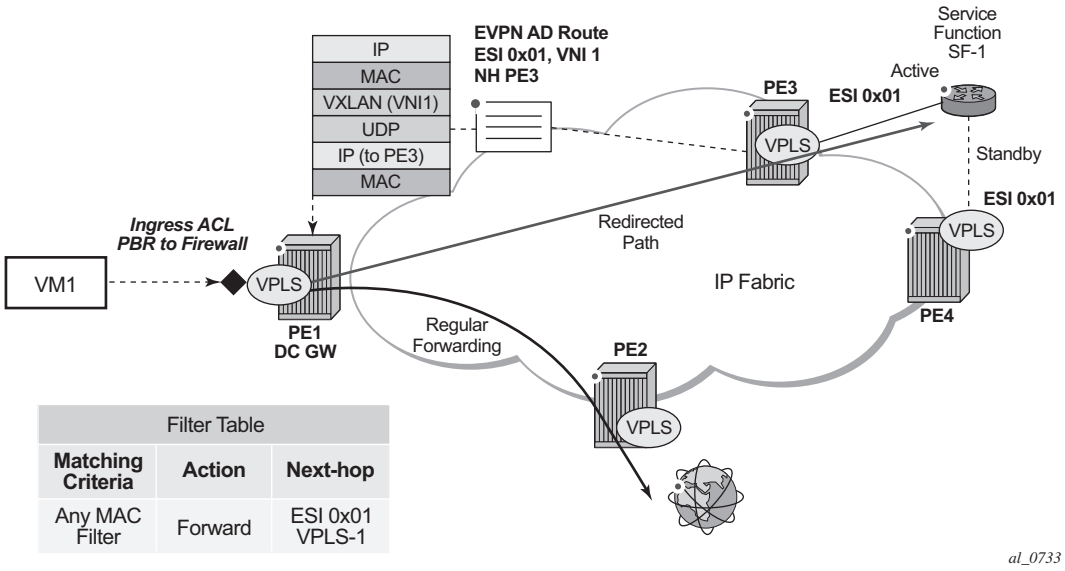

An operator must configure a PBF match/action filter policy entry in an IPv4 or MAC ingress access or network filter deployed on a VPLS interface using CLI/SNMP/ NETCONF management interfaces. The PBF target is the first service function in the chain (SF-1) that is identified by an Ethernet Segment Identifier.

In the example shown in Figure 152, the PBF filter will redirect the matching packets to ESI 0x01 in VPLS-1.

Note: Figure 152 represents ESI as '0x01' for simplicity; in reality, the ESI is a 10-byte number.

As soon as the redirection target is configured and associated with the vport connected to SF-1, the Nuage VSC (Virtual Services Controller, or the remote PE3 in the example) advertises the location of SF-1 via an Auto-Discovery Ethernet Tag route (route type 1) per-EVI. In this AD route, the ESI associated with SF-1 (ESI 0x01) is advertised along with the VTEP (PE3's IP) and VNI (VNI-1) identifying the vport where SF-1 is connected. PE1 will send all the frames matching the ingress filter to PE3's VTEP and VNI-1.

Note: When packets get to PE3, VNI-1 (the VNI advertised in the AD route) will indicate that a cut-through switching operation is needed to deliver the packets straight to the SF-1 vport, without the need for a regular MAC lookup.

The following filter configuration shows an example of PBF rule redirecting all the frames to an ESI.

```
A:PE1>config>filter>mac-filter# info
----------------------------------------------
            default-action forward
            entry 10 create
                action
                    forward esi ff:00:00:00:00:00:00:00:01 service-id 301
                exit
            exit
```

When the filter is properly applied to the VPLS service (VPLS-301 in this example), it will show 'Active' in the following show commands as long as the Auto-Discovery route for the ESI is received and imported.

```
A:PE1# show filter mac 1
===============================================================================
Mac Filter
===============================================================================
Filter Id   : 1                                 Applied         : Yes
Scope       : Template                          Def. Action     : Forward
Entries     : 1                                 Type            : normal
Description : (Not Specified)
-------------------------------------------------------------------------------
Filter Match Criteria : Mac
-------------------------------------------------------------------------------
Entry       : 10                                FrameType       : Ethernet
Description : (Not Specified)
Log Id      : n/a
Src Mac     : Undefined
Dest Mac    : Undefined
Dot1p       : Undefined                         Ethertype       : Undefined
DSAP        : Undefined                         SSAP            : Undefined
Snap-pid    : Undefined                         ESnap-oui-zero  : Undefined
Match action: Forward (ESI) Active
  ESI       : ff:00:00:00:00:00:00:00:01
  Svc Id    : 301
PBR Down Act: Forward (entry-default)
Ing. Matches: 3 pkts
Egr. Matches: 0 pkts
===============================================================================

A:PE1# show service id 301 es-pbr
===============================================================================
L2 ES PBR
===============================================================================
ESI                           Users      Status
                                         VTEP:VNI
-------------------------------------------------------------------------------
ff:00:00:00:00:00:00:00:01 1             Active
                                         192.0.2.72:7272
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

Details of the received AD route that resolves the filter forwarding are shown in the following '**show router bgp routes**' command.

```
A:PE1# show router bgp routes evpn auto-
disc esi ff:00:00:00:00:00:00:00:00:01
===============================================================================
 BGP Router ID:192.0.2.71       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup
 Origin codes  : i - IGP, e - EGP, ? - incomplete
===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.          ESI                             NextHop
      Tag                                                  Label
-------------------------------------------------------------------------------

u*>i  192.0.2.72:100       ff:00:00:00:00:00:00:00:00:01 192.0.2.72
      0                                                    VNI 7272

-------------------------------------------------------------------------------
Routes : 1
===============================================================
```

This AD route, when used for PBF redirection, is added to the list of EVPN-VXLAN bindings for the VPLS service and shown as 'L2 PBR' type:

```
A:PE1# show service id 301 vxlan
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 301
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address           Egress VNI     Num. MACs    Mcast   Oper State   L2 PBR
-------------------------------------------------------------------------------
192.0.2.69             301            1            Yes     Up           No
192.0.2.72             301            1            Yes     Up           No
192.0.2.72             7272           0            No      Up           Yes
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 3
-------------------------------------------------------------------------------
===============================================================================
```

If the AD route is withdrawn, the binding will disappear and the filter will be inactive again. The user can control whether the matching packets are dropped or forwarded if the PBF target cannot be resolved by BGP.

5.2.5.2 Policy Based Routing in VPRN Services for Nuage Service Chaining Integration in L2-DOMAIN-IRB Domains

Figure 153 shows the 7750 SR, 7450 ESS, and 7950 XRS Service Chaining integration with the Nuage VSP on L2-DOMAIN-IRB domains. In this example, the DC GW, PE1, is connected to an L2-DOMAIN-IRB that exists in the DC and must redirect the traffic to the Service Function SF-1 with IP address 10.10.10.1. The regular Layer 3 forwarding procedures would have taken the packets to PE2, as opposed to SF-1.

Figure 153 **PBR to ESI Function**

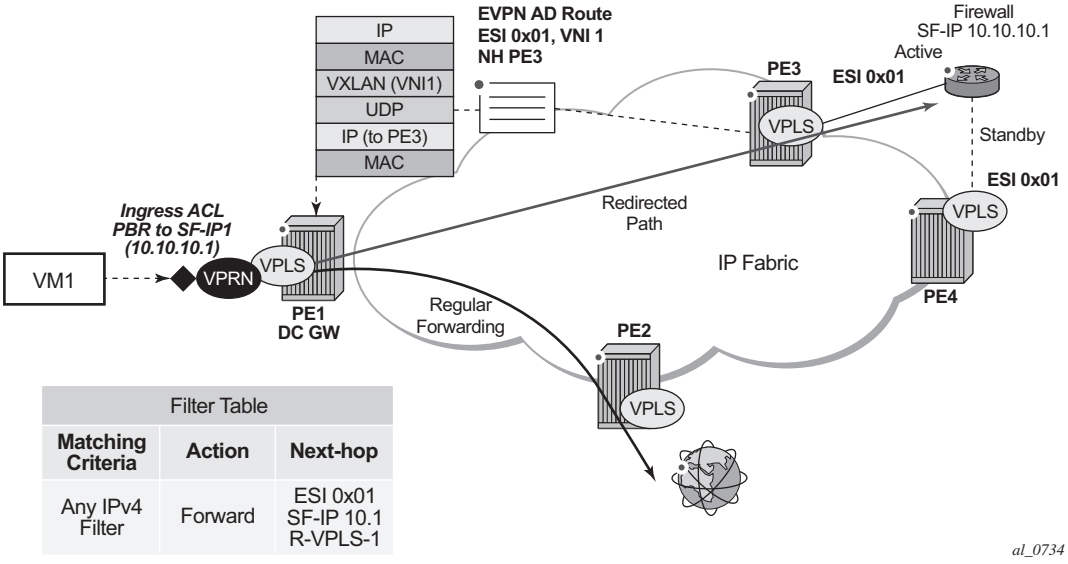

In this case, an operator must configure a PBR match/action filter policy entry in an IPv4 ingress access or network filter deployed on IES/VPRN interface using CLI, SNMP or NETCONF management interfaces. The PBR target identifies first service function in the chain (ESI 0x01 in Figure 153, identifying where the Service Function is connected and the IPv4 address of the SF) and EVPN VXLAN egress interface on the PE (VPRN routing instance and R-VPLS interface name). The BGP control plane together with ESI PBR configuration are used to forward the matching packets to the next-hop in the EVPN-VXLAN data center chain (through resolution to a VNI and VTEP). If the BGP control plane information is not available, the packets matching the ESI PBR entry will be, by default, forwarded using regular routing. Optionally, an operator can select to drop the packets when the ESI PBR target is not reachable.

The following filter configuration shows an example of a PBR rule redirecting all the matching packets to an ESI.

```
*A:PE1>config>filter>ip-filter# info
----------------------------------------------
```

```
            default-action forward
            entry 10 create
                match
                    dst-ip 10.10.10.253/32
                exit
                action
                    forward esi ff:00:00:00:00:21:5f:00:df:e5 sf-ip 10.10.10.1 vas-
interface "evi-301" router 300
                exit
                pbr-down-action-override filter-default-action
            exit
-----------------------------------------------
```

In this use case, the following are required in addition to the ESI: the **sf-ip** (10.10.10.1 in the example above), **router** instance (300), and **vas-interface**.

The **sf-ip** is used by the system to know which inner MAC DA it has to use when sending the redirected packets to the SF. The SF-IP will be resolved to the SF MAC following regular ARP procedures in EVPN-VXLAN.

The **router** instance may be the same as the one where the ingress filter is configured or may be different: for instance, the ingress PBR filter can be applied on an IES interface pointing at a VPRN router instances that is connected to the DC fabric.

The **vas-interface** refers to the R-VPLS interface name through which the SF can be found. The VPRN instance may have more than one R-VPLS interface, therefore, it is required to specify which R-VPLS interface to use.

When the filter is properly applied to the VPRN or IES service (VPRN-300 in this example), it will show 'Active' in the following show commands as long as the Auto-Discovery route for the ESI is received and imported and the SF-IP resolved to a MAC address.

```
*A:PE1# show filter ip 1

===============================================================================
IP Filter
===============================================================================
Filter Id    : 1                                Applied        : Yes
Scope        : Template                         Def. Action    : Forward
System filter: Unchained
Radius Ins Pt: n/a
CrCtl. Ins Pt: n/a
RadSh. Ins Pt: n/a
PccRl. Ins Pt: n/a
Entries      : 1
Description  : (Not Specified)
-------------------------------------------------------------------------------
Filter Match Criteria : IP
-------------------------------------------------------------------------------
Entry        : 10
Description  : (Not Specified)
```

```
Log Id       : n/a
Src. IP      : 0.0.0.0/0
Src. Port    : n/a
Dest. IP     : 172.16.0.253/32
Dest. Port   : n/a
Protocol     : Undefined                        Dscp           : Undefined
ICMP Type    : Undefined                        ICMP Code      : Undefined
Fragment     : Off                              Src Route Opt  : Off
Sampling     : Off                              Int. Sampling  : On
IP-Option    : 0/0                              Multiple Option: Off
TCP-syn      : Off                              TCP-ack        : Off
Option-pres  : Off
Egress PBR   : Undefined
Match action : Forward (ESI) Active
  ESI        : ff:00:00:00:00:21:5f:00:df:e5
  SF IP      : 10.10.10.1
  VAS If name: evi-301
  Router     : 300
PBR Down Act : Forward (filter-default-action) Ing. Matches : 3 pkts (318 bytes)
Egr. Matches : 0 pkts
===============================================================================

*A:PE1# show service id 300 es-pbr
===============================================================================
L3 ES PBR
===============================================================================
SF IP              ESI                                 Users Status
                   Interface                                 MAC
                                                             VTEP:VNI
-------------------------------------------------------------------------------
10.10.10.1         ff:00:00:00:00:21:5f:00:df:e5       1     Active
                   evi-301                                   d8:47:01:01:00:0a
                                                             192.0.2.71:7171
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

In the FDB for the R-VPLS 301, the MAC address is associated with the VTEP and VNI specified by the AD route, and not by the MAC/IP route anymore. When a PBR filter with a forward action to an ESI and SF-IP (Service Function IP) exists, a MAC route is auto-created by the system and this route has higher priority that the remote MAC, or IP routes for the MAC (see section BGP and EVPN route selection for EVPN routes).

The following shows that the AD route creates a new EVPN-VXLAN binding and the MAC address associated with the SF-IP uses that 'binding':

```
*A:PE1# show service id 301 vxlan
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 301
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address           Egress VNI     Num. MACs    Mcast   Oper State   L2 PBR
-------------------------------------------------------------------------------
```

```
192.0.2.69             301            1            Yes     Up          No
192.0.2.71             301            0            Yes     Up          No
192.0.2.71             7171           1            No      Up          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 3
-------------------------------------------------------------------------------
===============================================================================
*A:PE1# show service id 301 fdb detail
===============================================================================
Forwarding Database, Service 301
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
301       d8:45:ff:00:00:6a vxlan:                   EvpnS    06/15/15 21:55:27
                            192.0.2.69:301
301       d8:47:01:01:00:0a vxlan:                   EvpnS    06/15/15 22:32:56
                            192.0.2.71:7171
301       d8:48:ff:00:00:6a cpm                      Intf     06/15/15 21:54:12
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

For Layer 2, if the AD route is withdrawn or the SF-IP ARP not resolved, the filter will be inactive again. The user can control whether the matching packets are dropped or forwarded if the PBF target cannot be resolved by BGP.

5.3 EVPN for MPLS Tunnels

This section provides information about EVPN for MPLS tunnels.

5.3.1 BGP-EVPN Control Plane for MPLS Tunnels

Table 81 lists all the EVPN routes supported in 7750 SR, 7450 ESS, or 7950 XRS SR OS and their usage in EVPN-VXLAN, EVPN-MPLS, and PBB-EVPN.

Note: Route type 1 is not required in PBB-EVPN as per RFC 7623.

Table 81 **EVPN Routes and Usage**

EVPN Route	Usage	EVPN-VXLAN	EVPN-MPLS	PBB-EVPN
Type 1 - Ethernet Auto-Discovery route (A-D)	Mass-withdraw, ESI labels, Aliasing	—	Y	—
Type 2 - MAC/IP Advertisement route	MAC/IP advertisement, IP advertisement for ARP resolution	Y	Y	Y
Type 3 - Inclusive Multicast Ethernet Tag route	Flooding tree setup (BUM flooding)	Y	Y	Y
Type 4 - Ethernet Segment route	ES discovery and DF election	—	Y	Y
Type 5 - IP Prefix advertisement route	IP Routing	Y	Y	—

RFC 7432 describes the BGP-EVPN control plane for MPLS tunnels. If EVPN multi-homing is not required, two route types are needed to set up a basic EVI (EVPN Instance): MAC/IP Advertisement and the Inclusive Multicast Ethernet Tag routes. If multi-homing is required, the Ethernet Segment and the Auto-Discovery routes are also needed.

The route fields and extended communities for route types 2 and 3 are shown in Figure 145. BGP-EVPN Control Plane for VXLAN Overlay Tunnels The changes compared to their use in EVPN-VXLAN are described below.

EVPN Route Type 3 – Inclusive Multicast Ethernet Tag Route

As in EVPN-VXLAN, route type 3 is used for setting up the flooding tree (BUM flooding) for a specified VPLS service. The received inclusive multicast routes add entries to the VPLS flood list in the 7750 SR, 7450 ESS, and 7950 XRS. Ingress replication, p2mp mLDP, and composite tunnels are supported as tunnel types in route type 3 when BGP-EVPN MPLS is enabled

The following route values are used for EVPN-MPLS services:

- Route Distinguisher: taken from the RD of the VPLS service within the BGP context. The RD can be configured or derived from the **bgp-evpn evi** value.
- Ethernet Tag ID: 0.
- IP address length: always 32.
- Originating router's IP address: carries the system address (IPv4 only).
- PMSI attribute: the PMSI attribute can have different formats depending on the tunnel type enabled in the service.
  - Tunnel type = Ingress replication (6)

    The route is referred to as an Inclusive Multicast Ethernet Tag IR (IMET-IR) route and the PMSI Tunnel Attribute (PTA) fields are populated as follows:
    - Flags—Leaf not required.
    - MPLS label—Carries the MPLS label allocated for the service in the high-order 20 bits of the label field.

      Unless **bgp-evpn mpls ingress-replication-bum-label** is configured in the service, the MPLS label used will be the same as that used in the MAC/IP routes for the service.
    - Tunnel endpoint—Equal to the originating IP address.
  - Tunnel type=p2mp mLDP (2)

    The route is referred to as an IMET-P2MP route and its PTA fields are populated as follows.
    - Flags—Leaf not required.
    - MPLS label—0.
    - Tunnel endpoint—Includes the route node address and an opaque number. This is the tunnel identifier that the leaf-nodes will use to join the mLDP P2MP tree.
  - Tunnel type=Composite tunnel (130)

    The route is referred to as an IMET-P2MP-IR route and its PTA fields are populated as follows.
    - Flags—Leaf not required.
    - MPLS label 1— 0.

- Tunnel endpoint identifier will include the following:

  MPLS label2—Non-zero, downstream allocated label (like any other IR label). The leaf-nodes will use the label to set up an EVPN-MPLS destination to the root and add it to the default-multicast list.

  mLDP tunnel identifier—The route node address and an opaque number. This is the tunnel identifier that the leaf-nodes will use to join the mLDP P2MP tree.

IMET-P2MP-IR routes are used in EVIs with a few root nodes and a significant number of leaf-only PEs. In this scenario, a combination of P2MP and IR tunnels can be used in the network, such that the root nodes use P2MP tunnels to send broadcast, Unknown unicast, and Multicast traffic but the leaf-PE nodes use IR to send traffic to the roots. This use-case is documented in IETF Draft *draft-ietf-bess-evpn-etree* and the main advantage it offers is the significant savings in P2MP tunnels that the PE/P routers in the EVI need to handle (as opposed to a full mesh of P2MP tunnels among all the PEs in an EVI).

In this case, the root PEs will signal a special tunnel type in the PTA, indicating that they intend to transmit BUM traffic using an mLDP P2MP tunnel but they can also receive traffic over an IR evpn-mpls binding. An IMET route with this special "composite" tunnel type in the PTA is called an IMET-P2MP-IR route and the encoding of its PTA is shown in Figure 154.

Figure 154 **Composite p2mp mLDP and IR Tunnels—PTA**

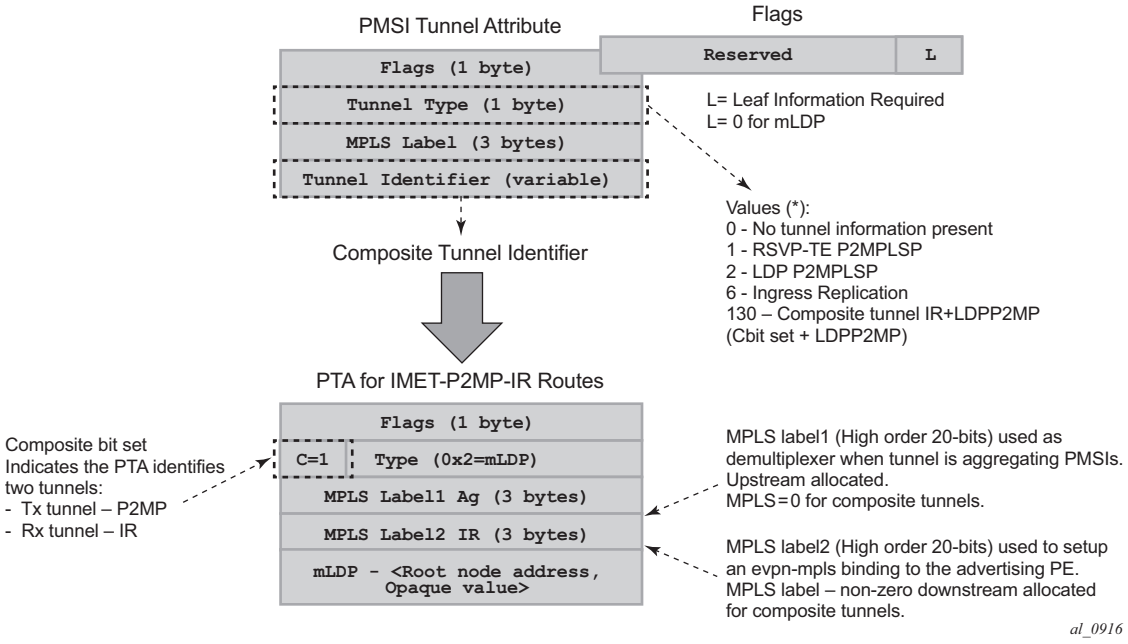

EVPN Route Type 2 - MAC/IP Advertisement Route

The 7750 SR, 7450 ESS, or 7950 XRS router generates this route type for advertising MAC addresses (and IP addresses if proxy-ARP/proxy-ND is enabled). The router generates MAC advertisement routes for the following:

- Learned MACs on SAPs or sdp-bindings—if mac-advertisement is enabled.
- Conditional static MACs—if mac-advertisement is enabled.

Note: The **unknown-mac-route** is not supported for EVPN-MPLS services.

The route type 2 generated by a router uses the following fields and values:

- Route Distinguisher: Taken from the RD of the VPLS service within the BGP context. The RD can be configured or derived from the **bgp-evpn evi** value.
- Ethernet Segment Identifier (ESI): Zero for MACs learned from single-homed CEs and different from zero for MACs learned from multi-homed CEs.
- Ethernet Tag ID: 0.
- MAC address length: Always 48.
- MAC Address learned or statically configured.
- IP address and IP address length:
  - It will be the IP address associated with the MAC being advertised with a length of 32 (or 128 for IPv6).
  - In general, any MAC route without IP will have IPL=0 (IP length) and the IP will be omitted.
  - When received, any IPL value not equal to zero, 32, or 128 will discard the route.
  - MPLS Label 1: Carries the MPLS label allocated by the system to the VPLS service. The label value is encoded in the high-order 20 bits of the field and will be the same label used in the routes type 3 for the same service unless **bgp-evpn mpls ingress-replication-bum-label** is configured in the service.
- MPLS Label 2: 0.
- The MAC Mobility extended community: Used for signaling the sequence number in case of mac moves and the sticky bit in case of advertising conditional static MACs. If a MAC route is received with a MAC mobility **ext-community**, the sequence number and the 'sticky' bit are considered for the route selection.

When EVPN multi-homing is enabled in the system, two more routes are required. Figure 155 shows the fields in routes type 1 and 4 and their associated extended communities.

Figure 155 **EVPN Routes Type 1 and 4**

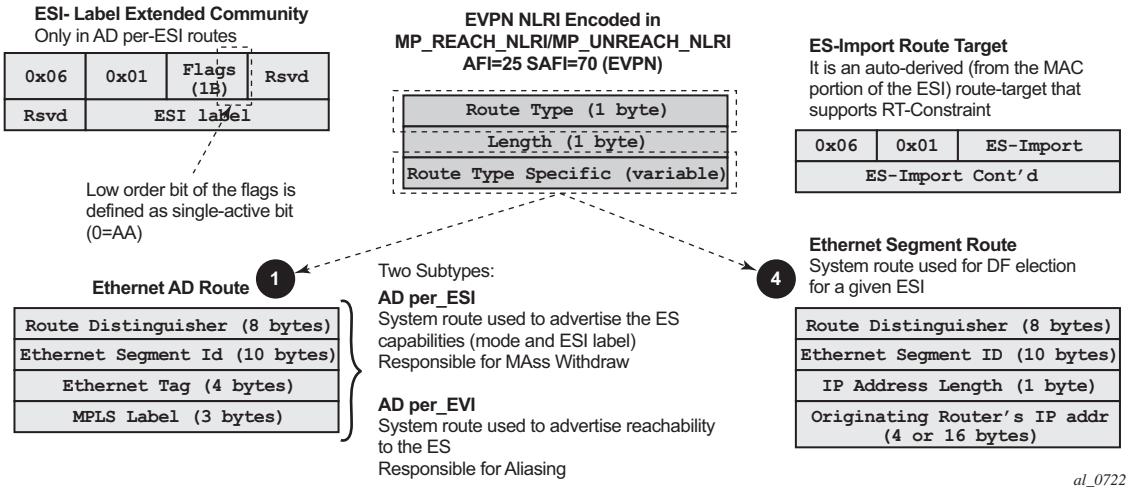

EVPN Route Type 1 - Ethernet Auto-discovery Route (AD route)

The 7750 SR, 7450 ESS, or 7950 XRS router generates this route type for advertising for multi-homing functions. The system can generate two types of AD routes:

- Ethernet AD route per-ESI (Ethernet Segment ID)
- Ethernet AD route per-EVI (EVPN Instance)

The Ethernet AD per-ESI route generated by a router uses the following fields and values:

- Route Distinguisher: Taken from the system level RD or service level RD.
- Ethernet Segment Identifier (ESI): Will contain a 10-byte identifier as configured in the system for a specified **ethernet-segment**.
- Ethernet Tag ID: MAX-ET (0xFFFFFFFF). This value is reserved and used only for AD routes per ESI.
- MPLS label: 0.
- ESI Label Extended community: Includes the single-active bit (0 for all-active and 1 for single-active) and ESI label for all-active multi-homing split-horizon.
- Route-target extended community: Taken from the service level RT or an RT-set for the services defined on the ethernet-segment.

The system can either send a separate Ethernet AD per-ESI route per service, or a few Ethernet AD per-ESI routes aggregating the route-targets for multiple services. While both alternatives will inter-operate, RFC 7432 states that the EVPN Auto-Discovery per-ES route must be sent with a set of route-targets corresponding to all the EVIs defined on the ethernet-segment. Either option can be enabled using the command: **config>service>system>bgp-evpn#ad-per-es-route-target <[evi-rt] | [evi-rt-set route-distinguisher <***ip-address***>]>**

The default option **ad-per-es-route-target evi-rt** configures the system to send a separate AD per-ES route per service. When enabled, the **evi-rt-set** option allows the aggregation of routes: A single AD per-ES route with the associated RD (ip-address:1) and a set of EVI route-targets will be advertised (to a maximum of 128). When the number of EVIs defined in the ethernet-segment is significant (hence the number of route-targets), the system will send more than one route. For example:

- AD per-ES route for evi-rt-set 1 will be sent with RD ip-address:1
- AD per-ES route for evi-rt-set 2 will be sent with RD ip-address:2

Note: When **evi-rt-set** is configured, no vsi-export policies are possible on the services defined on the ethernet-segment. If vsi-export policies are configured for a service, the system will send an individual AD per-ES route for that service. The maximum standard BGP update size is 4KB, with a maximum of 2KB for the route-target extended community attribute.

The Ethernet AD per-EVI route generated by a router uses the following fields and values:

- Route Distinguisher: Taken from the service level RD.
- Ethernet Segment Identifier (ESI): Will contain a 10-byte identifier as configured in the system for a specified ethernet-segment.
- Ethernet Tag ID: 0.
- MPLS label: Encodes the unicast label allocated for the service (high-order 20 bits).
- Route-target extended community: Taken from the service level RT.

Note: The AD per-EVI route is not sent with the ESI label Extended Community.

EVPN Route Type 4 - Ethernet Segment Route (ES route)

The router generates this route type for multi-homing ES discovery and DF (Designated Forwarder) election.

- Route Distinguisher: Taken from the service level RD.
- Ethernet Segment Identifier (ESI): Will contain a 10-byte identifier as configured in the system for a specified **ethernet-segment**.
- ES-import route-target community: The value is automatically derived from the MAC address portion of the ESI. This extended community is treated as a route-target and is supported by RT-constraint (route-target BGP family).

EVPN Route Type 5 - IP Prefix Route

IP Prefix Routes are also supported for MPLS tunnels. The route fields for route type 5 are shown in Figure 147. The 7750 SR, 7450 ESS (mixed mode), or 7950 XRS router will generate this route type for advertising IP prefixes in EVPN using the same fields that are described in section BGP-EVPN Control Plane for VXLAN Overlay Tunnels, with the following exceptions:

- MPLS Label—Carries the MPLS label allocated for the service
- This route will be sent with the RFC 5512 tunnel encapsulation extended community with the tunnel type value set to MPLS

RFC 5512 - BGP Tunnel Encapsulation Extended Community

The following routes are sent with the RFC 5512 BGP Encapsulation Extended Community: MAC/IP, Inclusive Multicast Ethernet Tag, and AD per-EVI routes. ES and AD per-ESI routes are not sent with this Extended Community.

The router processes the following BGP Tunnel Encapsulation tunnel values registered by IANA for RFC 5512:

- VXLAN encapsulation: 8.
- MPLS encapsulation: 10.

Any other tunnel value will make the route 'treat-as-withdraw'.

If the encapsulation value is MPLS, the BGP will validate the high-order 20-bits of the label field, ignoring the low-order 4 bits. If the encapsulation is VXLAN, the BGP will take the entire 24-bit value encoded in the MPLS label field as the VNI.

If the encapsulation extended community (as defined in RFC 5512) is not present in a received route, BGP will treat the route as an MPLS or VXLAN-based configuration of the **config>router>bgp>neighbor# def-recv-evpn-encap [mpls | vxlan]** command. The command is also available at the **bgp** and **group** levels.

5.3.2 EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS)

EVPN can be used in MPLS networks where PEs are interconnected through any type of tunnel, including RSVP-TE, Segment-Routing TE, LDP, BGP, Segment Routing IS-IS, Segment Routing OSPF, or MPLSoUDP. As with VPRN services, the selection of the tunnel to be used in a VPLS service (with BGP-EVPN MPLS enabled) is based on the **auto-bind-tunnel** command.

EVPN-MPLS is modeled similar to EVPN-VXLAN, that is, using a VPLS service where EVPN-MPLS 'bindings' can coexist with SAPs and SDP-bindings. The following shows an example of a VPLS service with EVPN-MPLS.

```
*A:PE-1>config>service>vpls# info
----------------------------------------------
  description "evpn-mpls-service"
  bgp
  bgp-evpn
    evi 10
    vxlan
      shutdown
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  sap 1/1/1:1 create
  exit
  spoke-sdp 1:1 create
```

The user will configure a **bgp-evpn** context where **vxlan** must be shutdown and **mpls no shutdown**. In addition to the **mpls no shutdown** command, the minimum set of commands to be configured to set up the EVPN-MPLS instance are the **evi** and the **auto-bind-tunnel resolution** commands. However, the user can configure some other options. The most relevant configuration options are described below.

evi {1..65535} — This EVPN identifier is unique in the system and will be used for the service-carving algorithm used for multi-homing (if configured) and auto-deriving route-target and route-distinguishers in the service. It can be used for EVPN-MPLS and EVPN-VXLAN services.

If this EVPN identifier is not specified, the value will be zero and no route-distinguisher or route-targets will be auto-derived from it. If specified and no other route-distinguisher/route-target are configured in the service:, then the following applies:

- The route-distinguisher is derived from: **<system_ip>:evi**
- The route-target is derived from: **<autonomous-system>:evi**

Note: When the vsi-import/export polices are configured, the route-target must be configured in the policies and those values take preference over the auto-derived route-targets. The operational route-target for a service will be displayed by the **show service id x bgp** command. If the **bgp-ad**>**vpls-id** is configured in the service, the **vpls-id** derived route-target takes precedence over the evi-derived route-target.

When the **evi** is configured, a **config>service>vpls>bgp** node (even empty) is required to allow the user to see the correct information on the **show service id 1 bgp** and **show service system bgp-route-distinguisher** commands.

Although not mandatory, if no multi-homing is configured, the configuration of an **evi** is enforced for EVPN services with SAPs/SDP-bindings in an **ethernet-segment**. See the 'EVPN multi-homing' section for more information about **ethernet-segments**.

The following options are specific to EVPN-MPLS (and defined on **bgp-evpn>mpls**):

- **control-word:** Required as per RFC 7432 to avoid frame disordering. The user can enable/disable it so that interoperability to other vendors can be guaranteed.
- **auto-bind-tunnel:** Allows the user to decide what type of MPLS transport tunnels will be used for a particular instance. The command will be used in the same way as it is used in VPRN services.

  For bgp-evpn mpls, '**bgp**' is explicitly added to the **resolution-filter** in EVPN ('**bgp**' is implicit in VPRNs).
- **force-vlan-vc-forwarding:** This command will allow the system to preserve the vlan-id and pbits of the service-delimiting qtag in a new tag added in the customer frame before sending it to the EVPN core.

Note: This command may be used in conjunction with the **sap ingress vlan-translation** command. If so, the configured translated vlan-id will be the vlan-id sent to the EVPN binds as opposed to the service-delimiting tag vlan-id. If the ingress SAP/binding is 'null'-encapsulated, the output vlan-id and pbits will be zero.

- **split-horizon-group:** This command allows the association of a user-created split horizon group to all the EVPN-MPLS destinations. See the EVPN and VPLS integration section for more information.
- **ecmp:** When this command is set to a value greater than 1, aliasing is activated to the remote PEs that are defined in the same all-active multi-homing ethernet-segment. See the EVPN multi-homing section for more information.

- **ingress-replication-bum-label:** This command is only enabled when the user wants the PE to advertise a label for BUM traffic (Inclusive Multicast routes) that is different from the label advertised for unicast traffic (with the MAC/IP routes). This is useful to avoid potential transient packet duplication in all-active multi-homing.

In addition to these options, the following bgp-evpn commands are also available for EVPN-MPLS services:

- **[no] mac-advertisement**
- **mac-duplication and settings**

When EVPN-MPLS is established among some PEs in the network, EVPN unicast and multicast 'bindings' are created on each PE to the remote EVPN destinations. A specified ingress PE will create:

- A unicast EVPN-MPLS destination binding to a remote egress PE as soon as a MAC/IP route is received from that egress PE.
- A multicast EVPN-MPLS destination binding to a remote egress PE, if and only if the egress PE advertises an Inclusive Multicast Ethernet Tag Route with a BUM label. That is only possible if the egress PE is configured with **ingress-replication-bum-label**.

Those bindings, as well as the MACs learned on them, can be checked through the following show commands. In the following example, the remote PE(192.0.2.69) is configured with **no ingress-replication-bum-label** and PE(192.0.2.70) is configured with **ingress-replication-bum-label**. Hence, Dut has a single EVPN-MPLS destination binding to PE(192.0.2.69) and two bindings (unicast and multicast) to PE(192.0.2.70).

```
*A:Dut# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                 Transport
-------------------------------------------------------------------------------
192.0.2.69      262118        1           Yes             06/11/2015 19:59:03
                ldp
192.0.2.70      262139        0           Yes             06/11/2015 19:59:03
                ldp
192.0.2.70      262140        1           No              06/11/2015 19:59:03
                ldp
192.0.2.72      262140        0           Yes             06/11/2015 19:59:03
                ldp
192.0.2.72      262141        1           No              06/11/2015 19:59:03
                ldp
192.0.2.73      262139        0           Yes             06/11/2015 19:59:03
                ldp
192.0.2.254     262142        0           Yes             06/11/2015 19:59:03
```

```
               bgp
-------------------------------------------------------------------------------
Number of entries : 7
-------------------------------------------------------------------------------
===============================================================================

*A:Dut# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:ca:fe:ca:fe:69 eMpls:                   EvpnS    06/11/15 21:53:48
                            192.0.2.69:262118
1         00:ca:fe:ca:fe:70 eMpls:                   EvpnS    06/11/15 19:59:57
                            192.0.2.70:262140
1         00:ca:fe:ca:fe:72 eMpls:                   EvpnS    06/11/15 19:59:57
                            192.0.2.72:262141
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

5.3.2.1 EVPN and VPLS Integration

The 7750 SR, 7450 ESS, or 7950 XRS router SR OS EVPN implementation supports *draft-ietf-bess-evpn-vpls-seamless-integ* so that EVPN-MPLS and VPLS can be integrated into the same network and within the same service. Since EVPN will not be deployed in green-field networks, this feature is useful for the integration between both technologies and even for the migration of VPLS services to EVPN-MPLS.

The following behavior enables the integration of EVPN and sdp-bindings in the same VPLS network:

a) Systems with EVPN endpoints and sdp-bindings to the same far-end bring down the sdp-bindings.

- The router will allow the establishment of an EVPN endpoint and a SDP-binding to the same far-end but the SDP-binding will be kept operationally down. Only the EVPN endpoint will be operationally up. This is true for spoke-sdps (manual and BGP-AD) and mesh-sdps. It is also possible between VXLAN and SDP-bindings.
- If there is an existing EVPN endpoint to a specified far-end and a spoke-SDP establishment is attempted, the spoke-SDP will be setup but kept down with an operational flag indicating that there is an EVPN route to the same far-end.

- If there is an existing spoke-SDP and a valid/used EVPN route arrives, the EVPN endpoint will be setup and the spoke-SDP will be brought down with an operational flag indicating that there is an EVPN route to the same far-end.
- In the case of an SDP-binding and EVPN endpoint to different far-end IPs on the same remote PE, both links will be up. This can happen if the SDP-binding is terminated in an IPv6 address or IPv4 address different from the system address where the EVPN endpoint is terminated.

b) The user can add spoke-SDPs and all the EVPN-MPLS endpoints in the same split horizon group (SHG).

- A CLI command is added under the **bgp-evpn**>**mpls**> context so that the EVPN-MPLS endpoints can be added to a split horizon group:
  - bgp-evpn>mpls> [no] split-horizon-group <group-name>
- The **bgp-evpn mpls split-horizon-group** must reference a user-configured split horizon group. User-configured split horizon groups can be configured within the service context. The same **group-name** can be associated with SAPs, spoke-SDPs, pw-templates, pw-template-bindings, and EVPN-MPLS endpoints.
- If the **split-horizon-group** command in **bgp-evpn**>**mpls**> is not used, the default split horizon group (that contains all the EVPN endpoints) is still used, but it will not be possible to refer to it on saps/spoke-sdps.
- SAPs and SDP-bindings that share the same split horizon group of the EVPN-MPLS provider-tunnel will be brought operationally down if the point-to-multipoint tunnel is operationally up.

c) The system disables the advertisement of MACs learned on spoke-sdps/ saps that are part of an EVPN split horizon group.

- When the saps and/or spoke-sdps (manual or BGP-AD-discovered) are configured within the same split horizon group as the EVPN endpoints, MAC addresses will still be learned on them, but they will not be advertised in EVPN.
- The preceding statement is also true if proxy-ARP/proxy-ND is enabled and an IP->MAC pair is learned on a SAP or SDP-binding that belongs to the EVPN split horizon group.
- The SAPs and/or spoke-SDPs added to an EVPN split horizon group should not be part of any EVPN multi-homed ES. If that happened, the PE would still advertise the AD per-EVI route for the SAP and/or spoke-SDP, attracting EVPN traffic that could not possibly be forwarded to that SAP and/or sdp-binding.
- Similar to the preceding statement, a split horizon group composed of SAPs/ sdp-bindings used in a BGP-MH site should not be configured under **bgp-evpn>mpls>split-horizon-group**. This misconfiguration would prevent traffic being forwarded from the EVPN to the BGP-MH site, regardless of the DF/NDF state.

Figure 156 shows an example of EVPN-VPLS integration.

Figure 156 **EVPN-VPLS Integration**

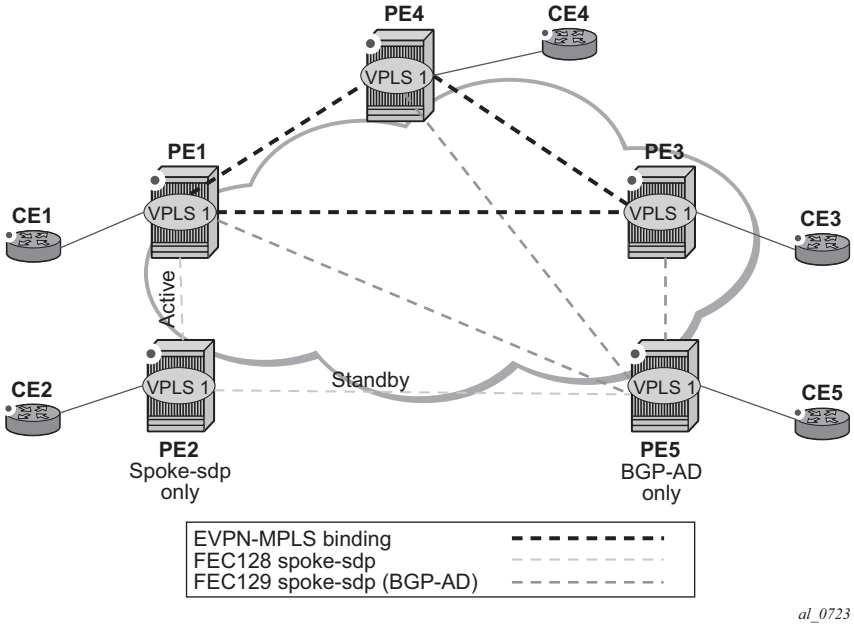

An example CLI configuration for PE1, PE5, and PE2 is provided below.

```
*A:PE1>config>service# info
----------------------------------------------
pw-template 1 create
vpls 1 customer 1 create
  split-horizon-group "SHG-1" create
  bgp
    route-target target:65000:1
    pw-template-binding 1 split-horizon-group SHG-1
  bgp-ad
    no shutdown
    vpls-id 65000:1
  bgp-evpn
    evi 1
    mpls
      no shutdown
      split-horizon-group SHG-1
  spoke-sdp 12:1 create
  exit
  sap 1/1/1:1 create
  exit

*A:PE5>config>service# info
----------------------------------------------
pw-template 1 create
vpls 1 customer 1 create
  bgp
    route-target target:65000:1
```

```
    pw-template-binding 1 split-horizon-group SHG-1 # auto-created SHG
  bgp-ad
    no shutdown
    vpls-id 65000:1
  spoke-sdp 52:1 create
  exit

*A:PE2>config>service# info
----------------------------------------------
vpls 1 customer 1 create
  end-point CORE create
    no suppress-standby-signaling
  spoke-sdp 21:1 end-point CORE
    precedence primary
  spoke-sdp 25:1 end-point CORE
```

- PE1, PE3, and PE4 have BGP-EVPN and BGP-AD enabled in VPLS-1. PE5 has BGP-AD enabled and PE2 has active/standby spoke-sdps to PE1 and PE5.

  In this configuration:

  - PE1, PE3, and PE4 will attempt to establish BGP-AD spoke-sdps, but they will be kept operationally down as long as there are EVPN endpoints active among them.
  - BGP-AD spoke-sdps and EVPN endpoints are instantiated within the same split horizon group, for example, SHG-1.
  - Manual spoke-sdps from PE1 and PE5 to PE2 are not part of SHG-1.

- EVPN MAC advertisements:
  - MACs learned on FEC128 spoke-sdps are advertised normally in EVPN.
  - MACs learned on FEC129 spoke-sdps are not advertised in EVPN (because they are part of SHG-1, which is the split horizon group used for **bgp-evpn>mpls**). This prevents any data plane MACs learned on the SHG from being advertised in EVPN.
- BUM operation on PE1:
  - When CE1 sends BUM, PE1 will flood to all the active bindings.
  - When CE2 sends BUM, PE2 will send it to PE1 (active spoke-sdp) and PE1 will flood to all the bindings and saps.
  - When CE5 sends BUM, PE5 will flood to the three EVPN PEs. PE1 will flood to the active spoke-sdp and saps, never to the EVPN PEs because they are part of the same SHG.

5.3.2.2 Auto-Derived Route-Distinguisher (RD) in Services with Multiple BGP Families

In a VPLS service, multiple BGP families and protocols can be enabled at the same time. When **bgp-evpn** is enabled, **bgp-ad** and **bgp-mh** are supported as well. A single RD is used per service and not per BGP family/protocol.

The following rules apply:

- The VPLS RD is selected based on the following precedence:
  - Manual RD or auto-rd always take precedence when configured.
  - If no manual/auto-rd configuration, the RD is derived from the **bgp-ad>vpls-id**.
  - If no manual/auto-rd/vpls-id configuration, the RD is derived from the **bgp-evpn>evi**, except for **bgp-mh**, which does not support evi-derived RD.
  - If no manual/auto-rd/vpls-id/evi configuration, there will not be RD and the service will fail.
- The selected RD (see above rules) will be displayed by the **Oper Route Dist** field of the **show service id bgp** command.
- The service supports dynamic RD changes, for instance, the CLI allows the vpls-id be changed dynamically, even if it is used to auto-derive the service RD for **bgp-ad**, **bgp-vpls**, or **bgp-mh**.

Note: When the RD changes, the active routes for that VPLS will be withdrawn and re-advertised with the new RD.

- If one of the mechanisms to derive the RD for a specified service is removed from the configuration, the system will select a new RD based on the above rules. For example, if the vpls-id is removed from the configuration, the routes will be withdrawn, the new RD selected from the evi, and the routes re-advertised with the new RD.

Note: This reconfiguration will fail if the new RD already exists in a different VPLS/epipe.

- Because the **vpls-id** takes precedence over the evi when deriving the RD automatically, adding **evpn** to an existing **bgp-ad** service will not impact the existing RD - this is important to support **bgp-ad** to **evpn** migration.

5.3.2.3 EVPN Multi-Homing in VPLS Services

EVPN multi-homing implementation is based on the concept of the **ethernet-segment**. An **ethernet-segment** is a logical structure that can be defined in one or more PEs and identifies the CE (or access network) multi-homed to the EVPN PEs. An **ethernet-segment** is associated with port, LAG, or SDP objects and is shared by all the services defined on those objects. In the case of virtual Ethernet segments, individual VID or VC-ID ranges can be associated to the port, LAG, or SDP objects defined in the **ethernet-segment**.

Each **ethernet-segment** has a unique identifier called **esi** (Ethernet Segment Identifier) that is 10 bytes long and is manually configured in the router.

Note: The **esi** is advertised in the control plane to all the PEs in an EVPN network; therefore, it is very important to ensure that the 10-byte **esi** value is unique throughout the entire network. Single-homed CEs are assumed to be connected to an ethernet-segment with esi = 0 (single-homed ethernet-segments are not explicitly configured).

This section describes the behavior of the EVPN multi-homing implementation in an EVPN-MPLS service.

5.3.2.3.1 EVPN All-Active Multi-Homing

As described in RFC 7432, all-active multi-homing is only supported on access LAG SAPs and it is mandatory that the CE is configured with a LAG to avoid duplicated packets to the network. LACP is optional.

Three different procedures are implemented in 7750 SR, 7450 ESS, and 7950 XRS SR OS to provide all-active multi-homing for a specified ethernet-segment:

- DF (Designated Forwarder) election
- Split-horizon
- Aliasing

Figure 157 shows the need for DF election in all-active multi-homing.

Figure 157 **DF Election**

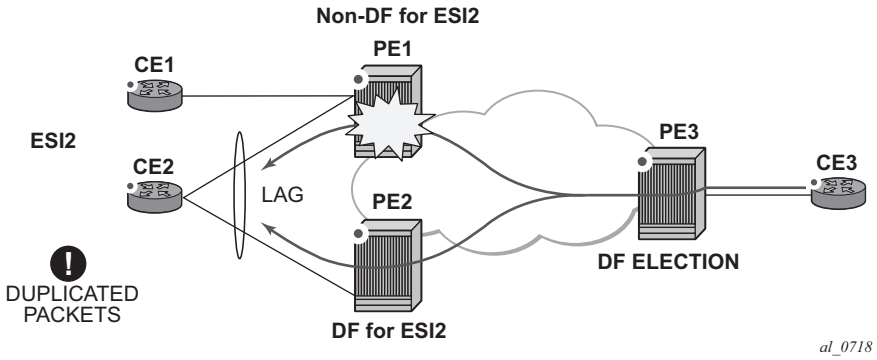

The DF election in EVPN all-active multi-homing avoids duplicate packets on the multi-homed CE. The DF election procedure is responsible for electing one DF PE per ESI per service; the rest of the PEs being non-DF for the ESI and service. Only the DF will forward BUM traffic from the EVPN network toward the ES SAPs (the multi-homed CE). The non-DF PEs will not forward BUM traffic to the local ethernet-segment SAPs.

Note: BUM traffic from the CE to the network and known unicast traffic in any direction is allowed on both the DF and non-DF PEs.

Figure 158 shows the EVPN split-horizon concept for all-active multi-homing.

Figure 158 **Split-Horizon**

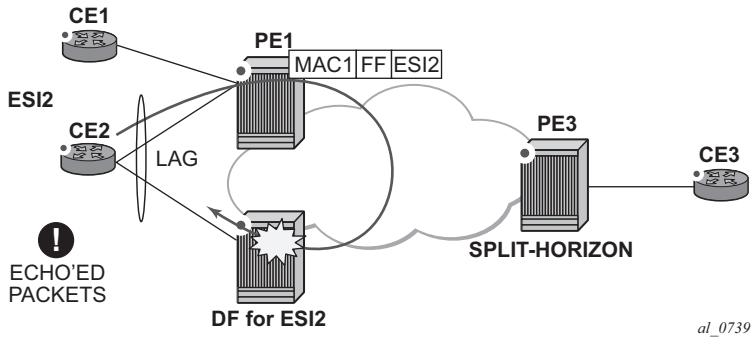

The EVPN split-horizon procedure ensures that the BUM traffic originated by the multi-homed PE and sent from the non-DF to the DF, is not replicated back to the CE (echoed packets on the CE). To avoid these echoed packets, the non-DF (PE1) will send all the BUM packets to the DF (PE2) with an indication of the source ethernet-segment. That indication is the ESI Label (ESI2 in the example), previously signaled by PE2 in the AD per-ESI route for the ethernet-segment. When PE2 receives an EVPN packet (after the EVPN label lookup), the PE2 will find the ESI label that will identify its local ethernet-segment ESI2. The BUM packet will be replicated to other local CEs but not to the ESI2 SAP.

Figure 159 shows the EVPN aliasing concept for all-active multi-homing.

Figure 159 **Aliasing**

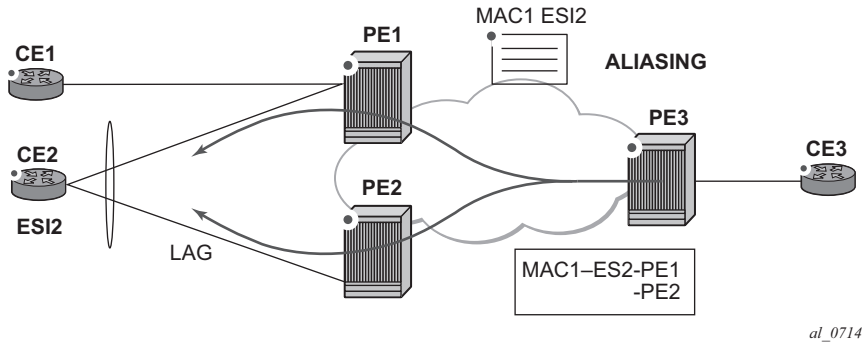

Because CE2 is multi-homed to PE1 and PE2 using an all-active ethernet-segment, 'aliasing' is the procedure by which PE3 can load-balance the known unicast traffic between PE1 and PE2, even if the destination MAC address was only advertised by PE1 as in the example. When PE3 installs MAC1 in the FDB, it will associate MAC1 not only with the advertising PE (PE1) but also with all the PEs advertising the same esi (ESI2) for the service. In this example, PE1 and PE2 advertise an AD per-EVI route for ESI2, therefore, the PE3 installs the two next-hops associated with MAC1.

Aliasing is enabled by configuring ECMP greater than 1 in the **bgp-evpn mpls** context.

All-Active Multi-Homing Service Model

The following shows an example PE1 configuration that provides all-active multi-homing to the CE2 shown in Figure 159.

```
*A:PE1>config>lag(1)# info
----------------------------------------------
  mode access
  encap-type dot1q
```

```
  port 1/1/2
  lacp active administrative-key 1 system-id 00:00:00:00:00:22
  no shutdown

*A:PE1>config>service>system>bgp-evpn# info
----------------------------------------------
  route-distinguisher 1.1.1.1:0
  ethernet-segment "ESI2" create
    esi 01:12:12:12:12:12:12:12:12:12
    multi-homing all-active
    service-carving
    lag 1
    no shutdown

*A:PE1>config>redundancy>evpn-multi-homing# info
----------------------------------------------
    boot-timer 120
    es-activation-timer 10

*A:PE1>config>service>vpls# info
----------------------------------------------
  description "evpn-mpls-service with all-active multihoming"
  bgp
  bgp-evpn
    evi 10
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  sap lag-1:1 create
  exit
```

In the same way, PE2 is configured as follows:

```
*A:PE1>config>lag(1)# info
----------------------------------------------
  mode access
  encap-type dot1q
  port 1/1/1
  lacp active administrative-key 1 system-id 00:00:00:00:00:22
  no shutdown

*A:PE1>config>service>system>bgp-evpn# info
----------------------------------------------
  route-distinguisher 1.1.1.1:0
  ethernet-segment "ESI12" create
    esi 01:12:12:12:12:12:12:12:12:12
    multi-homing all-active
    service-carving
    lag 1
    no shutdown

*A:PE1>config>redundancy>evpn-multi-homing# info
----------------------------------------------
    boot-timer 120
    es-activation-timer 10

*A:PE1>config>service>vpls# info
```

```
-----------------------------------------------
  description "evpn-mpls-service with all-active multihoming"
  bgp
    route-distinguisher 65001:60
    route-target target:65000:60
  bgp-evpn
    evi 10
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  sap lag-1:1 create
  exit
```

The preceding configuration will enable the all-active multi-homing procedures. The following must be considered:

- The **ethernet-segment** must be configured with a name and a 10-byte esi:
  - **config>service>system>bgp-evpn#ethernet-segment <es_name> create**
  - **config>service> system>bgp-evpn>ethernet-segment# esi <value>**
- When configuring the esi, the system enforces the 6 high-order octets after the type to be different from zero (so that the auto-derived route-target for the ES route is different from zero). Other than that, the entire esi value must be unique in the system.
- Only a LAG can be associated with the **ethernet-segment**. This LAG will be exclusively used for EVPN multi-homing. Other LAG ports in the system can be still used for MC-LAG and other services.
- When the LAG is configured on PE1 and PE2, the same **admin-key**, **system-priority**, and **system-id** must be configured on both PEs, so that CE2 responds as though it is connected to the same system.
- The same **ethernet-segment** may be used for EVPN-MPLS and PBB-EVPN services.

Note: The **source-bmac-lsb** attribute must be defined for PBB-EVPN (so that it will only be used in PBB-EVPN, and ignored by EVPN). Other than EVPN-MPLS and PBB-EVPN I-VPLS/Epipe services, no other Layer 2 services are allowed in the same **ethernet-segment** (regular VPLS or EVPN-VXLAN SAPs defined on the **ethernet-segment** will be kept operationally down).

- Only one sap per service can be part of the same **ethernet-segment**.

ES Discovery and DF Election Procedures

The ES (Ethernet Segment) discovery and DF election is implemented in three logical steps, as shown in Figure 160.

Figure 160 **ES Discovery and DF Election**

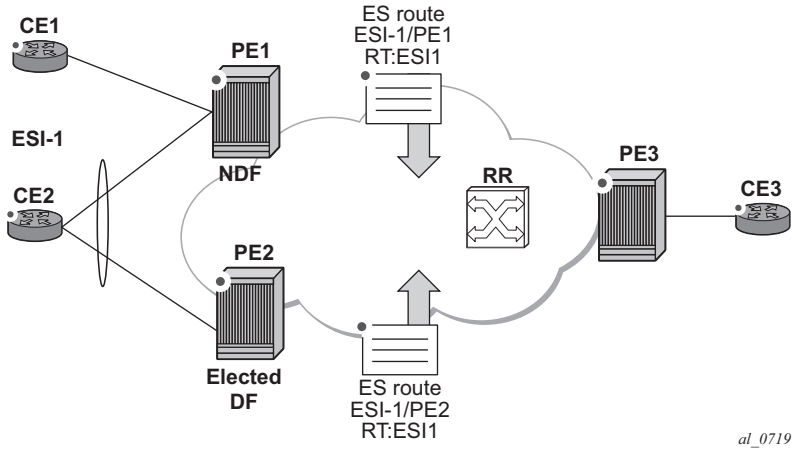

Step 1 - ES Advertisement and Discovery

Ethernet-segment ESI-1 is configured as per the previous section, with all the required parameters. When **ethernet-segment no shutdown** is executed, PE1 and PE2 will advertise an ES route for ESI-1. They will both include the route-target auto-derived from the MAC portion of the configured ESI. If the route-target address family is configured in the network, this will allow the RR to keep the dissemination of the ES routes under control.

In addition to the ES route, PE1 and PE2 will advertise AD per-ESI routes and AD per-EVI routes.

- AD per-ESI routes will announce the ethernet-segment capabilities, including the mode (single-active or all-active) as well as the ESI label for split-horizon.
- AD per-EVI routes are advertised so that PE3 knows what services (EVIs) are associated with the ESI. These routes are used by PE3 for its aliasing and backup procedures.

Step 2 - DF Election

When ES routes exchange between PE1 and PE2 is complete, both run the DF election for all the services in the **ethernet-segment**.

PE1 and PE2 elect a Designated Forwarder (DF) per <ESI, service>. The default DF election mechanism in 7750 SR, 7450 ESS, and 7950 XRS SR OS is **service-carving** (as per RFC 7432). The following applies when enabled on a specified PE:

- An ordered list of PE IPs where ESI-1 resides is built. The IPs are gotten from the Origin IP fields of all the ES routes received for ESI-1, as well as the local system address. The lowest IP will be considered ordinal '0' in the list.
- The local IP can only be considered a "candidate" after successful **ethernet-segment no shutdown** for a specified service.

Note: The remote PE IPs must be present in the local PE's RTM so that they can participate in the DF election.

- A PE will only consider a specified remote IP address as candidate for the DF election algorithm for a specified service if, as well as the ES route, the corresponding AD routes per-ESI and per-EVI for that PE have been received and properly activated.
- All the remote PEs receiving the AD per-ES routes (for example, PE3), will interpret that ESI-1 is all-active if all the PEs send their AD per-ES routes with the single-active bit = 0. Otherwise, if at least one PE sends an AD route per-ESI with the single-active flag set or the local ESI configuration is single-active, the ESI will behave as single-active.
- An **es-activation-timer** can be configured at the **redundancy>bgp-evpn-multi-homing>es-activation-timer** level or at the **service>system>bgp-evpn>eth-seg>es-activation-timer** level. This timer, which is 3 seconds by default, delays the transition from non-DF to DF for a specified service, after the DF election has run.
  - This use of the **es-activation-timer** is different from zero and minimizes the risks of loops and packet duplication due to "transient" multiple DFs.
  - The same **es-activation-timer** should be configured in all the PEs that are part of the same ESI. It is up to the user to configure either a long timer to minimize the risks of loops/duplication or even **es-activation-timer=0** to speed up the convergence for non-DF to DF transitions. When the user configures a specific value, the value configured at ES level supersedes the configured global value.
- The DF election is triggered by the following events:
  - **config>service>system>bgp-evpn>eth-seg# no shutdown** triggers the DF election for all the services in the ESI.
  - Reception of a new update/withdrawal of an ES route (containing an ESI configured locally) triggers the DF election for all the services in the ESI.

- Reception of a new update/withdrawal of an AD per-ES route (containing an ESI configured locally) triggers the DF election for all the services associated with the list of route-targets received along with the route.
- Reception of a new update of an AD per-ES route with a change in the ESI-label extended community (single-active bit or MPLS label) triggers the DF election for all the services associated with the list of route-targets received along with the route.
- Reception of a new update/withdrawal of an AD route per-EVI (containing an ESI configured locally) triggers the DF election for that service.

- When the PE boots up, the boot-timer will allow the necessary time for the control plane protocols to come up before bringing up the ethernet-segment and running the DF algorithm. The boot-timer is configured at system level - config>redundancy>bgp-evpn-multi-homing# boot-timer - and should use a value long enough to allow the IOMs and BGP sessions to come up before exchanging ES routes and running the DF election for each EVI/ISID.
  - The system will not advertise ES routes until the boot timer expires. This will guarantee that the peer ES PEs don't run the DF election either until the PE is ready to become the DF if it needs to.
  - The following show command displays the configured boot-timer as well as the remaining timer if the system is still in boot-stage.

```
A:PE1# show redundancy bgp-evpn-multi-homing
===============================================================================
Redundancy BGP EVPN Multi-homing Information
===============================================================================
Boot-Timer              : 10 secs
Boot-Timer Remaining    : 0 secs
ES Activation Timer     : 3 secs
===============================================================================
```

- When **service-carving mode auto** is configured (default mode), the DF election algorithm will run the function [V(evi) mod N(peers) = i(ordinal)] to identify the DF for a specified service and ESI, as described in the following example:
  - As shown in Figure 160, PE1 and PE2 are configured with ESI-1. Given that V(10) mod N(2) = 0, PE1 will be elected DF for VPLS-10 (because its IP address is lower than PE2's and it is the first PE in the candidate list).

Note: The algorithm takes the configured **evi** in the service as opposed to the service-id itself. The **evi** for a service must match in all the PEs that are part of the ESI. This guarantees that the election algorithm is consistent across all the PEs of the ESI. The **evi** must be always configured in a service with saps/sdp-bindings that are created in an ES.

- A **manual** service-carving option is allowed so that the user can manually configure for which evi identifiers the PE is primary: **service-carving mode manual / manual service <evi> to <evi>.**

– The system will be the PE forwarding/multicasting traffic for the **evi** identifiers included in the configuration. The PE will be secondary (non-DF) for the non-specified **evis**.
– If a range is configured but the service-carving is not mode manual, then the range has no effect.
– Only two PEs are supported when service-carving mode manual is configured. If a third PE is configured with service-carving mode manual for an ESI, the two non-primary PEs will remain non-DF regardless of the primary status.
– For example, as shown in Figure 160: if PE1 is configured with service-carving manual evi 1 to 100 and PE2 with service-carving manual evi 101 to 200, then PE1 will be the primary PE for service VPLS 10 and PE2 the secondary PE.

- When service-carving is disabled, the lowest originator IP will win the election for a specified service and ESI:

config>service>system>bgp-evpn>ethernet-segment> mode off

The following show command displays the **ethernet-segment** configuration and DF status for all the EVIs and ISIDs (if PBB-EVPN is enabled) configured in the **ethernet-segment**.

```
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-1" all
===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-1
Admin State             : Up                  Oper State          : Up
ESI                     : 01:00:00:00:00:71:00:00:00:01
Multi-homing            : allActive           Oper Multi-homing   : allActive
Source BMAC LSB         : 71-71
ES BMac Tbl Size        : 8                   ES BMac Entries     : 1
Lag Id                  : 1
ES Activation Timer     : 0 secs
Exp/Imp Route-Target    : target:00:00:00:00:71:00

Svc Carving             : auto
ES SHG Label            : 262142
===============================================================================
===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   no
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
```

```
-------------------------------------------------------------------------------
1                                   192.0.2.69
1                                   192.0.2.72
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
ISID Information
===============================================================================
ISID                SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
20001               20001               0                   no
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
ISID                                DF Address
-------------------------------------------------------------------------------
20001                               192.0.2.69
20001                               192.0.2.72
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
BMAC Information
===============================================================================
SvcId                               BMacAddress
-------------------------------------------------------------------------------
20000                               00:00:00:00:71:71
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

Step 3 - DF and Non-DF Service Behavior

Based on the result of the DF election or the manual service-carving, the control plane on the non-DF (PE1) will instruct the data path to remove the LAG SAP (associated with the ESI) from the default flooding list for BM traffic (unknown unicast traffic may still be sent if the EVI label is a unicast label and the source MAC address is not associated to the ESI). On PE1 and PE2, both LAG SAPs will learn the same MAC address (coming from the CE). For instance, in the following show commands, 00:ca:ca:ba:ce:03 is learned on both PE1 and PE2 access LAG (on ESI-1). However, PE1 learns the MAC as 'Learned' whereas PE2 learns it as 'Evpn'. This is due to the CE2 hashing the traffic for that source MAC to PE1. PE2 learns the MAC through EVPN but it associates the MAC to the ESI SAP, because the MAC belongs to the ESI.

```
*A:PE1# show service id 1 fdb detail
```

```
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:ca:ca:ba:ce:03 sap:lag-1:1              L/0      06/11/15 00:14:47
1         00:ca:fe:ca:fe:70 eMpls:                   EvpnS    06/11/15 00:09:06
                            192.0.2.70:262140
1         00:ca:fe:ca:fe:72 eMpls:                   EvpnS    06/11/15 00:09:39
                            192.0.2.72:262141
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================

*A:PE2# show service id 1 fdb detail
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:ca:ca:ba:ce:03 sap:lag-1:1              Evpn     06/11/15 00:14:47
1         00:ca:fe:ca:fe:69 eMpls:                   EvpnS    06/11/15 00:09:40
                            192.0.2.69:262141
1         00:ca:fe:ca:fe:70 eMpls:                   EvpnS    06/11/15 00:09:40
                            192.0.2.70:262140
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

When PE1 (non-DF) and PE2 (DF) exchange BUM packets for **evi 1**, all those packets will be sent including the ESI label at the bottom of the stack (in both directions). The ESI label advertised by each PE for ESI-1 can be displayed by the following command:

```
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-1"
===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-1
Admin State             : Up                 Oper State         : Up
ESI                     : 01:00:00:00:00:71:00:00:00:01
Multi-homing            : allActive          Oper Multi-homing  : allActive
Source BMAC LSB         : 71-71
ES BMac Tbl Size        : 8                  ES BMac Entries    : 1
Lag Id                  : 1
ES Activation Timer     : 0 secs
Exp/Imp Route-Target    : target:00:00:00:00:71:00

Svc Carving             : auto
ES SHG Label            : 262142
===============================================================================
```

```
*A:PE2# show service system bgp-evpn ethernet-segment name "ESI-1"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-1
Admin State             : Up                  Oper State          : Up
ESI                     : 01:00:00:00:00:71:00:00:00:01
Multi-homing            : allActive           Oper Multi-homing   : allActive
Source BMAC LSB         : 71-71
ES BMac Tbl Size        : 8                   ES BMac Entries     : 0
Lag Id                  : 1
ES Activation Timer     : 20 secs
Exp/Imp Route-Target    : target:00:00:00:00:71:00

Svc Carving             : auto
ES SHG Label            : 262142
===============================================================================
```

Aliasing

Following the example in Figure 160, if the service configuration on PE3 has ECMP > 1, PE3 will add PE1 and PE2 to the list of next-hops for ESI-1. As soon as PE3 receives a MAC for ESI-1, it will start load-balancing between PE1 and PE2 the flows to the remote ESI CE. The following command shows the FDB in PE3.

Note: mac 00:ca:ca:ba:ce:03 is associated with the ethernet-segment eES:01:00:00:00:00:71:00:00:00:01 (esi configured on PE1 and PE2 for ESI-1).

```
*A:PE3# show service id 1 fdb detail
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:ca:ca:ba:ce:03 eES:                     Evpn     06/11/15 00:14:47
                            01:00:00:00:00:71:00:00:00:01
1         00:ca:fe:ca:fe:69 eMpls:                   EvpnS    06/11/15 00:09:18
                            192.0.2.69:262141
1         00:ca:fe:ca:fe:70 eMpls:                   EvpnS    06/11/15 00:09:18
                            192.0.2.70:262140
1         00:ca:fe:ca:fe:72 eMpls:                   EvpnS    06/11/15 00:09:39
                            192.0.2.72:262141
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

The following command shows all the EVPN-MPLS destination bindings on PE3, including the ES destination bindings.

The ethernet-segment eES:01:00:00:00:00:71:00:00:00:01 is resolved to PE1 and PE2 addresses:

```
*A:PE3# show service id 1 evpn-mpls
===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                 Transport
-------------------------------------------------------------------------------
192.0.2.69      262140        0           Yes             06/10/2015 14:33:30
                ldp
192.0.2.69      262141        1           No              06/10/2015 14:33:30
                ldp
192.0.2.70      262139        0           Yes             06/10/2015 14:33:30
                ldp
192.0.2.70      262140        1           No              06/10/2015 14:33:30
                ldp
192.0.2.72      262140        0           Yes             06/10/2015 14:33:30
                ldp
192.0.2.72      262141        1           No              06/10/2015 14:33:30
                ldp
192.0.2.73      262139        0           Yes             06/10/2015 14:33:30
                ldp
192.0.2.254     262142        0           Yes             06/10/2015 14:33:30
                bgp
-------------------------------------------------------------------------------
Number of entries : 8
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       TEP Address     Egr Label   Last Change
                                                 Transport
-------------------------------------------------------------------------------
01:00:00:00:00:71:00:00:00:01 192.0.2.69      262141      06/10/2015 14:33:30
                                              ldp
01:00:00:00:00:71:00:00:00:01 192.0.2.72      262141      06/10/2015 14:33:30
                                              ldp
01:74:13:00:74:13:00:00:74:13 192.0.2.73      262140      06/10/2015 14:33:30
                                              ldp
-------------------------------------------------------------------------------
Number of entries : 3
-------------------------------------------------------------------------------
===============================================================================
```

PE3 will perform aliasing for all the MACs associated with that ESI. This is possible because PE1 is configured with ecmp parameter >1:

```
*A:PE3>config>service>vpls# info
----------------------------------------------
            bgp
            exit
```

```
bgp-evpn
    evi 1
    vxlan
        shutdown
    exit
    mpls
        ecmp 4
        auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
exit
proxy-arp
    shutdown
exit
stp
    shutdown
exit
sap 1/1/1:2 create
exit
no shutdown
```

Network Failures and Convergence for All-Active Multi-Homing

Figure 161 shows the behavior on the remote PEs (PE3) when there is an **ethernet-segment** failure.

Figure 161 **All-Active Multi-Homing ES Failure**

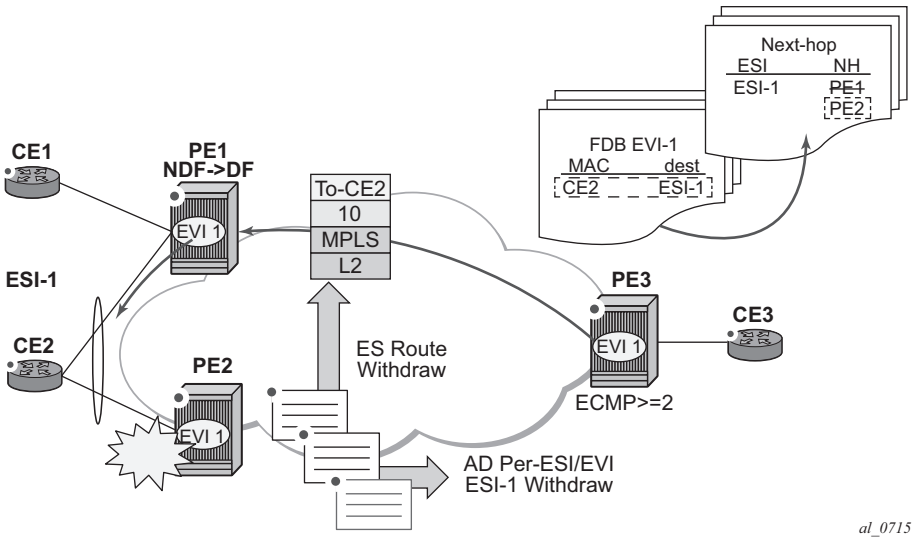

The unicast traffic behavior on PE3 is as follows:

1. PE3 can only forward MAC DA = CE2 to both PE1 and PE2 when the MAC advertisement route from PE1 (or PE2) and the set of Ethernet AD per-ES routes and Ethernet AD per-EVI routes from PE1 and PE2 are active at PE3.
2. If there was a failure between CE2 and PE2, PE2 would withdraw its set of Ethernet AD and ES routes, then PE3 would forward traffic destined to CE2 to PE1 only. PE3 does not need to wait for the withdrawal of the individual MAC.
3. The same behavior would be followed if the failure had been at PE1.
4. If after (2), PE2 withdraws its MAC advertisement route, then PE3 treats traffic to MAC DA = CE2 as unknown unicast, unless the MAC had been previously advertised by PE1.

For BUM traffic, the following events would trigger a DF election on a PE and only the DF would forward BUM traffic after the **esi-activation-timer** expiration (if there was a transition from non-DF to DF).

1. Reception of ES route update (local ES shutdown/no shutdown or remote route)
2. New AD-ES route update/withdraw
3. New AD-EVI route update/withdraw
4. Local ES port/SAP/service shutdown
5. Service carving range change (affecting the evi)
6. Multi-homing mode change (single/all active to all/single-active)

Logical Failures on Ethernet Segments and Black-Holes

Be aware of the effects triggered by certain 'failure scenarios'; some of these scenarios are shown in Figure 162:

Figure 162 **Black-hole Caused by SAP/SVC Shutdown**

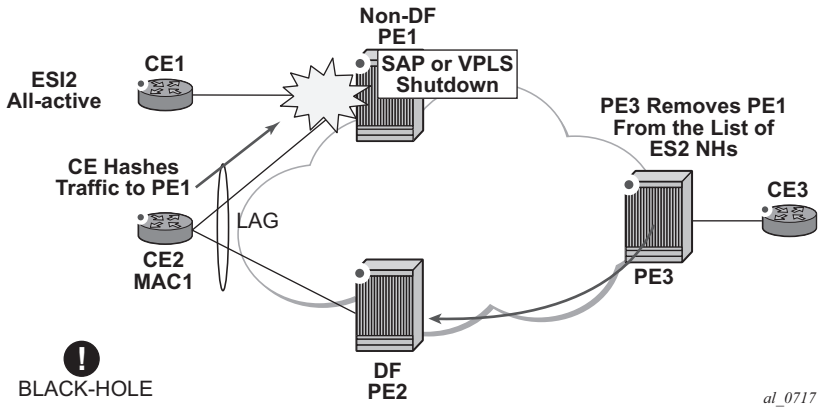

If an individual VPLS service is **shutdown** in PE1 (the example is also valid for PE2), the corresponding LAG SAP will go *operationally down*. This event will trigger the withdrawal of the AD per-EVI route for that particular SAP. PE3 will remove PE1 of its list of aliased next-hops and PE2 will take over as DF (if it was not the DF already). However, this will not prevent the network from black-holing the traffic that CE2 'hashes' to the link to PE1. Traffic sent from CE2 to PE2 or traffic from the rest of the CEs to CE2 will be unaffected, so this situation is not easily detected on the CE.

The same result occurs if the ES SAP is administratively **shutdown** instead of the service.

Note: When **bgp-evpn mpls shutdown** is executed, the sap associated with the ES will be brought operationally down (**StandbyforMHprotocol**) and so will the entire service if there are no other saps or sdp-bindings in the service. However, if there are other saps/sdp-bindings, the service will remain operationally up.

Transient Issues Due to MAC Route Delays

Some situations may cause potential transient issues to occur. These are shown in Figure 163 and explained below.

Figure 163 **Transient Issues Caused by "slow" MAC Learning**

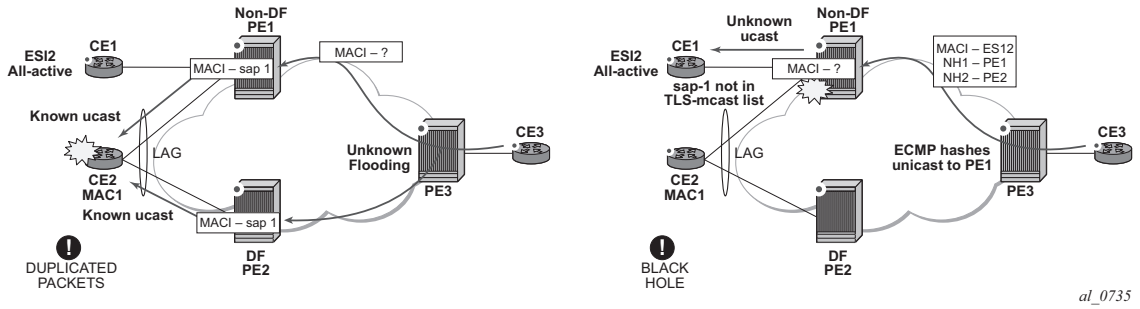

Transient packet duplication caused by delay in PE3 to learn MAC1:

This scenario is illustrated by the diagram on the left in Figure 163. In an all-active multi-homing scenario, if a specified MAC address is not yet learned in a remote PE, but is known in the two PEs of the ES, for example, PE1 and PE2, the latter PEs might send duplicated packets to the CE.

In an all-active multi-homing scenario, if a specified MAC address (for example, MAC1), is not learned yet in a remote PE (for example, PE3), but it is known in the two PEs of the ES (for example, PE1 and PE2), the latter PEs might send duplicated packets to the CE.

This issue is solved by the use of **ingress-replication-bum-label** in PE1 and PE2. If configured, PE1/PE2 will know that the received packet is an unknown unicast packet, therefore, the NDF (PE1) will not send the packets to the CE and there will not be duplication.

Note: Even without the **ingress-replication-bum-label,** this is only a transient situation that would be solved as soon as MAC1 is learned in PE3.

Transient black-hole caused by delay in PE1 to learn MAC1:

This case is illustrated by the diagram on the right in Figure 163. In an all-active multi-homing scenario, MAC1 is known in PE3 and aliasing is applied to MAC1. However, MAC1 is not known yet in PE1, the NDF for the ES. If PE3 hashing picks up PE1 as the destination of the aliased MAC1, the packets will be black-holed. This case is solved on the NDF by not blocking unknown unicast traffic that arrives with a unicast label (possible if PE1 and PE2 are configured using **ingress-replication-bum-label**).

As soon as PE1 learns MAC1, the black-hole will be resolved even if **ingress-replication-bum-label** is not used.

5.3.2.3.2 EVPN Single-Active Multi-Homing

The 7750 SR, 7450 ESS, and 7950 XRS SR OS supports single-active multi-homing on access LAG SAPs, regular SAPs, and spoke-SDPs for a specified VPLS service. For LAG SAPs, the CE will be configured with a different LAG to each PE in the ethernet-segment (as opposed to a single LAG in an all-active multi-homing).

The following SR OS procedures support EVPN single-active multi-homing for a specified ethernet-segment:

- DF (Designated Forwarder) election

  As in all-active multi-homing, DF election in single-active multi-homing determines the forwarding for BUM traffic from the EVPN network to the ethernet-segment CE. Also, in single-active multi-homing, DF election also determines the forwarding of any traffic (unicast/BUM) and in any direction (to/from the CE).
- Backup PE

In single-active multi-homing, the remote PEs do not perform aliasing to the PEs in the ethernet-segment. The remote PEs identify the DF based on the MAC routes and send the unicast flows for the ethernet-segment to the PE in the DF and program a backup PE as an alternative next-hop for the remote ESI in case of failure.

This RFC 7432 procedure is known as 'Backup PE' and is shown in Figure 164 for PE3.

Figure 164 **Backup PE**

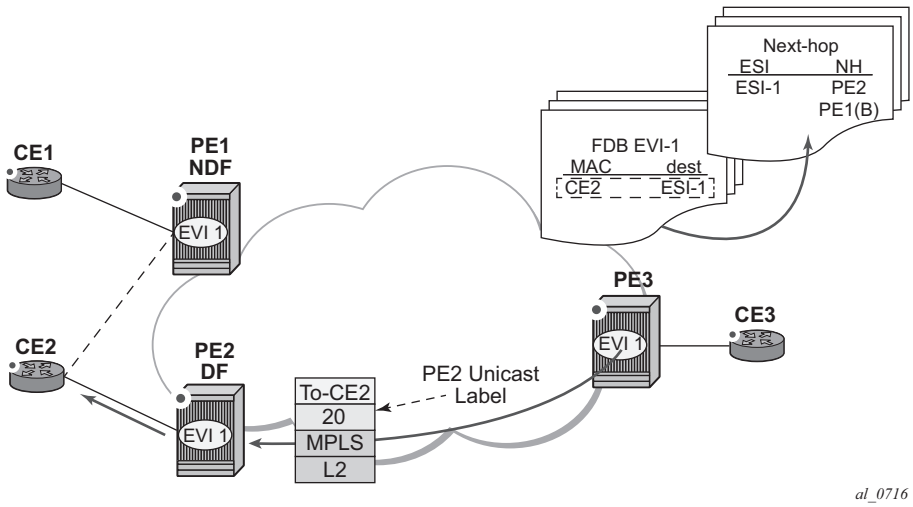

Single-Active Multi-Homing Service Model

The following shows an example of PE1 configuration that provides single-active multi-homing to CE2, as shown in Figure 164.

```
*A:PE1>config>service>system>bgp-evpn# info
----------------------------------------------
  route-distinguisher 1.1.1.1:0
  ethernet-segment "ESI2" create
    esi 01:12:12:12:12:12:12:12:12:12
    multi-homing single-active
    service-carving
    sdp 1
    no shutdown

*A:PE1>config>redundancy>evpn-multi-homing# info
----------------------------------------------
    boot-timer 120
    es-activation-timer 10

*A:PE1>config>service>vpls# info
----------------------------------------------
  description "evpn-mpls-service with single-active multihoming"
```

```
  bgp
  bgp-evpn
    evi 10
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  spoke-sdp 1:1 create
  exit
```

The PE2 example configuration for this scenario is as follows:

```
*A:PE1>config>service>system>bgp-evpn# info
----------------------------------------------
  route-distinguisher 1.1.1.1:0
  ethernet-segment "ESI2" create
    esi 01:12:12:12:12:12:12:12:12
    multi-homing single-active
    service-carving
    sdp 2
    no shutdown

*A:PE1>config>redundancy>evpn-multi-homing# info
----------------------------------------------
    boot-timer 120
    es-activation-timer 10

*A:PE1>config>service>vpls# info
----------------------------------------------
  description "evpn-mpls-service with single-active multihoming"
  bgp
  bgp-evpn
    evi 10
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  spoke-sdp 2:1 create
  exit
```

In single-active multi-homing, the non-DF PEs for a specified ESI will block unicast and BUM traffic in both directions (upstream and downstream) on the object associated with the ESI. Other than that, single-active multi-homing is similar to all-active multi-homing with the following differences:

- The **ethernet-segment** will be configured for single-active: **service>system>bgp-evpn>ethernet-segment>multi-homing single-active**.
- The advertisement of the ESI-label in a per-ESI AD route is optional for **single-activ**e ethernet-segments. The user can control the no advertisement of the ESI label by using the following command: **service>system>bgp-evpn>ethernet-segment>multi-homing single-active no-esi-label**. By default, the ESI label is used for single-active ESs too.

- For single-active multi-homing, the ethernet-segment can be associated with a **port** and **sdp**, as well as a **lag-id**, as shown in Figure 164, where:
  - **port** would be used for single-active sap redundancy without the need for lag.
  - **sdp** would be used for single-active spoke-sdp redundancy.
  - **lag** would be used for single-active LAG redundancy

Note: In this case, key, system-id, and system-priority must be different on the PEs that are part of the ethernet-segment).

- For single-active multi-homing, when the PE is non-DF for the service, the saps/spoke-sdps on the ethernet-segment will be down and show **StandByForMHProtocol** as the reason.
- From a service perspective, single-active multi-homing can provide redundancy to CEs (MHD, Multi-Homed Devices) or networks (MHN, Multi-Homed Networks) with the following setup:
  - **LAG with or without LACP**

    In this case, the multi-homed ports on the CE will be part of the different LAGs (a LAG per multi-homed PE will be used in the CE). The non-DF PE for each service can signal that the sap is operationally down if eth-cfm fault-propagation-enable {use-if-tlv|suspend-ccm} is configured.
  - **Regular Ethernet 802.1q/ad ports**

    In this case, the multi-homed ports on the CE/network will not be part of any LAG. Eth-cfm can also be used for non-DF indication to the multi-homed device/network.
  - **Active-standby PWs**

    In this case, the multi-homed CE/network is connected to the PEs through an MPLS network and an active/standby spoke-sdp per service. The non-DF PE for each service will make use of the LDP PW status bits to signal that the spoke-sdp is operationally down on the PE side.

ES and DF Election Procedures

In all-active multi-homing, the non-DF keeps the SAP up, although it removes it from the default flooding list. In the single-active multi-homing implementation the non-DF will bring the SAP or SDP-binding operationally down. Refer to the ES Discovery and DF Election Procedures for more information.

The following **show** commands display the status of the single-active ESI-7413 in the non-DF. The associated spoke-SDP is operationally down and it signals PW Status standby to the multi-homed CE:

```
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-7413"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-7413
Admin State             : Up                 Oper State         : Up
ESI                     : 01:74:13:00:74:13:00:00:74:13
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
Source BMAC LSB         : <none>
Sdp Id                  : 4
ES Activation Timer     : 0 secs
Exp/Imp Route-Target    : target:74:13:00:74:13:00

Svc Carving             : auto
ES SHG Label            : 262141
===============================================================================
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-7413" evi 1
===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1             1             0                   no  06/11/2015 20:05:32
===============================================================================
===============================================================================
DF Candidates                           Time Added
-------------------------------------------------------------------------------
192.0.2.70                              06/11/2015 20:05:20
192.0.2.73                              06/11/2015 20:05:32
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
*A:PE1# show service id 1 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
Name              : (Not Specified)
Description       : (Not Specified)

<snip>
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1:1                              q-tag        9000    9000    Up   Up
sdp:4:13 S(192.0.2.74)                   Spok         0       8978    Up   Down
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

```
*A:PE1# show service id 1 all | match Pw
Local Pw Bits      : pwFwdingStandby
Peer Pw Bits       : None

*A:PE1# show service id 1 all | match Flag
Flags              : StandbyForMHProtocol
Flags              : None
```

Backup PE Function

A remote PE (PE3 in Figure 164) will import the AD routes per ESI, where the single-active flag is set. PE3 will interpret that the ethernet-segment is single-active if at least one PE sends an AD route per-ESI with the single-active flag set. MACs for a specified service and ESI will be learned from a single PE, that is, the DF for that <ESI, EVI>.

The remote PE will install a single EVPN-MPLS destination (TEP, label) for a received MAC address and a backup next-hop to the PE for which the AD routes per-ESI and per-EVI are received. For instance, in the following command, 00:ca:ca:ba:ca:06 is associated with the remote **ethernet-segment eES 01:74:13:00:74:13:00:00:74:13**. That eES is resolved to PE(192.0.2.73), which is the DF on the ES.

```
*A:PE3# show service id 1 fdb detail
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:ca:ca:ba:ca:02 sap:1/1/1:2              L/0      06/12/15 00:33:39
1         00:ca:ca:ba:ca:06 eES:                     Evpn     06/12/15 00:33:39
                            01:74:13:00:74:13:00:00:74:13
1         00:ca:fe:ca:fe:69 eMpls:                   EvpnS    06/11/15 21:53:47
                            192.0.2.69:262118
1         00:ca:fe:ca:fe:70 eMpls:                   EvpnS    06/11/15 19:59:57
                            192.0.2.70:262140
1         00:ca:fe:ca:fe:72 eMpls:                   EvpnS    06/11/15 19:59:57
                            192.0.2.72:262141
-------------------------------------------------------------------------------
No. of MAC Entries: 5
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================

*A:PE3# show service id 1 evpn-mpls
===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                 Transport
-------------------------------------------------------------------------------
```

```
192.0.2.69      262118        1           Yes             06/11/2015 19:59:03
                ldp
192.0.2.70      262139        0           Yes             06/11/2015 19:59:03
                ldp
192.0.2.70      262140        1           No              06/11/2015 19:59:03
                ldp
192.0.2.72      262140        0           Yes             06/11/2015 19:59:03
                ldp
192.0.2.72      262141        1           No              06/11/2015 19:59:03
                ldp
192.0.2.73      262139        0           Yes             06/11/2015 19:59:03
                ldp
192.0.2.254     262142        0           Yes             06/11/2015 19:59:03
                bgp
-------------------------------------------------------------------------------
Number of entries : 7
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       TEP Address     Egr Label   Last Change
                                                 Transport
-------------------------------------------------------------------------------
01:74:13:00:74:13:00:00:74:13 192.0.2.73      262140      06/11/2015 19:59:03
                                              ldp
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

If PE3 sees only two single-active PEs in the same ESI, the second PE will be the backup PE. Upon receiving an AD per-ES/per-EVI route withdrawal for the ESI from the primary PE, the PE3 will start sending the unicast traffic to the backup PE immediately.

If PE3 receives AD routes for the same ESI and EVI from more than two PEs, the PE will not install any backup route in the data path. Upon receiving an AD per-ES/per-EVI route withdrawal for the ESI, it will flush the MACs associated with the ESI.

Network Failures and Convergence for Single-Active Multi-Homing

Figure 165 shows the remote PE (PE3) behavior when there is an ethernet-segment failure.

Figure 165 **Single-Active Multi-Homing ES Failure**

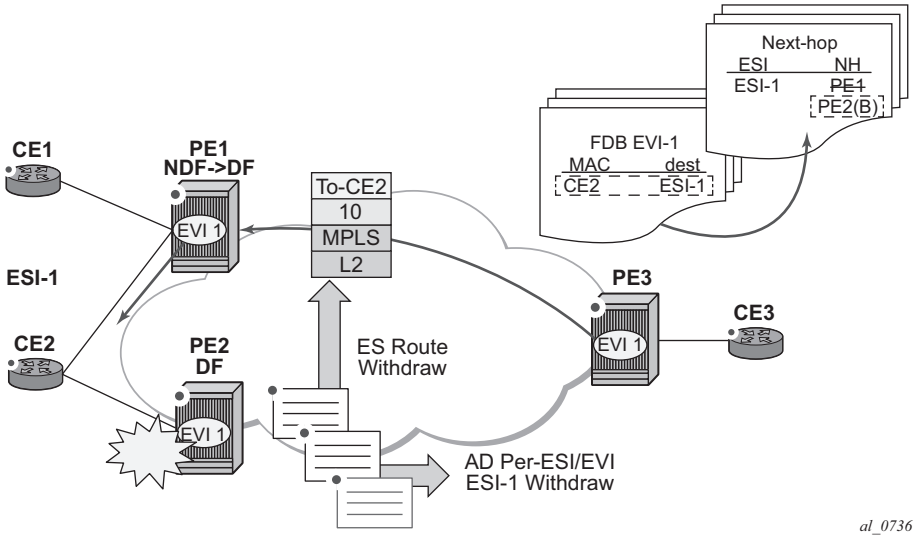

The PE3 behavior for unicast traffic is as follows:

1. PE3 forwards MAC DA = CE2 to PE2 when the MAC Advertisement Route came from PE2 and the set of Ethernet AD per-ES routes and Ethernet AD per-EVI routes from PE1 and PE2 are active at PE3.
2. If there was a failure between CE2 and PE2, PE2 would withdraw its set of Ethernet AD and ES routes, then PE3 would immediately forward the traffic destined for CE2 to PE1 only (the backup PE). PE3 does not need to wait for the withdrawal of the individual MAC.
3. After the (2) PE2 withdraws its MAC advertisement route, PE3 will treat traffic to MAC DA = CE2 as unknown unicast, unless the MAC has been previously advertised by PE1.

Also, a DF election on PE1 is triggered. In general, a DF election is triggered by the same events as for all-active multi-homing. In this case, the DF will forward traffic to CE2 when the **esi-activation-timer** expiration occurs (the timer kicks in when there is a transition from non-DF to DF).

5.3.3 P2MP mLDP tunnels for BUM traffic in EVPN-MPLS Services

P2MP mLDP tunnels for BUM traffic in EVPN-MPLS services are supported and enabled through the use of the provider-tunnel context. If EVPN-MPLS takes ownership over the provider-tunnel, **bgp-ad** is still supported in the service but it will not generate BGP updates, including the PMSI Tunnel Attribute. The following CLI example shows an EVPN-MPLS service that uses P2MP mLDP LSPs for BUM traffic.

```
*A:PE-1>config>service>vpls(vpls or b-vpls)# info
----------------------------------------------
  description "evpn-mpls-service with p2mp mLDP"
  bgp-evpn
    evi 10
    no ingress-repl-inc-mcast-advertisement
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  provider-tunnel
    inclusive
      owner bgp-evpn-mpls
      root-and-leaf
      mldp
      no shutdown
      exit
    exit
  sap 1/1/1:1 create
  exit
  spoke-sdp 1:1 create
 exit
```

When **provider-tunnel inclusive** is used in EVPN-MPLS services, the following commands can be used in the same way as for BGP-AD or BGP-VPLS services:

- **data-delay-interval**
- **root-and-leaf**
- **mldp**
- **shutdown**

The following commands are used by **provider-tunnel** in BGP-EVPN MPLS services:

- **[no] ingress-repl-inc-mcast-advertisement**

  This command allows you to control the advertisement of IMET-IR and IMET-P2MP-IR routes for the service. See BGP-EVPN Control Plane for MPLS Tunnels for a description of the IMET routes. The following considerations apply:

– If configured as **no ingress-repl-inc-mcast-advertisement**, the system will not send the IMET-IR or IMET-P2MP-IR routes, regardless of the service being enabled for bgp-evpn mpls or bgp-evpn vxlan.

– If configured as **ingress-repl-inc-mcast-advertisement** and the PE is **root-and-leaf**, the system will send an IMET-P2MP-IR route.

– If configured as **ingress-repl-inc-mcast-advertisement** and the PE is **no root-and-leaf**, the system will send an IMET-IR route.

– Default value is **ingress-repl-inc-mcast-advertisement**.

- **[no] owner {bgp-ad|bgp-vpls|bgp-evpn-mpls}**

  The owner of the provider tunnel must be configured. The default value is **no owner**. The following considerations apply:

  – Only one of the protocols will support a provider tunnel in the service and it must be explicitly configured.

  – **bgp-vpls** and **bgp-evpn** are mutually-exclusive.

  – While **bgp-ad** and **bgp-evpn** can coexist in the same service, only **bgp-evpn** can be the provider-tunnel owner in such cases.

Figure 166 shows the use of P2MP mLDP tunnels in an EVI with a root node and a few leaf-only nodes.

Figure 166 **EVPN Services with p2mp mLDP—Control Plane**

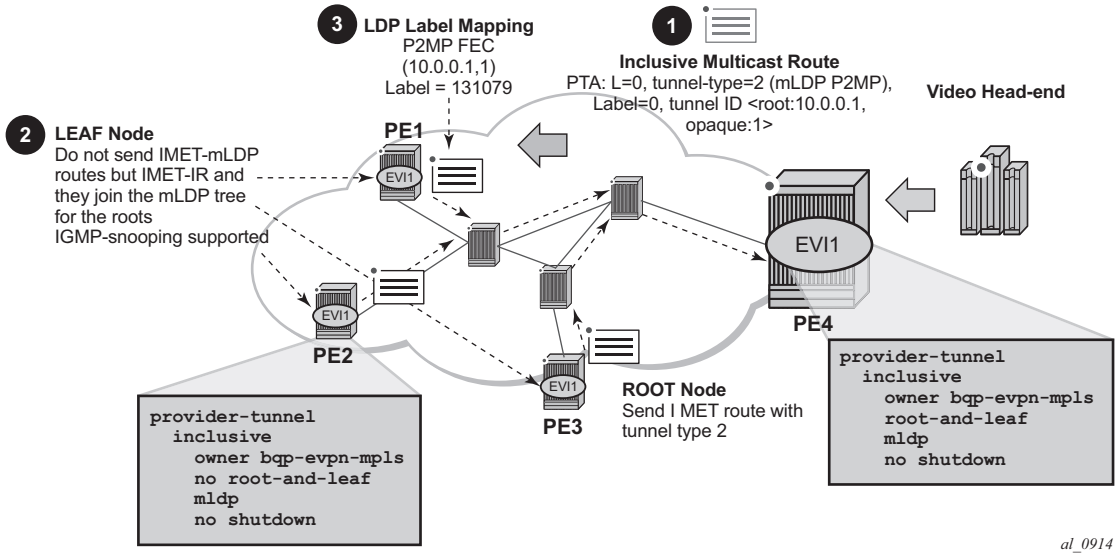

Consider the use-case of a root-and-leaf PE4 where the other nodes are configured as leaf-only nodes (**no root-and-leaf**). This scenario is handled as follows:

1. If **ingress-repl-inc-mcast-advertisement** is configured, then as soon as the **bgp-evpn mpls** option is enabled, the PE4 sends an IMET-P2MP route (tunnel type mLDP), or optionally, an IMET-P2MP-IR route (tunnel type composite). IMET-P2MP-IR routes allow leaf-only nodes to create EVPN-MPLS multicast destinations and send BUM traffic to the root.
2. If **ingress-repl-inc-mcast-advertisement** is configured, PE1/2/3 will not send IMET-P2MP routes; only IMET-IR routes will be sent.
   - The **root-and-leaf** node will import the IMET-IR routes from the leaf nodes but it will only send BUM traffic to the P2MP tunnel as long as it is active.
   - If the P2MP tunnel goes operationally down, the **root-and-leaf** node will start sending BUM traffic to the evpn-mpls multicast destinations
3. When PE1/2/3 receive and import the IMET-P2MP or IMET-P2MP-IR from PE4, they will join the mLDP P2MP tree signaled by PE4. They will issue an LDP label-mapping message including the corresponding P2MP FEC.

As described in IETF Draft *draft-ietf-bess-evpn-etree*, mLDP and Ingress Replication (IR) can work in the same network for the same service; that is, EVI1 can have some nodes using mLDP (for example, PE1) and others using IR (for example, PE2). For scaling, this is significantly important in services that consist of a pair of root nodes sending BUM in P2MP tunnels and hundreds of leaf-nodes that only need to send BUM traffic to the roots. By using IMET-P2MP-IR routes from the roots, the operator makes sure the leaf-only nodes can send BUM traffic to the root nodes without the need to set up P2MP tunnels from the leaf nodes.

When both static and dynamic P2MP mLDP tunnels are used on the same router, Nokia recommends that the static tunnels use a tunnel ID lower than 8193. If a tunnel ID is statically configured with a value equal to or greater than 8193, BGP-EVPN may attempt to use the same tunnel ID for services with **enabled provider-tunnel**, and fail to set up an mLDP tunnel.

Inter-AS option C or seamless-MPLS models for non-segmented mLDP trees are supported with EVPN for BUM traffic. The leaf PE that joins an mLDP EVPN root PE supports Recursive and Basic Opaque FEC elements (types 7 and 1, respectively). Therefore, packet forwarding is handled as follows:

- The ABR or ASBR may leak the root IP address into the leaf PE IGP, which allows the leaf PE to issue a Basic opaque FEC to join the root.
- The ABR or ASBR may distribute the root IP using BGP label-ipv4, which results in the leaf PE issuing a Recursive opaque FEC to join the root.

For more information about mLDP opaque FECs, refer to the *7450 ESS, 7750 SR, and 7950 XRS Layer 3 Services Guide: Internet Enhanced Services and Virtual Private Routed Network Services* and the *7450 ESS, 7750 SR, and 7950 XRS MPLS Guide*.

All-active multihoming and single-active with an ESI label multihoming are supported in EVPN-MPLS services together with P2MP mLDP tunnels. Both use an upstream-allocated ESI label, as described in *RFC 7432* section 8.3.1.2, which is popped at the leaf PEs, resulting in the requirement that, in addition to the root PE, all EVPN-MPLS P2MP leaf PEs must support this capability (including the PEs not connected to the multihoming Ethernet segment).

5.3.4 EVPN-VPWS for MPLS Tunnels

This section contains information about EVPN-VPWS for MPLS tunnels.

5.3.4.1 BGP-EVPN Control Plane for EVPN-VPWS

EVPN-VPWS uses route-type 1 and route-type 4; it does not use route-types 2, 3 or 5. Figure 167 shows the encoding of the required extensions for the Ethernet A-D per-EVI routes. The encoding follows the guidelines described in *draft-ief-bess-evpn-vpws*.

Figure 167 **EVPN-VPWS BGP Extensions**

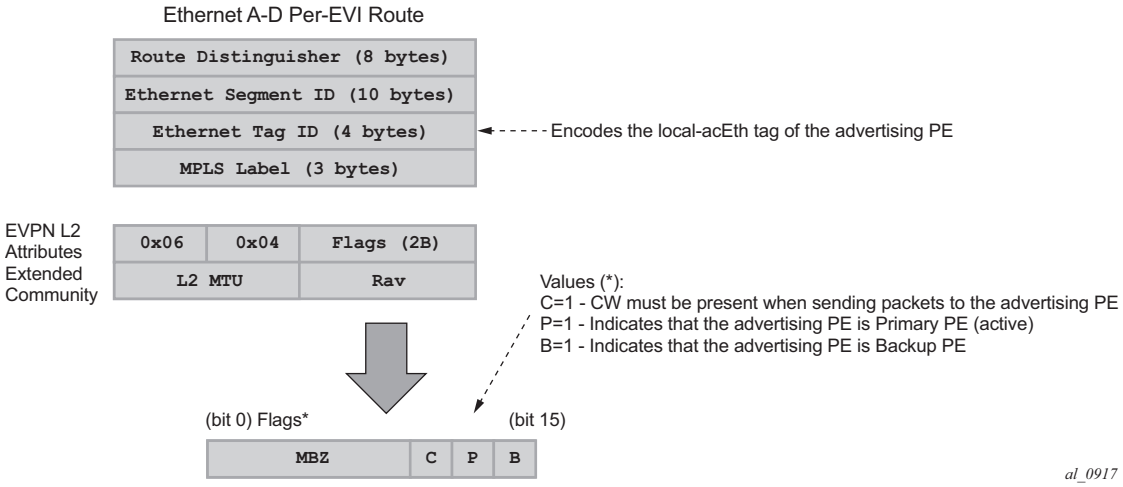

Assuming that the advertising PE has an access SAP-SDP or Spoke-SDP that is not part of an Ethernet Segment (ES), the PE populates the fields of the AD per-EVI route with the following values.

- Ethernet Tag ID field is encoded with the value configured by the user in the **service>bgp-evpn>local-ac-name>eth-tag <***value***>** command.

- RD and MPLS label values are encoded as specified in RFC 7432.
- ESI is 0.
- The route is sent along an EVPN L2 attributes extended community, as specified in IETF Draft *draft-ietf-bess-evpn-vpws*, where:
  - type and subtype are 0x06 and 0x04 as allocated by IANA
  - flag C is set if **control-word** is configured in the service
  - P and B flags are zero
  - L2 MTU is encoded with **service-mtu** configured in the Epipe service

If the advertising PE has an access SAP-SDP or Spoke-SDP that is part of an ES, the AD per-EVI route is sent with the information described in the preceding list, with the following minor differences.

- The ESI encodes the corresponding non-zero value.
- The P and B flags are set in the following cases.
  - All-active multi-homing
    - All PEs that are part of the ES always set the P flag.
    - The B flag is never set in the all-active multi-homing ES case.
  - Single-active multi-homing
    - Only the DF PE sets the P bit for an EVI and the remaining PEs send it as P=0.
    - Only the backup DF PE sets the B bit.

      If more than two PEs are present in the same single-active ES, the backup PE is the winner of a second DF election (excluding the DF). The remaining non-DF PEs send B=0.

Also, ES and AD per-ES routes are advertised and processed for the Ethernet-Segment, as described in RFC7432 Ethernet Segments. The ESI label sent with the AD per-ES route is used by BUM traffic on VPLS services; it is not used for Epipe traffic.

5.3.4.2 EVPN for MPLS Tunnels in Epipe Services (EVPN-VPWS)

BGP-EVPN can be enabled in Epipe services with either SAPs or spoke-SDPs at the access, as shown in Figure 168.

Figure 168 **EVPN-MPLS VPWS**

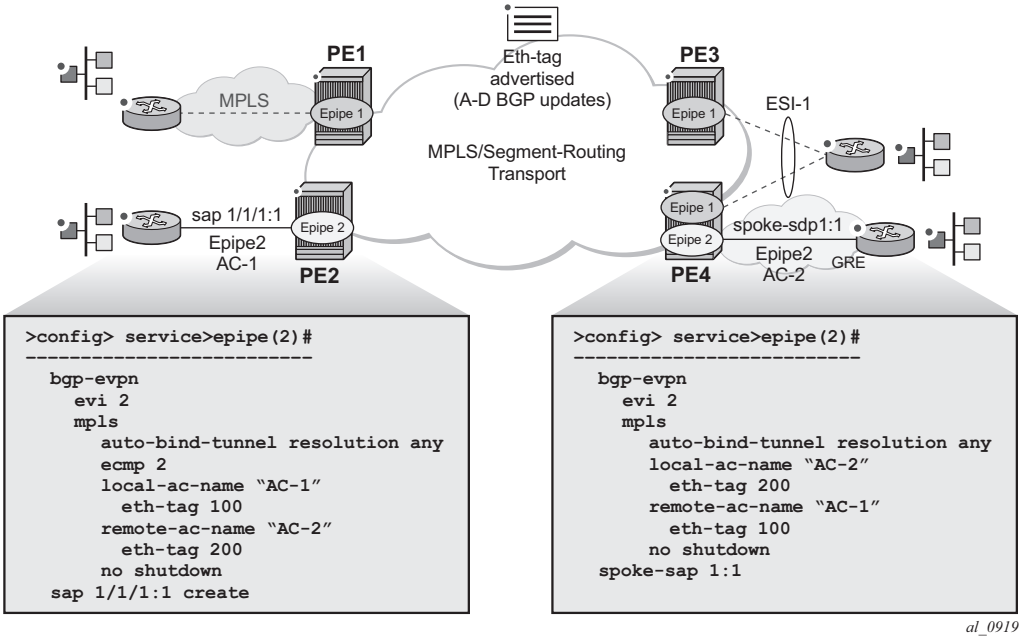

EVPN-VPWS is supported in MPLS networks that also run EVPN-MPLS in VPLS services. From a control plane perspective, EVPN-VPWS is a simplified point-to-point version of RFC 7432 for E-Line services for the following reasons.

- EVPN-VPWS does not use inclusive multicast, MAC/IP routes or IP-Prefix routes.
- A-D Ethernet per EVI routes are used to advertise the local attachment circuit identifiers at each side of the VPWS instance. The attachment circuit identifiers are configured as local and remote ethernet tags. When an AD per EVI route is imported and the ethernet-tag matches the configured remote ethernet-tag, an EVPN destination is created for the Epipe.

In the following configuration example, Epipe 2 is an EVPN-VPWS service between PE2 and PE4 (as shown in Figure 168):

```
PE2>config>service>epipe(2)#
-----------------------
  bgp-evpn
    evi 2
    mpls
      auto-bind-tunnel resolution any
      ecmp 2
      local-ac-name "AC-1"
        eth-tag 100
      remote-ac-name "AC-2"
        eth-tag 200
```

```
      no shutdown
  sap 1/1/1:1 create
PE4>config>service>epipe(2)#
-----------------------
  bgp-evpn
    evi 2
    mpls
      auto-bind-tunnel resolution any
      local-ac-name "AC-2"
        eth-tag 200
      remote-ac-name "AC-1"
        eth-tag 100
      no shutdown
 spoke-sdp 1:1
```

The following considerations apply to the example configuration.

- The **evi** is used to auto-derive the route-target/route-distinguisher of the service. The **evi** values must be unique in the system no matter, irrespecitve of the type of service they are assigned (EPIPE or VPLS).
- Support for the following **bgp-evpn** commands in Epipe services is the same as in VPLS services:
  - **mpls auto-bind-tunnel**
  - **mpls control-word**
  - **mpls entropy-label**
  - **mpls force-vlan-vc-forwarding**
  - **mpls shutdown**
- The following **bgp-evpn** commands identify the local and remote attachment circuits, with the configured eth-tags encoded in the advertised and received A-D Ethernet per-EVI routes:
  - **local-ac-name <***name***>**
  - **local-ac-name <***name***> eth-tag <**tag-value**>;** where tag-value **[1.. 16777215]**
  - **remote-ac-name <***name***>**
  - **remote-ac-name <***name***> eth-tag <**tag-value**>;** where tag-value **[1.. 16777215]**
  - Changes on remote eth-tags are allowed without shutting down **bgp-evpn mpls** or the Epipe service. The **local-ac eth-tag** value cannot be changed without **bgp-evpn mpls shutdown**.
  - Both local and remote **eth-tags** are mandatory to bring up the Epipe service.

EVPN-VPWS Epipes can also be configured with the following characteristics.

- Access attachment circuits can be SAPs or spoke-SDPs. Only manually configured spoke-SDPs are supported; BGP-VPWS and **endpoints** are not supported. The **vc-switching** configuration is not supported on **bgp-evpn** enabled pipes.
- EVPN-VPWS Epipes support **control-word** and **entropy-label**.

  When **bgp-evpn>mpls>control-word** is configured, the PE will set the C bit in its AD per-EVI advertisement and send the control-word in the data path. In this case, the PE will also expect the control-word to be received. If there is a mismatch between the received control-word and the configured control-word, the system will not setup the EVPN destination; as a result, the service will not come up.
- EVPN-VPWS Epipes can advertise the Layer 2 (service) MTU and check its consistency as follows:
  - The advertised MTU value will be taken from the configured **service-mtu** in the Epipe service.
  - The received L2 MTU will be checked and compared with the local value. In case of a mismatch between the received MTU and the configured **service-mtu**, the system will not setup the EVPN destination; as a result, the service will not come up.
  - The system will not check the network port MTU value.
  - If the received L2 MTU value is 0, the MTU will be ignored.

5.3.4.3 Using A/S PW and MC-LAG with EVPN-VPWS Epipes

The use of A/S PW (for access spoke-SDPs) and MC-LAG (for access SAPs) provides an alternative redundant solution for EVPN-VPWS that do not use the EVPN multi-homing procedures described in IETF Draft *draft-ietf-bess-evpn-vpws*. Figure 169 shows the use of both mechanisms in a single Epipe.

Figure 169 **A/S PW and MC-LAG Support on EVPN-VPWS**

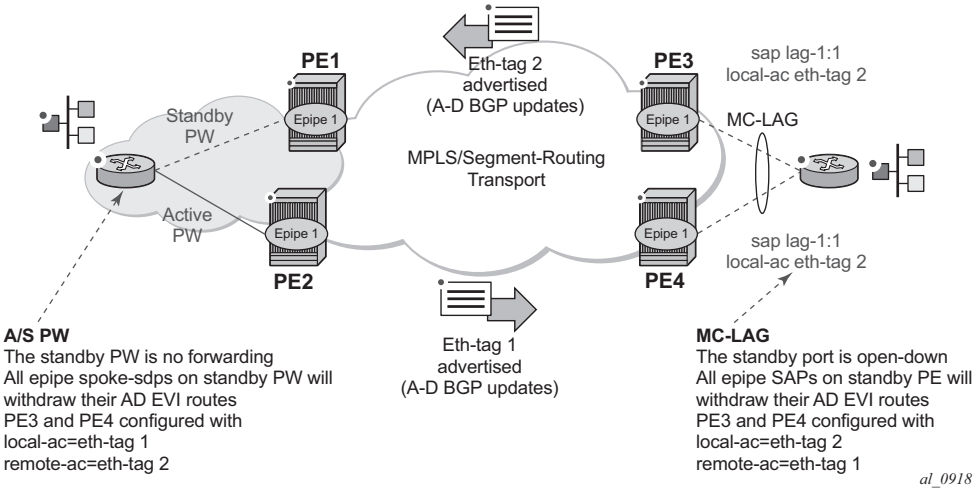

In Figure 169, an A/S PW connects the CE to PE1 and PE2 (left-hand side of the diagram), and an MC-LAG connects the CE to PE3 and PE4 (right-hand side of the diagram). As EVPN multi-homing is not used, there are no AD per-ES routes or ES routes in this example. The redundancy is handled as follows:

- PE1 and PE2 are configured with EPIPE-1, where a spoke-SDP connects the service in each PE to the access CE. The **local-ac eth-tag** is 1 and the **remote-ac eth-tag** is 2 (in PE1/PE2).
- PE3 and PE4 are configured with EPIPE-1, where each PE has a lag SAP that belongs to a previously configured MC-LAG construct. The **local-ac eth-tag** is 2 and the **remote-ac eth-tag** is 1.
- An endpoint and A/S PW is configured on the CE on the left-hand side of the diagram. PE1/PE2 are able to advertise **eth-tag** 1 based on the oper-status or the forwarding status of the spoke-SDP.

  For example, if PE1 receives a standby PW status indication from the CE and the previous status was forward, it will withdraw the AD EVI route for eth-tag 1. If PE2 receives a forward PW status indication and the previous status was standby or down, it will advertise the AD EVI route for eth-tag 1.
- The user can configure MC-LAG for access SAPs using the example configuration of PE3 and PE4 shown in Figure 169. In this case, the MC-LAG will determine which of the two chassis is active or standby.

  If PE4 becomes the standby chassis, the entire LAG port will be brought down. As a result, the SAP will go operationally down and PE4 will withdraw any previous AD EVI route for eth-tag 2.

  If PE3 becomes the active chassis, the LAG port becomes operationally up. As a result, the SAP and the PE3 will advertise the AD per-EVI route for eth-tag 2.

5.3.4.4 EVPN Multi-homing for EVPN-VPWS Services

EVPN multi-homing is supported for EVPN-VPWS Epipe services with the following considerations.

- Single-active and all-active multi-homing is supported for SAPs and spoke-SDPs.
- Ethernet Segments (ES) can be shared between the Epipe and VPLS services for LAGs, ports, and SDPs.
- A split-horizon function is not required because there is no traffic between the Designated Forwarder (DF) and the non-DF for Epipe services. As a result the ESI label is never used. For this reason, the **ethernet-segment multi-homing single-active no-esi-label** and **ethernet-segment source-bmac-lsb** commands do not affect Epipe services.
- The local Ethernet Tag values must match on all PEs that are part of the same ES, regardless of the multi-homing mode. The PEs in the ES use the AD per-EVI routes from the peer PEs to validate the PEs as DF election candidates for a specific EVI.

The DF election for Epipes that is defined in an all-active multi-homing ES is not relevant because all PEs in the ES behave in the same way as follows:

- All PEs send P=1 on the AD per-EVI routes.
- All PEs can send upstream and downstream traffic, regardless of whether the traffic is unicast, multicast, or broadcast (all traffic is treated as unicast in the Epipe services).

Therefore, the following **tools** command shows “N/A” when all-active multi-homing is configured.

```
*A:PE-2# tools dump service system bgp-evpn ethernet-segment "ESI-12" evi 6000 df
[03/18/2016 20:31:35] All Active VPWS - DF N/A
```

Aliasing is supported for traffic sent to an Ethernet-Segment destination. If **ecmp** is enabled on the ingress PE, per-flow load-balancing is performed to all PEs that advertise P=1. The PEs that advertise P=0 are not considered as next hops for an ES destination.

Note: The ingress PE will load-balance the traffic if shared-queuing or ingress policing is enabled on the access SAPs.

Although DF election is not relevant for Epipes in an all-active multi-homing ES, it is essential for the following forwarding and backup functions in a single-active multi-homing ES.

- The PE elected as DF is the primary PE for the Ethernet Segment in the Epipe. The primary PE unblocks the SAP or Spoke-SDP for upstream and downstream traffic; the remaining PEs in the ES bring their ES SAPs or Spoke-SDPs operationally down.
- The DF candidate list is built from the PEs sending ES routes for the same ES and is pruned for a specific service, depending on the availability of the AD per-ES and per-EVI routes.
- When the SAP/Spoke-SDPs (part of the ES) come up, the AD per-EVI routes are sent with P=B=0. The remote PEs do not send traffic at this stage.

  The remote PEs do not start sending traffic until the DF election process is completed and the **es-activation-timer** has expired, and the PEs advertise AD per-EVI routes with P and B bits different from zero.
- The backup PE function is supported as defined in IETF Draft *draft-ietf-bess-evpn-vpws*. The primary PE, backup, or none status is signaled by the PEs (part of the same single-active MH ES) in the P or B flags of the EVPN L2 attributes extended community. Figure 170 shows the advertisement and use of the primary, backup, or none indication by the PEs in the ES.

Figure 170 **EVPN-VPWS Single-active Multi-homing**

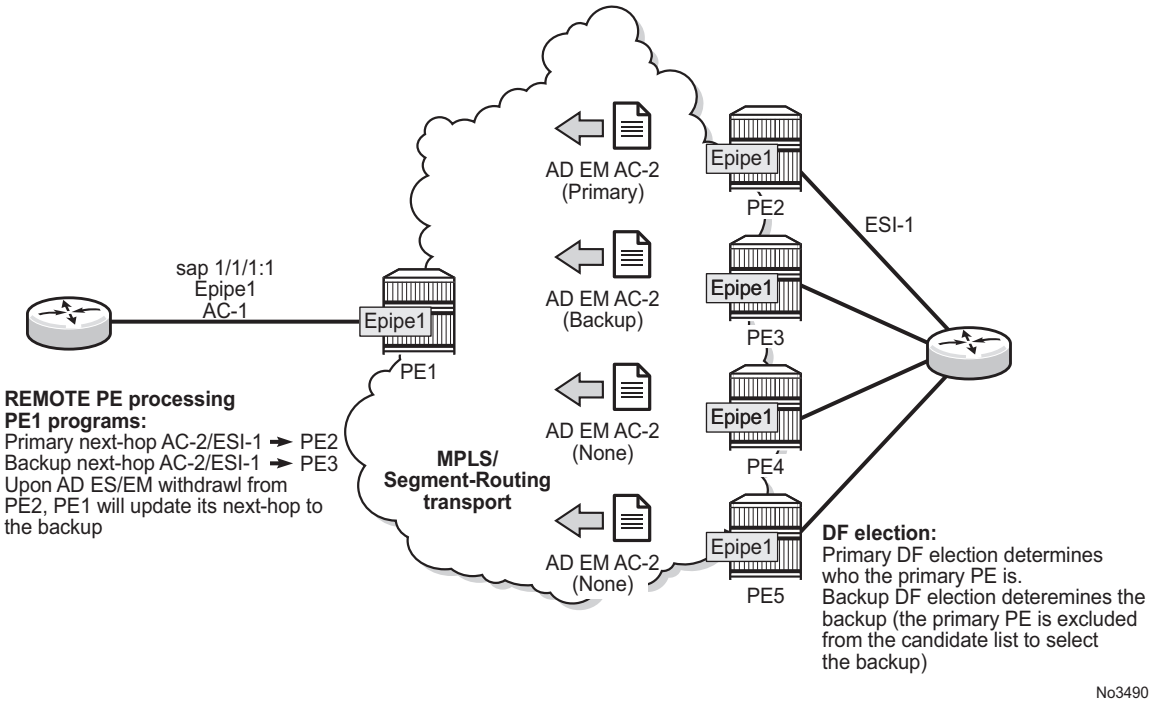

As specified in RFC7432, the remote PEs in VPLS services have knowledge of the primary PE in the remote single-active ES, based on the advertisement of the MAC/IP routes (because only the DF will learn and advertise MAC/IP routes).

Because there are no MAC or IP routes in EVPN-VPWS, the remote PEs can forward the traffic based on the P/B bits. The process is described in the following list; all PE references used see the example in Figure 170.

- The DF PE for an EVI (PE2 in Figure 170) sends P=1 and B=0.
- For each ES or EVI, a second DF election is run among the PEs in the backup candidate list to elect the backup PE. The backup PE sends P=0 and B=1 (PE3 in Figure 170).
- All remaining multi-homing PEs send P=B=0 (PE4 and PE5).
- At the remote PEs (PE1 in Figure 170), the P and B flags are used to identify the primary and backup PEs within the ES destination. The traffic is then sent to the primary PE, provided that it is active.
  - When a remote PE receives the withdrawal of an Ethernet A-D per ES (or EVI) route from the primary PE, the remote PE immediately switches the traffic to the backup PE for the affected EVIs.
  - The backup PE takes over immediately without waiting for the **es-activation-timer** to bring up its SAP or spoke-SDP.
  - The **bgp-evpn mpls ecmp** setting also governs the forwarding in single-active multi-homing, regardless of the single-active multi-homing bit in the AD per-ES route received at the remote PE (PE1 in Figure 170).
    - PE1 always sends the traffic to the primary remote PE (the owner of the P=1 bit). In case of multiple primary PEs and ecmp>1, PE1 will load-balance the traffic to all the primary PEs, regardless of the multi-homing mode.
    - If the last primary PE withdraws its AD per-EVI or ES route, PE1 sends the traffic to the backup PE or PEs. In case of multiple backup PEs and ecmp>1, PE1 load-balances the traffic to the backup PEs.

5.3.5 EVPN for MPLS Tunnels in Routed VPLS Services

EVPN-MPLS and IP-prefix advertisement (enabled by the **ip-route-advertisement** command) are fully supported in routed VPLS services and provide the same feature-set as EVPN-VXLAN. The following capabilities are supported in a service where **bgp-evpn mpls** is enabled:

- R-VPLS with VRRP support on the VPRN interfaces
- R-VPLS support including **ip-route-advertisement** with regular interfaces

This includes the advertisement and process of ip-prefix routes defined in IETF Draft *draft-ietf-bess-evpn-prefix-advertisement* with the appropriate encoding for EVPN-MPLS.

- R-VPLS support including **ip-route-advertisement** with **evpn-tunnel** interfaces
- R-VPLS with IPv6 support on the VPRN IP interface

5.3.5.1 EVPN-MPLS Multi-Homing and Passive VRRP

SAP and spoke-SDP based Ethernet Segments are supported on R-VPLS services where **bgp-evpn mpls** is enabled.

Figure 171 shows an example of EVPN-MPLS multi-homing in R-VPLS services, with the following assumptions.

- There are two subnets for a specific customer (for example, EVI1 and EVI2 in Figure 171), and a VPRN is instantiated in all the PEs for efficient inter-subnet forwarding.
- A "backhaul" R-VPLS with **evpn-tunnel** mode enabled is used in the core to interconnect all the VPRNs. EVPN IP-prefix routes are used to exchange the prefixes corresponding to the two subnets.
- An all-active ES is configured for EVI1 on PE1 and PE2.
- A single-active ES is configured for EVI2 on PE3 and PE4.

Figure 171 **EVPN-MPLS Multi-Homing in R-VPLS Services**

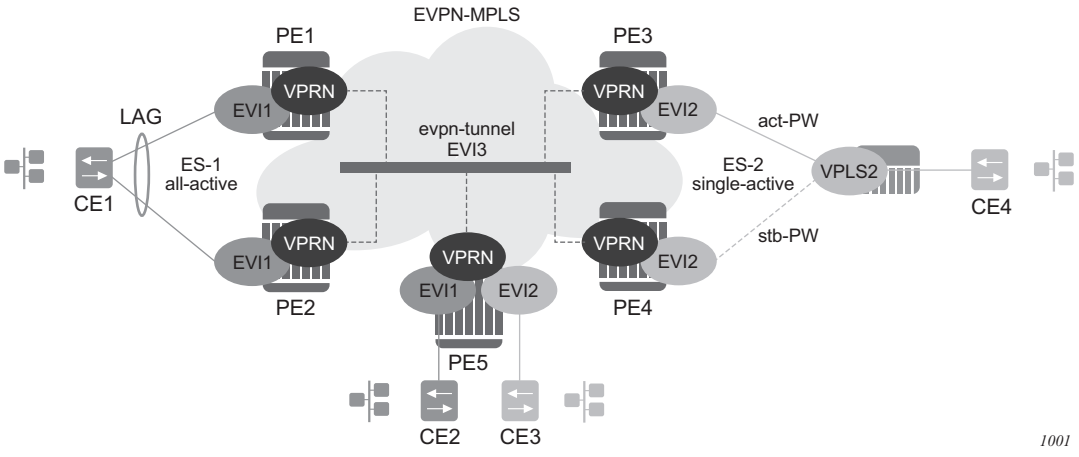

In the example in Figure 171, the hosts connected to CE1 and CE4 could use regular VRRP for default gateway redundancy; however, this may not be the most efficient way to provide upstream routing.

For example, if PE1 and PE2 are using regular VRRP, the upstream traffic from CE1 may be hashed to the backup IRB VRRP interface, instead of being hashed to the master. The same thing may occur for single-active multi-homing and regular VRRP for PE3 and PE4. The traffic from CE4 will be sent to PE3, while PE4 may be the VRRP master. In that case, PE3 will have to send the traffic to PE4, instead of route it directly.

In both cases, unnecessary bandwidth between the PEs is used to get to the master IRB interface. In addition, VRRP scaling is limited if aggressive keepalive timers are used.

Because of these issues, passive VRRP is recommended as the best method when EVPN-MPLS multi-homing is used in combination with R-VPLS redundant interfaces.

Passive VRRP is a VRRP setting in which the transmission and reception of keepalive messages is completely suppressed, and therefore the VPRN interface always behaves as the master. Passive VRRP is enabled by adding the **passive** keyword to the VRRP instance at creation, as shown in the following examples:

config service vprn 1 interface int-1 vrrp 1 passive

config service vprn 1 interface int-1 ipv6 vrrp 1 passive

For example, if PE1, PE2, and PE5 in Figure 171 use passive VRRP, even if each individual R-VPLS interface has a different MAC/IP address, because they share the same VRRP instance 1 and the same backup IP, the three PEs will own the same virtual MAC and virtual IP address (for example, 00-00-5E-00-00-01 and 10.0.0.254). The virtual MAC is auto-derived from 00-00-5E-00-00-VRID per RFC 3768. The following is the expected behavior when passive VRRP is used in this example.

- All R-VPLS IRB interfaces for EVI1 have their own physical MAC/IP address; they will also own the same default gateway virtual MAC and IP address.
- All EVI1 hosts have a unique configured default gateway; for example, 10.0.0.254.
- When CE1 or CE2 send upstream traffic to a remote subnet, the packets are routed by the closest PE because the virtual MAC is always local to the PE.

  For example, the packets from CE1 hashed to PE1 will be routed at PE1. The packets from CE1 hashed to PE2 will be routed directly at PE2.
- Downstream packets (for example, packets from CE3 to CE1), are routed directly by the PE to CE1, regardless of the PE to which PE5 routed the packets.

  For example, the packets from CE3 sent to PE1 will be routed at PE1. The packets from CE3 sent to PE2 will be routed at PE2.
- In case of ES failure in one of the PEs, the traffic will be forwarded by the available PE.

For example, if the packets routed by PE5 arrive at PE1 and the link to CE1 is down, then PE1 will send the packets to PE2. PE2 will forward the packets to CE1 even if the MAC source address of the packets matches PE2's virtual MAC address. Virtual MACs bypass the R-VPLS interface MAC protection.

The following list summarizes the advantages of using passive VRRP mode versus regular VRRP for EVPN-MPLS multi-homing in R-VPLS services.

- Passive VRRP does not require multiple VRRP instances to achieve default gateway load-balancing. Only one instance per R-VPLS, hence only one default gateway, is needed for all the hosts.
- The convergence time for link/node failures is not impacted by the VRRP convergence, as all the nodes in the VRRP instance are master routers.
- Passive VRRP scales better than VRRP, as it does not use keepalive or BFD messages to detect failures and allow the backup to take over.

5.3.6 PBB-EVPN

This section contains information on PBB-EVPN.

5.3.6.1 BGP-EVPN Control Plane for PBB-EVPN

PBB-EVPN uses a reduced subset of the routes and procedures described in RFC 7432. The supported routes are:

- ES routes
- MAC/IP routes
- Inclusive Multicast Ethernet Tag routes.

5.3.6.1.1 EVPN Route Type 3 - Inclusive Multicast Ethernet Tag Route

This route is used to advertise the ISIDs that belong to I-VPLS services as well as the default multicast tree. PBB-epipe ISIDs are not advertised in Inclusive Multicast routes. The following fields are used:

- Route Distinguisher: Taken from the RD of the B-VPLS service within the BGP context. The RD can be configured or derived from the **bgp-evpn evi** value.
- Ethernet Tag ID: Encodes the ISID for a specified I-VPLS.

- IP address length: Always 32.
- Originating router's IP address: Carries the system address (IPv4 only).
- PMSI attribute:
  - Tunnel type = Ingress replication (6).
  - Flags = Leaf no required.
  - MPLS label: Carries the MPLS label allocated for the service in the high-order 20 bits of the label field.

Note: This label will be the same label used in the BMAC routes for the same B-VPLS service unless **bgp-evpn mpls ingress-replication-bum-label** is configured in the B-VPLS service.

  - Tunnel endpoint = equal to the originating IP address.

Note: The mLDP P2MP tunnel type is supported on PBB-EPVN services, but it can be used in the default multicast tree only.

5.3.6.1.2 EVPN Route Type 2 - MAC/IP Advertisement Route (or BMAC Routes)

The 7750 SR, 7450 ESS, or 7950 XRS will generate this route type for advertising BMAC addresses for the following:

- Learned MACs on B-SAPs or B-SDP-bindings - if mac-advertisement is enabled.
- Conditional static MACs - if mac-advertisement is enabled.
- B-VPLS shared-BMACs (**source-bmacs**) and dedicated-BMACs (**es-bmacs**).

The route type 2 generated by the router uses the following fields and values:

- Route Distinguisher—Taken from the RD of the VPLS service within the BGP context. The RD can be configured or derived from the **bgp-evpn evi** value.
- Ethernet Segment Identifier (ESI):
  - ESI = 0 for the advertisement of source-bmac, es-bmacs, sap-bmacs, or sdp-bmacs if no multi-homing or single-active multi-homing is used.
  - ESI=MAX-ESI (0xFF..FF) in the advertisement of es-bmacs used for all-active multi-homing.
  - ESI different from zero or MAX-ESI for learned BMACs on B-SAPs/SDP-bindings if EVPN multi-homing is used on B-VPLS SAPs and SDP-bindings.

- Ethernet Tag ID: 0.

  Note: A different Ethernet Tag value may be used only when **send-bvpls-evpn-flush** is enabled.
- MAC address length: Always 48.
- BMAC Address learned, configured, or system-generated.
- IP address length zero and IP address omitted.
- MPLS Label 1: carries the MPLS label allocated by the system to the B-VPLS service. The label value is encoded in the high-order 20 bits of the field and will be the same label used in the routes type 3 for the same service unless bgp-evpn mpls ingress-replication-bum-label is configured in the service.
- The MAC Mobility extended community:
  - The mac mobility extended community is used in PBB-EVPN for CMAC flush purposes if per ISID load balancing (single-active multi-homing) is used and a source-bmac is used for traffic coming from the ESI.

    If there is a failure in one of the ES links, CMAC flush through the withdrawal of the BMAC cannot be done (other ESIs are still working); therefore, the mac mobility extended community is used to signal CMAC flush to the remote PEs.
  - When a dedicated es-bmac per ESI is used, the mac flush can be based on the withdrawal of the BMAC from the failing node.
  - es-bmacs will be advertised as static (sticky bit set).
  - Source-bmacs will be advertised as static MACs (sticky bit set). In the case of an update, if advertised to indicate that CMAC flush is needed, the mac mobility extended community will be added to the BMAC route including a higher sequence number (than the one previously advertised) in addition to the sticky bit.

5.3.6.1.3 EVPN Route Type 4 - Ethernet Segment Route

This route type is used for DF election as described in section BGP-EVPN Control Plane for MPLS Tunnels.

Note: The EVPN route type 1—Ethernet Auto Discovery route is not used in PBB-EVPN.

5.3.6.2 PBB-EVPN for I-VPLS and PBB Epipe Services

The 7750 SR, 7450 ESS, and 7950 XRS SR OS implementation of PBB-EVPN reuses the existing PBB-VPLS model, where N I-VPLS (or Epipe) services can be linked to a B-VPLS service. BGP-EVPN is enabled in the B-VPLS and the B-VPLS becomes an EVI (EVPN Instance). Figure 172 shows the PBB-EVPN model in the SR OS.

Figure 172 **PBB-EVPN for I-VPLS and PFF Epipe Services**

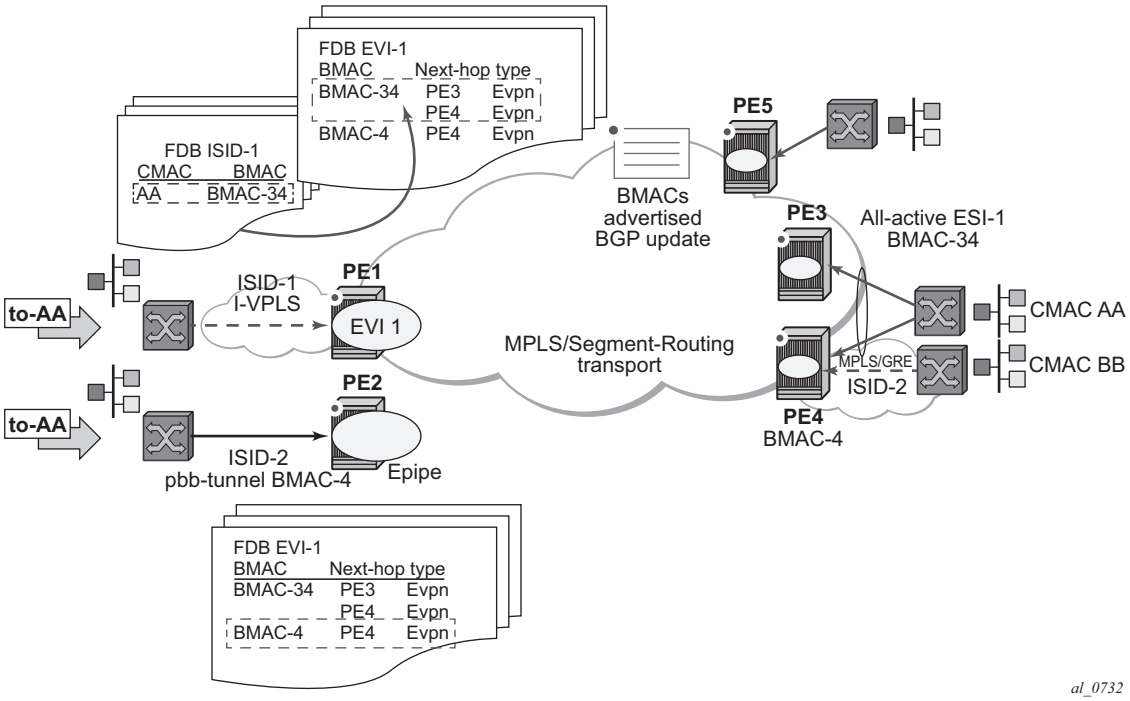

Each PE in the B-VPLS domain will advertise its **source-bmac** as either configured in **(b)vpls>pbb>source-bmac** or auto-derived from the chassis mac. The remote PEs will install the advertised BMACs in the B-VPLS FDB. If a specified PE is configured with an **ethernet-segment** associated with an I-VPLS or PBB Epipe, it may also advertise an **es-bmac** for the ethernet-segment.

In the example shown in Figure 172, when a frame with MAC DA = AA gets to PE1, a mac lookup is performed on the I-VPLS FDB and BMAC-34 is found. A BMAC lookup on the B-VPLS FDB will yield the next-hop (or next-hops if the destination is in an all-active ethernet-segment) to which the frame is sent. As in PBB-VPLS, the frame will be encapsulated with the corresponding PBB header. A label specified by EVPN for the B-VPLS and the MPLS transport label are also added.

If the lookup on the I-VPLS FDB fails, the system will send the frame encapsulated into a PBB packet with BMAC DA = Group BMAC for the ISID. That packet will be distributed to all the PEs where the ISID is defined and will contain the EVPN label distributed by the Inclusive Multicast routes for that ISID, as well as the transport label.

For PBB-Epipes, all the traffic is sent in a unicast PBB packet to the BMAC configured in the **pbb-tunnel**.

The following CLI output shows an example of the configuration of an I-VPLS, PBB-Epipe, and their corresponding B-VPLS.

```
*A:PE-1>config#

service vpls 1 b-vpls create
  description "pbb-evpn-service"
  service-mtu 2000
  pbb
    source-bmac 00:00:00:00:00:03
  bgp
  bgp-evpn
    evi 1
    vxlan
      shutdown
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  sap 1/1/1:1 create
  exit
  spoke-sdp 1:1 create


*A:PE-1>config#

service vpls 101 i-vpls create
  pbb
    backbone-vpls 1
  sap 1/2/1:101 create
  spoke-sdp 1:102 create

*A:PE-1>config#

service epipe 102 create
  pbb
    tunnel 1 backbone-dest-mac 00:00:00:00:00:01 isid 102
  sap 1/2/1:102 create
```

Configure the bgp-evpn context as described in section EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS).

Some EVPN configuration options are not relevant to PBB-EVPN and are not supported when bgp-evpn is configured in a B-VPLS; these are as follows:

- bgp-evpn> [no] ip-route-advertisement

- bgp-evpn> [no] unknown-mac-route
- bgp-evpn> vxlan [no] shutdown
- bgp-evpn>mpls>force-vlan-vc-forwarding

When **bgp-evpn>mpls no shutdown** is added to a specified B-VPLS instance, the following considerations apply:

- BGP-AD is supported along with EVPN in the same B-VPLS instance.
- The following B-VPLS and BGP-EVPN commands are fully supported:
  - **vpls>backbone-vpls**
  - **vpls>backbone-vpls>send-flush-on-bvpls-failure**
  - **vpls>backbone-vpls>source-bmac**
  - **vpls>backbone-vpls>use-sap-bmac**
  - **vpls>backbone-vpls>use-es-bmac** (See PBB-EVPN Multi-Homing in I-VPLS and PBB Epipe Services for more information)
  - **vpls>isid-policies**
  - **vpls>static-mac**
  - **vpls>SAP or SDP-binding>static-isid**
  - **bgp-evpn>mac-advertisement -** this command will only have affect on the 'learned' BMACs on saps or sdp-bindings and not on the system BMAC or sap/es-bmacs being advertised.
  - **bgp-evpn>mac-duplication** and settings.
  - **bgp-evpn>mpls>auto-bind-tunnel** and options.
  - **bgp-evpn>mpls>ecmp**
  - **bgp-evpn>mpls>control-word**
  - **bgp-evpn>evi**
  - **bgp-evpn>mpls>ingress-replication-bum-label**

5.3.6.2.1 Flood Containment for I-VPLS Services

In general, PBB technologies in the 7750 SR, 7450 ESS, or 7950 XRS SR OS support a way to contain the flooding for a specified I-VPLS ISID, so that BUM traffic for that ISID only reaches the PEs where the ISID is locally defined. Each PE will create an MFIB per I-VPLS ISID on the B-VPLS instance. That MFIB supports SAP or SDP-bindings endpoints that can be populated by:

- MMRP in regular PBB-VPLS
- IS-IS in SPBM

In PBB-EVPN, B-VPLS EVPN endpoints can be added to the MFIBs using EVPN Inclusive Multicast Ethernet Tag routes.

The example in Figure 173 shows how the MFIBs are populated in PBB-EVPN.

Figure 173 **PBB-EVPN and I-VPLS Flooding Containment**

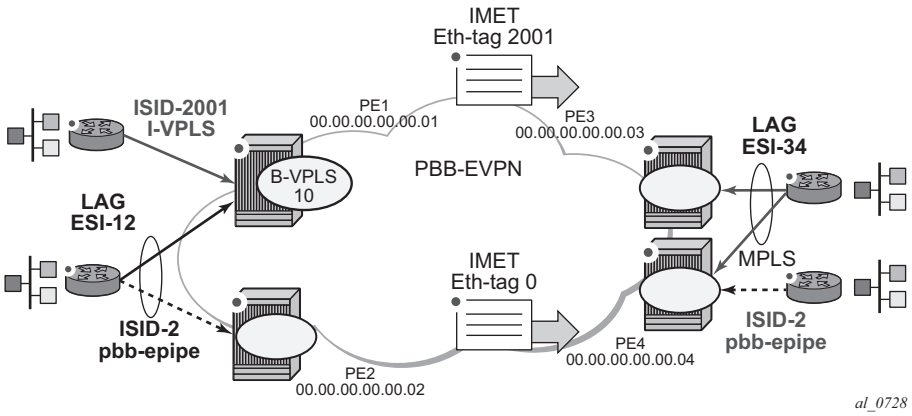

When the B-VPLS 10 is enabled, PE1 will advertise as follows:

- A BMAC route containing PE1's system BMAC (00:01 as configured in **pbb>source-bmac**) along with an MPLS label.
- An Inclusive Multicast Ethernet Tag route (IMET route) with Ethernet-tag = 0 that will allow the remote B-VPLS 10 instances to add an entry for PE1 in the default multicast list.

Note: The MPLS label that will be advertised for the MAC routes and the inclusive multicast routes for a specified B-VPLS can be the same label or a different label. As in regular EVPN-MPLS, this will depend on the **[no] ingress-replication-bum-label** command.

When I-VPLS 2001 (ISID 2001) is enabled as per the CLI in the preceding section, PE1 will advertise as follows:

- An additional inclusive multicast route with Ethernet-tag = 2001. This will allow the remote PEs to create an MFIB for the corresponding ISID 2001 and add the corresponding EVPN binding entry to the MFIB.

  This default behavior can be modified by the configured **isid-policy**. For instance, for ISIDs 1-2000, configure as follows:

```
isid-policy
    entry 10 create
      no advertise-local
      range 1 to 2000
      use-def-mcast
```

This configuration has the following effect for the ISID range:

- **no advertise-local** instructs the system to not advertise the local active ISIDs contained in the 1 to 2001 range.
- **use-def-mcast** instructs the system to use the default flooding list as opposed to the MFIB.

The ISID flooding behavior on B-VPLS saps and sdp-bindings is as follows:

- B-VPLS saps and sdp-bindings are only added to the TLS-multicast list and not to the MFIB list (unless **static-isids** are configured, which is only possible for saps/sdp-bindings and not BGP-AD spoke-sdps).

  As a result, if the system needs to flood ISID BUM traffic and the ISID is also defined in remote PEs connected through saps or spoke-sdps without **static-isids**, then an **isid-policy** must be configured for the ISID so that the ISID uses the default multicast list.
- When an **isid-policy** is configured and a range of ISIDs use the default multicast list, the remote PBB-EVPN PEs will be added to the default multicast list as long as they advertise an IMET route with an ISID included in the policy's ISID range. PEs advertising IMET routes with Ethernet-tag = 0 are also added to the default multicast list (7750 SR, 7450 ESS, or 7950 XRS SR OS behavior).
- The B-VPLS 10 also allows the ISID flooding to legacy PBB networks via B-SAPs or B-SDPs. The legacy PBB network BMACs will be dynamically learned on those saps/binds or statically configured through the use of conditional **static-macs**. The use of **static-isids** is required so that non-local ISIDs are advertised.

```
sap 1/1/1:1 create
exit
spoke-sdp 1:1 create
  static-mac
      mac 00:fe:ca:fe:ca:fe create sap 1/1/1:1 monitor fwd-status
  static-isid
      range 1 isid 3000 to 5000 create
```

Note: The configuration of PBB-Epipes does not trigger any IMET advertisement.

5.3.6.2.2 PBB-EVPN and PBB-VPLS Integration

The 7750 SR, 7450 ESS, and 7950 XRS SR OS EVPN implementation supports *draft-ietf-bess-evpn-vpls-seamless-integ* so that PBB-EVPN and PBB-VPLS can be integrated into the same network and within the same B-VPLS service.

All the concepts described in section EVPN and VPLS Integration are also supported in B-VPLS services so that B-VPLS SAP or SDP-bindings can be integrated with PBB-EVPN destination bindings. The features described in that section also facilitate a smooth migration from B-VPLS SDP-bindings to PBB-EVPN destination bindings.

5.3.6.2.3 PBB-EVPN Multi-Homing in I-VPLS and PBB Epipe Services

The 7750 SR, 7450 ESS, and 7950 XRS SR OS PBB-EVPN implementation supports all-active and single-active multi-homing for I-VPLS and PBB Epipe services.

PBB-EVPN multi-homing reuses the **ethernet-segment** concept described in section EVPN Multi-Homing in VPLS Services. However, unlike EVPN-MPLS, PBB-EVPN does not use AD routes; it uses BMACs for split-horizon checks and aliasing.

System BMAC Assignment in PBB-EVPN

RFC 7623 describes two types of BMAC assignments that a PE can implement:

- Shared BMAC addresses that can be used for single-homed CEs and a number of multi-homed CEs connected to ethernet-segments.
- Dedicated BMAC addresses per ethernet-segment.

In this document and in 7750 SR, 7450 ESS, and 7950 XRS SR OS terminology:

- A *shared-bmac* (in IETF) is a **source-bmac** as configured in **service>(b)vpls>pbb>source-bmac**
- A *dedicated-bmac* per ES (in IETF) is an **es-bmac** as configured in **service>pbb>use-es-bmac**

BMAC selection and use depends on the multi-homing model; for single-active mode, the type of BMAC will impact the flooding in the network as follows:

- All-active multi-homing requires **es-bmacs**.
- Single-active multi-homing can use **es-bmacs** or **source-bmacs**.
  - The use of **source-bmacs** minimizes the number of BMACs being advertised but has a larger impact on CMAC flush upon ES failures.
  - The use of **es-bmacs** optimizes the CMAC flush upon ES failures at the expense of advertising more BMACs.

PBB-EVPN all-active multi-homing service model

Figure 174 shows the use of all-active multi-homing in the 7750 SR, 7450 ESS, and 7950 XRS SR OS PBB-EVPN implementation.

Figure 174 **PBB-EVPN All-Active Multi-Homing**

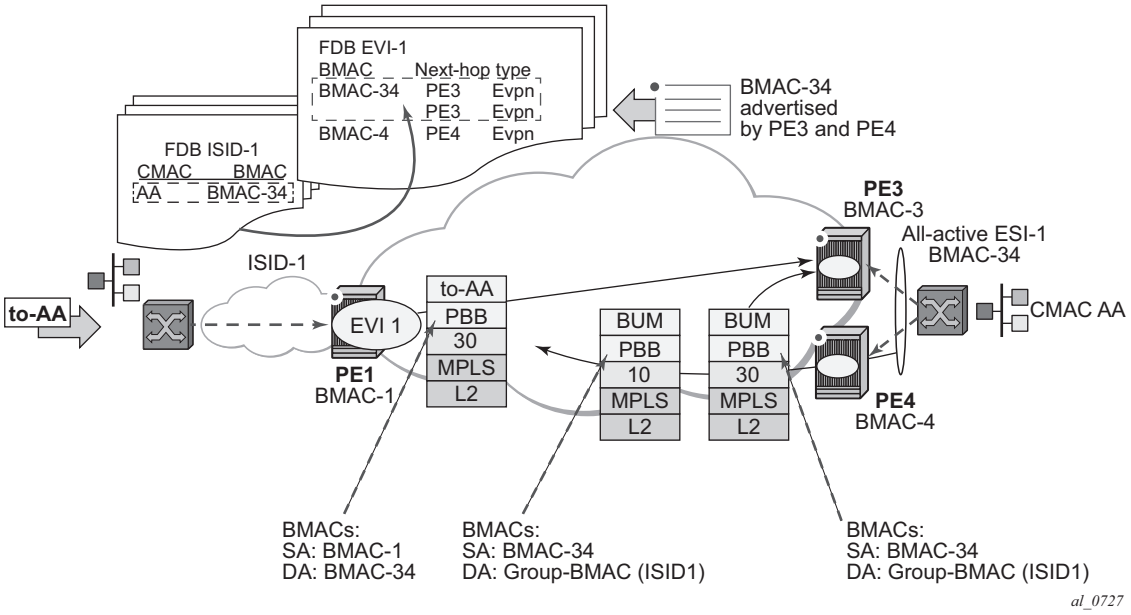

For example, the following shows the ESI-1 and all-active configuration in PE3 and PE4. As in EVPN-MPLS, all-active multi-homing is only possible if a LAG is used at the CE. All-active multi-homing uses es-bmacs, that is, each ESI will be assigned a dedicated BMAC. All the PEs part of the ES will source traffic using the same **es-bmac**.

In Figure 174 and the following configuration, the **es-bmac** used by PE3 and PE4 will be BMAC-34, i.e. 00:00:00:00:00:34. The **es-bmac** for a specified **ethernet-segment** is configured by the **source-bmac-lsb** along with the **(b-)vpls>pbb>use-es-bmac** command.

Configuration in PE3:

```
*A:PE3>config>lag(1)# info
----------------------------------------------
  mode access
  encap-type dot1q
  port 1/1/1
  lacp active administrative-key 32768
  no shutdown

*A:PE3>config>service>system>bgp-evpn# info
----------------------------------------------
```

```
  route-distinguisher 3.3.3.3:0
  ethernet-segment ESI-1 create
    esi 00:34:34:34:34:34:34:34:34:34
    multi-homing all-active
    service-carving auto
    lag 1
    source-bmac-lsb 00:34 es-bmac-table-size 8
    no shutdown

*A:PE3>config>service>vpls 1(b-vpls)# info
----------------------------------------------
  bgp
  bgp-evpn
    evi 1
    mpls
      no shutdown
      ecmp 2
      auto-bind-tunnel resolution any
  pbb
    source-bmac 00:00:00:00:00:03
    use-es-bmac

*A:PE3>config>service>vpls (i-vpls)# info
----------------------------------------------
  pbb
    backbone-vpls 1
  sap lag-1:101 create

*A:PE1>config>service>epipe (pbb)# info
----------------------------------------------
  pbb
    tunnel 1 backbone-dest-mac 00:00:00:00:00:01 isid 102
  sap lag-1:102 create
```

Configuration in PE4:

```
*A:PE4>config>lag(1)# info
----------------------------------------------
  mode access
  encap-type dot1q
  port 1/1/1
  lacp active administrative-key 32768
  no shutdown

*A:PE4>config>service>system>bgp-evpn# info
----------------------------------------------
  route-distinguisher 4.4.4.4:0
  ethernet-segment ESI-1 create
    esi 00:34:34:34:34:34:34:34:34:34
    multi-homing all-active
    service-carving auto
    lag 1
    source-bmac-lsb 00:34 es-bmac-table-size 8
    no shutdown


*A:PE4>config>service>vpls 1(b-vpls)# info
```

```
-----------------------------------------------
  bgp
  bgp-evpn
    evi 1
    mpls
      no shutdown
      ecmp 2
      auto-bind-tunnel resolution any
  pbb
    source-bmac 00:00:00:00:00:04
    use-es-bmac

*A:PE4>config>service>vpls (i-vpls)# info
-----------------------------------------------
  pbb
    backbone-vpls 1
  sap lag-1:101 create

*A:PE4>config>service>epipe (pbb)# info
-----------------------------------------------
  pbb
    tunnel 1 backbone-dest-mac 00:00:00:00:00:01 isid 102
  sap lag-1:102 create
```

The above configuration will enable the all-active multi-homing procedures for PBB-EVPN.

Note: The **ethernet-segment ESI-1** can also be used for regular VPLS services.

The following considerations apply when the ESI is used for PBB-EVPN.

- **ESI association:** Only LAG is supported for all-active multi-homing. The following commands are used for the LAG to ESI association:
  - **config>service>system>bgp-evpn>ethernet-segment# lag <id>**
  - **config>service>system>bgp-evpn>ethernet-segment# source-bmac-lsb <MAC-lsb> [es-bmac-table-size <size>]**
  - Where:
    - The same ESI may be used for EVPN and PBB-EVPN services.
    - For PBB-EVPN services, the **source-bmac-lsb** attribute is mandatory and ignored for EVPN-MPLS services.
    - The **source-bmac-lsb** attribute must be set to a specific 2-byte value. The value must match on all the PEs part of the same ESI, for example, PE3 and PE4 for ESI-1. This means that the configured **pbb>source-bmac** on the two PEs for B-VPLS 1 must have the same 4 most significant bytes.

- The **es-bmac-table-size** parameter modifies the default value (8) for the maximum number of virtual BMACs that can be associated with the **ethernet-segment**, i.e. **es-bmacs**. When the **source-bmac-lsb** is configured, the associated **es-bmac-table-size** is reserved out of the total FDB space.
- When **multi-homing all-active** is configured within the **ethernet-segment**, only a LAG can be associated with it. The association of a port or an sdp will be restricted by the CLI.

- If **service-carving** is configured in the ESI, the DF election algorithm will be a modulo function of the ISID and the number of PEs part of the ESI, as opposed to a modulo function of evi and number of PEs (used for EVPN-MPLS).
- A **service-carving mode manual** option is added so that the user can control what PE is DF for a specified ISID. The PE will be DF for the configured ISIDs and non-DF for the non-configured ISIDs.
- **DF election:** An all-active Designated Forwarder (DF) election is also carried out for PBB-EVPN. In this case, the DF election defines which of the PEs of the ESI for a specified I-VPLS is the one able to send the downstream BUM traffic to the CE. Only one DF per ESI is allowed in the I-VPLS service, and the non-DF will only block BUM traffic and in the downstream direction.
- **Split-horizon function:** In PBB-EVPN, the split-horizon function to avoid echoed packets on the CE is based on an ingress lookup of the ES BMAC (as opposed to the ESI label in EVPN-MPLS). In Figure 174 PE3 sends packets using BMAC SA = BMAC-34. PE4 does not send those packets back to the CE because BMAC-34 is identified as the **es-bmac** for ESI-1.
- **Aliasing:** In PBB-EVPN, aliasing is based on the ES BMAC sent by all the PEs part of the same ESI. See the following section for more information. In Figure 174 PE1 performs load balancing between PE3 and PE4 when sending unicast flows to BMAC-34 (es-bmac for ESI-1).

In the configuration above, a PBB-Epipe is configured in PE3 and PE4, both pointing at the same remote **pbb tunnel backbone-dest-mac**. On the remote PE, i.e. PE1, the configuration of the PBB-Epipe will point at the **es-bmac**:

```
*A:PE1>config>service>epipe (pbb)# info
----------------------------------------------
  pbb
    tunnel 1 backbone-dest-mac 00:00:00:00:00:34 isid 102
  sap 1/1/1:102 create
```

When PBB-Epipes are used in combination with all-active multi-homing, Nokia recommends using **bgp-evpn mpls ingress-replication-bum-label** in the PEs where the **ethernet-segment** is created, that is in PE3 and PE4. This guarantees that in case of flooding in the B-VPLS service for the PBB Epipe, only the DF will forward the traffic to the CE.

Note: PBB-Epipe traffic always uses BMAC DA = unicast; therefore, the DF cannot check whether the inner frame is unknown unicast or not based on the group BMAC. Therefore, the use of an EVPN BUM label is highly recommended.

Aliasing for PBB-epipes with all-active multi-homing only works if shared-queuing or ingress policing is enabled on the ingress PE epipe. In any other case, the IOM will send the traffic to a single destination (no ECMP will be used in spite of the **bgp-evpn mpls ecmp** setting).

All-active multi-homed **es-bmacs** are treated by the remote PEs as **eES:MAX-ESI BMACs**. The following example shows the FDB in B-VPLS 1 in PE1 as shown in Figure 174:

```
*A:PE1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:00:00:00:00:03 eMpls:                   EvpnS    06/12/15 15:35:39
                            192.0.2.3:262138
1         00:00:00:00:00:04 eMpls:                   EvpnS    06/12/15 15:42:52
                            192.0.2.4:262130
1         00:00:00:00:00:34 eES:                     EvpnS    06/12/15 15:35:57
                            MAX-ESI
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

The **show service id evpn-mpls** on PE1 shows that the remote **es-bmac**, i.e. 00:00:00:00:00:34, has two associated next-hops, i.e. PE3 and PE4:

```
*A:PE1# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                 Transport
-------------------------------------------------------------------------------
192.0.2.3       262138        1           Yes             06/12/2015 15:34:48
                ldp
192.0.2.4       262130        1           Yes             06/12/2015 15:34:48
                ldp
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

```
===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                     TEP Address     Egr Label   Last Change
                                               Transport
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN-MPLS ES BMAC Dest
===============================================================================
VBMacAddr                   TEP Address     Egr Label     Last Change
                                             Transport
-------------------------------------------------------------------------------
00:00:00:00:00:34           192.0.2.3       262138        06/12/2015 15:34:48
                                            ldp
00:00:00:00:00:34           192.0.2.4       262130        06/12/2015 15:34:48
                                            ldp
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Network failures and convergence for all-active multi-homing

ES failures are resolved by the PEs withdrawing the **es-bmac**. The remote PEs will withdraw the route and update their list of next-hops for a specified **es-bmac**.

No mac-flush of the I-VPLS FDB tables is required as long as the **es-bmac** is still in the FDB.

When the route corresponding to the last next-hop for a specified **es-bmac** is withdrawn, the es-bmac will be flushed from the B-VPLS FDB and all the CMACs associated with it will be flushed too.

The following events will trigger a withdrawal of the **es-bmac** and the corresponding next-hop update in the remote PEs:

- B-VPLS transition to operationally down status.
- Change of **pbb>source-bmac**.
- Change of **es-bmac** (or removal of **pbb use-es-bmac**).
- Ethernet-segment transition to operationally down status.

Note: Individual saps going operationally down in an ES will not generate any BGP withdrawal or indication so that the remote nodes can flush their CMACs. This is solved in EVPN-MPLS by the use of AD routes per EVI; however, there is nothing similar in PBB-EVPN for indicating a partial failure in an ESI.

PBB-EVPN Single-Active Multi-Homing Service Model

In single-active multi-homing, the non-DF PEs for a specified ESI will block unicast and BUM traffic in both directions (upstream and downstream) on the object associated with the ESI. Other than that, single-active multi-homing will follow the same service model defined in the section 'PBB-EVPN all-active multi-homing service model' with the following differences:

- The **ethernet-segment** will be configured for **single-active**: **service>system>bgp-evpn>ethernet-segment>multi-homing single-active**.
- For single-active multi-homing, the **ethernet-segment** can be associated with a port and sdp, as well as a **lag**.
- From a service perspective, single-active multi-homing can provide redundancy to the following services and access types:
  - I-VPLS LAG and regular SAPs
  - I-VPLS active/standby spoke-sdps
  - EVPN single-active multi-homing is supported for PBB-Epipes only in two-node scenarios with local switching.
- While all-active multi-homing only uses **es-bmac** assignment to the ES, single-active multi-homing can use source-bmac or **es-bmac** assignment. The system allows the following user choices per B-VPLS and ES:
  - A dedicated **es-bmac** per ES can be used. In that case, the **pbb>use-es-bmac** command will be configured in the B-VPLS and the same procedures explained in PBB-EVPN all-active multi-homing service model will follow with one difference. While in all-active multi-homing all the PEs part of the ESI will source the PBB packets with the same source es-bmac, single-active multi-homing requires the use of a different **es-bmac** per PE.
  - A non-dedicated **source-bmac** can be used. In this case, the user will not configure **pbb>use-es-bmac** and the regular **source-bmac** will be used for the traffic. A different **source-bmac** has to be advertised per PE.
  - The use of **source-bmacs** or **es-bmacs** for single-active multi-homed ESIs has a different impact on CMAC flushing, as shown in Figure 175.

Figure 175 **Source-Bmac Versus Es-Bmac CMAC Flushing**

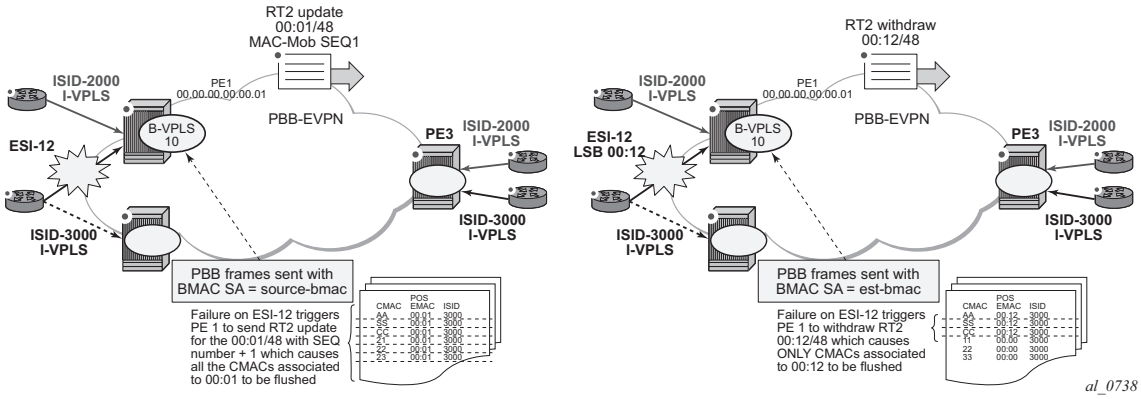

– If **es-bmacs** are used as shown in the representation on the right in Figure 175, a less-impacting CMAC flush is achieved, therefore, minimizing the flooding after ESI failures. In case of ESI failure, PE1 will withdraw the **es-bmac** 00:12 and the remote PE3 will only flush the CMACs associated with that **es-bmac** (only the CMACs behind the CE are flushed).

– If **source-bmacs** are used, as shown on the left-hand side of Figure Figure 175, in case of ES failure, a BGP update with higher sequence number will be issued by PE1 and the remote PE3 will flush all the CMACs associated with the **source-bmac**. Therefore, all the CMACs behind the PE's B-VPLS will be flushed, as opposed to only the CMACs behind the ESI's CE.

- As in EVPN-MPLS, the non-DF status can be notified to the access CE or network:
  – LAG with or without LACP: In this case, the multi-homed ports on the CE will not be part of the same LAG. The non-DF PE for each service may signal that the LAG sap is operationally down by using **eth-cfm fault-propagation-enable {use-if-tlv|suspend-ccm}**.
  – Regular Ethernet 802.1q/ad ports: In this case, the multi-homed ports on the CE/network will not be part of any LAG. The non-DF PE for each service will signal that the sap is operationally down by using **eth-cfm fault-propagation-enable {use-if-tlv|suspend-ccm}**.
  – Active-standby PWs: in this case, the multihomed CE/network is connected to the PEs through an MPLS network and an active/standby spoke-sdp per service. The non-DF PE for each service will make use of the LDP PW status bits to signal that the spoke-sdp is standby at the PE side. Nokia recommends that the CE suppresses the signaling of PW status standby.

Network Failures and Convergence for Single-Active Multihoming

ESI failures are resolved depending on the BMAC address assignment chosen by the user:

- If the BMAC address assignment is based on the use of **es-bmacs**, DF and non-DFs will send the **es-bmac/ESI=0** for a specified ESI. Each PE will have a different **es-bmac** for the same ESI (as opposed to the same **es-bmac** on all the PEs for all-active). In case of an ESI failure on a PE:
  - The PE will withdraw its **es-bmac** route triggering a mac-flush of all the CMACs associated with it in the remote PEs.
- If the BMAC address assignment is based on the use of **source-bmac**, DF and non-DFs will advertise their respective **source-bmacs**. In case of an ES failure:
  - The PE will re-advertise its **source-bmac** with a higher sequence number (the new DF will not re-advertise its s**ource-bmac**).
  - The far-end PEs will interpret a **source-bmac** advertisement with a different sequence number as a flush-all-from-me message from the PE detecting the failure. They will flush all the CMACs associated with that BMAC in all the ISID services.

The following events will trigger a CMAC flush notification. A 'CMAC flush notification' means the withdrawal of a specified BMAC or the update of BMAC with a higher sequence number (SQN). Both BGP messages will make the remote PEs flush all the CMACs associated with the indicated BMAC:

- B-VPLS transition to operationally down status. This will trigger the withdrawal of the associated BMACs, regardless of the **use-es-bmac** setting.
- Change of **pbb>source-bmac**. This will trigger the withdrawal and re-advertisement of the **source-bmac**, causing the corresponding CMAC flush in the remote PEs.
- Change of **es-bmac** (removal of **pbb use-es-bmac**). This will trigger the withdrawal of the **es-bmac** and re-advertisement of the new **es-bmac**.
- Ethernet-Segment (ES) transition to operationally down or admin-down status. This will trigger an **es-bmac** withdrawal (if **use-es-bmac** is used) or an update of the source-bmac with a higher SQN (if n**o use-es-bmac** is used).
- Service Carving Range change for the ES. This will trigger an **es-bmac** update with higher SQN (if **use-es-bmac** is used) or an update of the source-bmac with a higher SQN (if **no use-es-bmac** is used).
- Change in the number of candidate PEs for the ES. This will trigger an **es-bmac** update with higher SQN (if **use-es-bmac** is used) or an update of the source-bmac with a higher SQN (if **no use-es-bmac** is used).

- In an ESI, individual saps/sdp-bindings or individual I-VPLS going operationally down will not generate any BGP withdrawal or indication so that the remote nodes can flush their CMACs. This is solved in EVPN-MPLS by the use of AD routes per EVI; however, there is nothing similar in PBB-EVPN for indicating a partial failure in an ESI.

PBB-Epipes and EVPN Multi-Homing

EVPN multi-homing is supported with PBB-EVPN Epipes, but only in a limited number of scenarios. In general, the following applies to PBB-EVPN Epipes:

- PBB-EVPN Epipes don't support spoke-sdps that are associated with EVPN Ethernet Segments (ES).
- PBB-EVPN Epipes support all-active EVPN multi-homing as long as no local-switching is required in the Epipe instance where the ES is defined.
- PBB-EVPN Epipes support single-active EVPN multi-homing only in a two-node case scenario.

Figure 176 shows the EVPN MH support in a three-node scenario.

Figure 176 **PBB-EVPN MH in a Three-Node Scenario**

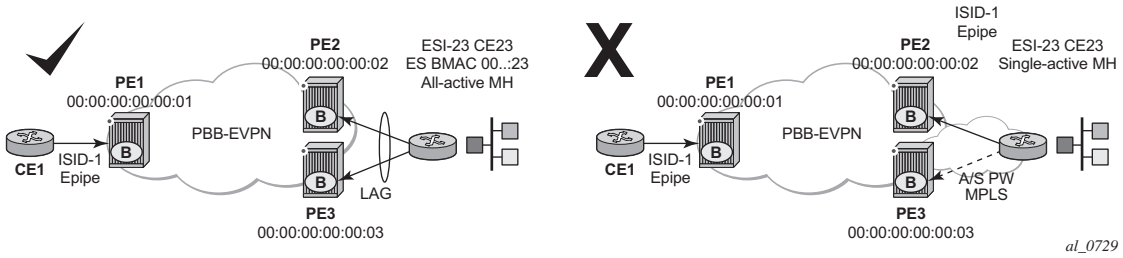

EVPN MH support in a three-node scenario has the following characteristics:

- All-active EVPN multi-homing is fully supported (diagram on the left in Figure 176). CE1 might also be multi-homed to other PEs, as long as those PEs are not PE2 or PE3. In this case, PE1 Epipe's **pbb-tunnel** would be configured with the remote ES BMAC.
- Single-active EVPN multi-homing is not supported in a three (or more)-node scenario (diagram on the right in Figure 176). Since PE1's Epipe **pbb-tunnel** can only point at a single remote BMAC and single-active multi-homing requires the use of separate BMACs on PE2 and PE3, the scenario is not possible and not supported regardless of the ES association to port/LAG/sdps.
- Regardless of the EVPN multi-homing type, the CLI prevents the user from adding a spoke-sdp to an Epipe, if the corresponding SDP is part of an ES.

Figure 177 shows the EVPN MH support in a two-node scenario.

Figure 177 **PBB-EVPN MH in a Two-Node Scenario**

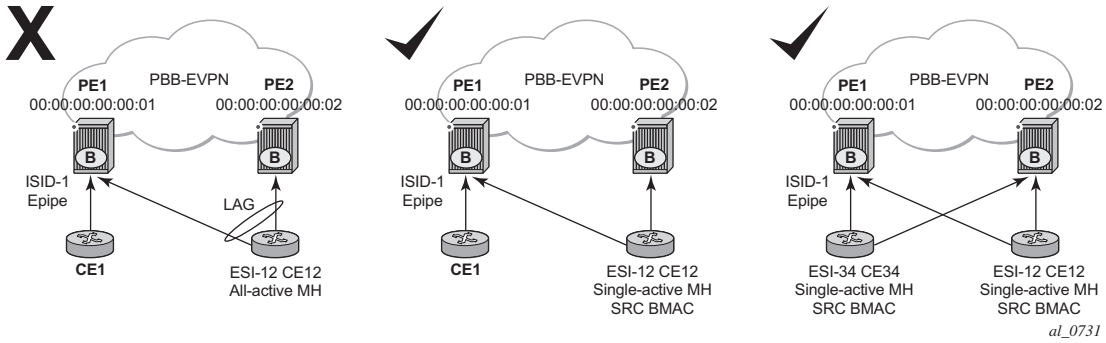

EVPN MH support in a two-node scenario has the following characteristics, as shown in Figure 177:

- All-active multi-homing is not supported for redundancy in this scenario because PE1's **pbb-tunnel** cannot point at a locally defined ES-BMAC. This is represented in the left-most scenario in Figure 177.
- Single-active multi-homing is supported for redundancy in a two-node three or four SAP scenario, as displayed by the two right-most scenarios in Figure 177).

  In these two cases, the Epipe **pbb-tunnel** will be configured with the source BMAC of the remote PE node.

  When two saps are active in the same Epipe, local-switching is used to exchange frames between the CEs.

5.3.6.2.4 PBB-EVPN and Use of P2MP mLDP Tunnels for Default Multicast List

P2MP mLDP tunnels can also be used in PBB-EVPN services. The use of provider-tunnel inclusive MLDP is only supported in the B-VPLS default multicast list; that is, no per-ISID IMET-P2MP routes are supported. IMET-P2MP routes in a B-VPLS are always advertised with Eth-tag zero. All-active EVPN multihoming is supported in PBB-EVPN services together with P2MP mLDP tunnels; however, single-active multihoming is not supported. This capability is only required on the P2MP root PEs within PBB-EVPN services using all-active multihoming.

B-VPLS supports the use of MFIBs for ISIDs using ingress replication. The following considerations apply when **provider-tunnel** is enabled in a B-VPLS service.

- Local I-VPLS or static-ISIDs configured on the B-VPLS will generate IMET-IR routes and MFIBs will be created per ISID by default.

- The default IMET-P2MP or IMET-P2MP-IR route sent with ethernet-tag = 0 will be issued depending on the **ingress-repl-inc-mcast-advertisement** command.
- The following considerations apply if an **isid-policy** is configured in the B-VPLS.
  - A range of ISIDs configured with **use-def-mcast** will make use of the P2MP tree, assuming the node is configured as **root-and-leaf**.
  - A range of ISIDs configured with **advertise-local** will make the system advertise IMET-IR routes for the local ISIDs included in the range.

The following example CLI output shows a range of ISIDs (1000-2000) that use the P2MP tree and the system does not advertise the IMET-IR routes for those ISIDs. Other local ISIDs will advertise the IMET-IR and will use the MFIB to forward BUM packets to the EVPN-MPLS destinations created by the IMET-IR routes.

```
*A:PE-1>config>service>vpls(b-vpls)# info
----------------------------------------------
  service-mtu 2000
  bgp-evpn
    evi 10
    mpls
      no shutdown
      auto-bind-tunnel resolution any
  isid-policy
    entry 10 create
      use-def-mcast
      no advertise-local
      range 1000 to 2000
      exit
    exit
  provider-tunnel
    inclusive
      owner bgp-evpn-mpls
      root-and-leaf
      mldp
      no shutdown
      exit
    exit
  sap 1/1/1:1 create
  exit
  spoke-sdp 1:1 create
 exit
```

5.3.6.2.5 PBB-EVPN ISID-Based CMAC-Flush

SR OS supports ISID-based CMAC-flush procedures for PBB-EVPN I-VPLS services where no single-active Ethernet segments are used. SR OS also supports CMAC-flush procedure where other redundancy mechanisms, such as BGP-MH, need these procedures to avoid black-holes caused by a SAP or spoke-SDP failure.

The CMAC-flush procedures are enabled on the I-VPLS service using the **config**>**service**>**vpls**>**pbb**>**send-bvpls-evpn-flush** CLI command. The feature can be disabled on a per-SAP or per-spoke-SDP basis by using the **disable-send-bvpls-evpn-flush** command in the **config**>**service**>**vpls**>**sap** or **config**>**service**>**vpls**>**spoke-sdp** context.

With the feature enabled on an I-VPLS service and a SAP or spoke-SDP, if there is a SAP or spoke-SDP failure, the router sends a CMAC-flush notification for the corresponding BMAC and ISID. The router receiving the notification flushes all the CMACs associated with the indicated BMAC and ISID when the **config**>**service**>**vpls**>**bgp-evpn**>**accept-ivpls-evpn-flush** command is enabled for the B-VPLS service.

The CMAC-flush notification consists of an EVPN BMAC route that is encoded as follows: the ISID to be flushed is encoded in the Ethernet Tag field and the sequence number is incremented with respect to the previously advertised route.

If **send-bvpls-evpn-flush** is configured on an I-VPLS with SAPs or spoke-SDPs, one of the following rules must be observed.

a. The **disable-send-bvpls-evpn-flush** option is configured on the SAPs or spoke-SDPs.
b. The SAPs or spoke-SDPs are not on an ES.
c. The SAPs or spoke-SDPs are on an ES or vES with **no src-bmac-lsb** enabled.
d. The **no use-es-bmac** is enabled on the B-VPLS.

ISID-based CMAC-flush can be enabled in I-VPLS services with ES or vES. If enabled, the expected interaction between the RFC 7623-based CMAC-flush and the ISID-based CMAC-flush is as follows.

- If **send-bvpls-evpn-flush** is enabled in an I-VPLS service, the ISID-based CMAC-flush overrides (replaces) the RFC 7623-based CMAC-flushing.
- For each ES, vES, or B-VPLS, the system checks for at least one I-VPLS service that does not have **send-bvpls-evpn-flush** enabled.
  - If ISID-based CMAC-flush is enabled for all I-VPLS services, RFC 7623-based CMAC-flushing is not triggered; only ISID-based CMAC-flush notifications are generated.
  - If at least one I-VPLS service is found with no ISID-based CMAC-flush enabled, then RFC 7623-based CMAC-flushing notifications are triggered based on ES events.

    ISID-based CMAC-flush notifications are also generated for I-VPLS services that have **send-bvpls-evpn-flush** enabled.

Figure 178 shows an example where the ISID-based CMAC-flush prevents black-hole situations for a CE that is using BGP-MH as the redundancy mechanism in the I-VPLS with an ISID of 3000.

Figure 178 **Per-ISID CMAC-Flush Following a SAP Failure**

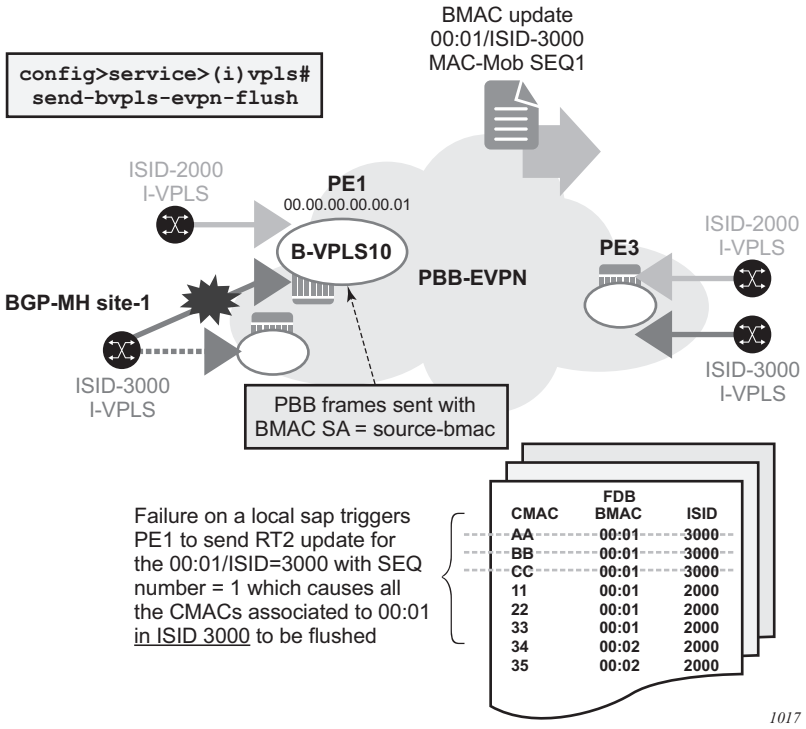

When **send-bvpls-evpn-flush** is enabled, the I-VPLS service is ready to send per-ISID CMAC-flush messages in the form of BMAC/ISID routes. The first BMAC/ISID route for an I-VPLS service is sent with sequence number zero; subsequent updates for the same route increment the sequence number. A BMAC/ISID route for the I-VPLS is advertised or withdrawn in the following cases:

- during I-VPLS **send-bvpls-evpn-flush** configuration and deconfiguration
- during I-VPLS association and disassociation from the B-VPLS service
- during I-VPLS operational status change (up/down)
- during B-VPLS operational status change (up/down)
- during B-VPLS **bgp-evpn mpls** status change (no shutdown/shutdown)
- during B-VPLS operational source BMAC change

- if **no disable-send-bvpls-evpn-flush** is configured for a SAP or spoke-SDP, upon a failure on that SAP or spoke-SDP, the system will send a per-ISID CMAC-flush message; that is, a BMAC/ISID route update with an incremented sequence number

If the user explicitly configures **disable-send-bvpls-evpn-flush** for a SAP or spoke-SDP, the system will not send per-ISID CMAC-flush messages for failures on that SAP or spoke-SDP.

The B-VPLS on the receiving node must be configured with **bgp-evpn**>**accept-ivpls-evpn-flush** to accept and process CMAC-flush non-zero Ethernet-tag MAC routes. If the **accept-ivpls-evpn-flush** command is enabled (the command is disabled by default), the node accepts non-zero Ethernet-tag MAC routes (BMAC/ISID routes) and processes them. When a new BMAC/ISID update (with an incremented sequence number) for an existing route is received, the router will flush all the CMACs associated with that BMAC and ISID. The BMAC/ISID route withdrawals will also cause a CMAC-flush.

Note: Only BMAC routes with the Ethernet Tag field set to zero are considered for BMAC installation in the FDB.

The following CLI example shows the commands that enable the CMAC-flush feature on PE1 and PE3.

```
*A:PE-1>config>service>vpls(i-vpls)# info
----------------------------------------------
  pbb
    backbone-vpls 10
    send-bvpls-evpn-flush
    exit
  exit
  bgp
    route-distinguisher 65000:1
    vsi-export "vsi_export"
    vsi-import "vsi_import"
  exit
  site "CE-1" create
    site-id 1
    sap lag-1:1
    site-activation-timer 3
    no shutdown
  exit
  sap lag-1:1 create
    no disable-send-bvpls-evpn-flush
    no shutdown
    exit
<snip>
*A:PE-3>config>service>vpls(b-vpls 10)# info
----------------------------------------------
<snip>
```

```
bgp-evpn
  accept-ivpls-evpn-flush
```

In the preceding example, with **send-bvpls-evpn-flush** enabled on the I-VPLS service of PE1, a BMAC/ISID route (for pbb source-bmac address BMAC 00:..:01 and ISID 3000) is advertised. If the SAP goes operationally down, PE1 will send an update of the source BMAC address (00:..:01) for ISID 3000 with a higher sequence number.

With **accept-ivpls-evpn-flush** enabled on PE3's B-VPLS service, PE3 flushes all CMACs associated with BMAC 00:01 and ISID 3000. The CMACs associated with other BMACs or ISIDs are retained in PE3's FDB.

5.3.6.2.6 PBB-EVPN ISID-based Route Targets

Routers with PBB-EVPN services use the following route types to advertise the ISID of a specific service.

- Inclusive Multicast Ethernet Tag routes (IMET-ISID routes) are used to auto-discover ISIDs in the PBB-EVPN network. The routes encode the service ISID in the Ethernet Tag field.
- BMAC-ISID routes are only used when ISID-based CMAC-flush is configured. The routes encode the ISID in the Ethernet Tag field.

Although the preceding routes are only relevant for routers where the advertised ISID is configured, they are sent with the B-VPLS route-target by default. As a result, the routes are unnecessarily disseminated to all the routers in the B-VPLS network.

SR OS supports the use of per-ISID or group of ISID route-targets, which limits the dissemination of IMET-ISID or BMAC-ISID routes for a specific ISID to the PEs where the ISID is configured.

The **config**>**service**>**(b-)vpls**>**isid-route-target**>**isid-range** *from* [**to** *to*] [**auto-rt** | **route-target** *rt*] command allows the user to determine whether the IMET-ISID and BMAC-ISID routes are sent with the B-VPLS route-target (default option, **no** command), or a route-target specific to the ISID or range of ISIDs.

The following configuration example shows how to configure ISID ranges as **auto-rt** or with a specific **route-target**.

```
*A:PE-3>config>service>(b-)vpls>bgp-evpn#
isid-route-target
[no] isid-range <from> [to <to>] {auto-rt|route-target <rt>}
/* For example:
*A:PE-3>config>service>(b-)vpls>bgp-evpn#
isid-route-target
 isid-range 1000 to 1999 auto-rt
```

```
isid-range 2000 route-target target:65000:2000
```

The **auto-rt** option auto-derives a route-target per ISID in the following format:

<2-byte-as-number>:<4-byte-value>

Where: 4-byte-value = 0x30+ISID, as described in RFC 7623. Figure 179 shows the format of the **auto-rt** option.

Figure 179 **PBB-EVPN auto-rt ISID-based Route Target Format**

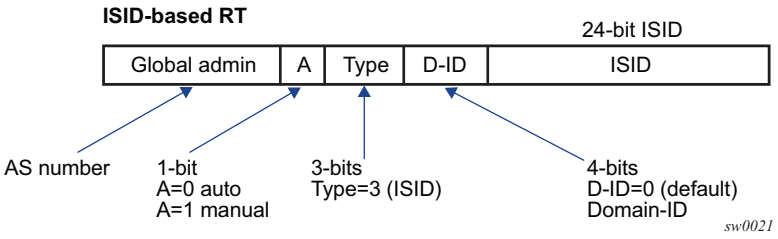

Where:

- If it is 2 bytes, then the AS number is obtained from the **config>router>autonomous-system** command. If the AS number exceeds the 2 byte limit, then the low order 16-bit value is used.
- A = 0 for auto-derivation
- Type = 3, which corresponds to an ISID-type route-target
- ISID is the 24-bit ISID
- The type and sub-type are 0x00 and 0x02.

If **isid-route-target** is enabled, the export and import directions for IMET-ISID and BMAC-ISID route processing are modified as follows.

- Exported IMET-ISID and BMAC-ISID routes
  - For local I-VPLS ISIDs and static ISIDs, IMET-ISID routes are sent individually with an ISID-based route-target (and without a B-VPLS route-target) unless the ISID is contained in an ISID policy for which **no advertise-local** is configured.
  - If both **isid-route-target** and **send-bvpls-evpn-flush** options are enabled for an I-VPLS, the BMAC-ISID route is also sent with the ISID-based route-target and no B-VPLS route-target.
  - The **isid-route-target** command affects the IMET-ISID and BMAC-ISID routes only. The BMAC-0, IMET-0 (BMAC and IMET routes with Ethernet Tag == 0), and ES routes are not impacted by the command.
- Imported IMET-ISID and BMAC-ISID routes

- Upon enabling **isid-route-target** for a specific I-VPLS, the BGP starts importing IMET-ISID routes with ISID-based route-targets, and (assuming the **bgp-evpn accept-ivpls-evpn-flush** option is enabled) BMAC-ISID routes with ISID-based route-targets.
- The new ISID-based RTs are added for import operations when the I-VPLS is associated to the B-VPLS service (and not based on the I-VPLS operational status), or when the **static-isid** is added.
- The system does not maintain a mapping of the route-targets and ISIDs for the imported routes. For example, if I-VPLS 1 and 2 are configured with the **isid-route-target** option and IMET-ISID=2 route is received with a **route-target** corresponding to ISID=1, then BGP imports the route and the router processes it.
- The router does not check the format of the received auto-derived route-targets. The route is imported as long as the route-target is on the list of RTs for the B-VPLS.

• If the **isid-route-target** option is configured for one or more I-VPLS services, the **vsi-import** and **vsi-export** policies are blocked in the B-VPLS. BGP peer import and export policies are still allowed. Matching on the export ISID-based route-target is supported.

5.3.7 Virtual Ethernet Segments

SR OS supports virtual Ethernet Segments (vES) for EVPN multi-homing in accordance with *draft-sajassi-bess-evpn-virtual-eth-segment*.

Regular Ethernet segments can only be associated to ports, LAGs, and SDPs, which satisfies the redundancy requirements for CEs that are directly connected to the ES PEs by a port, LAG, or SDP. However, this implementation does not work when an aggregation network exists between the CEs and the ES PEs, which requires different ESs to be defined for the port or LAG of the SDP.

Figure 180 shows an example of how CE1 and CE2 use all-active multi-homing to the EVPN-MPLS network despite the third-party Ethernet aggregation network to which they are connected.

Figure 180 **All-Active Multi-Homing on vES**

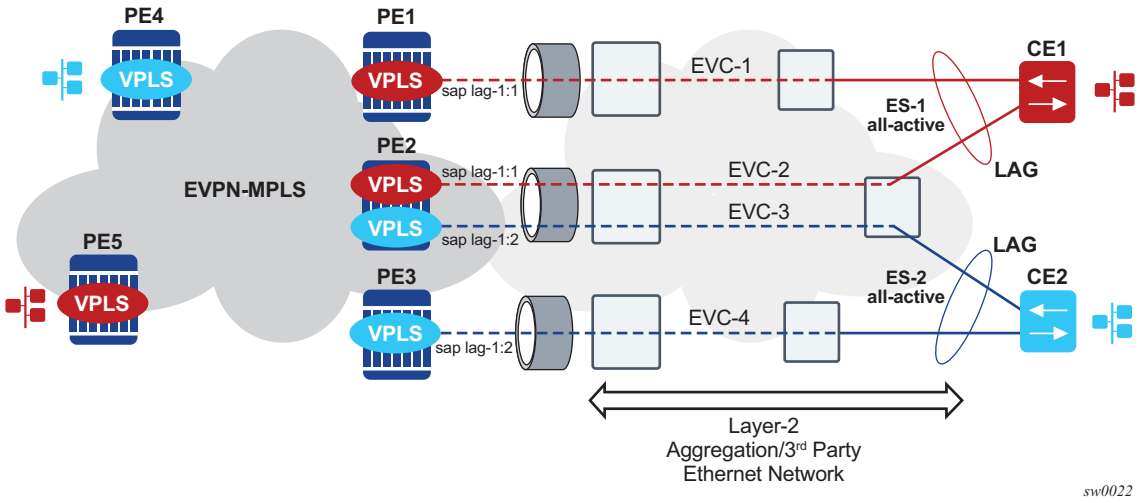

The ES association can be made in a more granular way by creating a vES. A vES can be associated to the following:

- Qtag-ranges on dot1q ports or LAGs
- S-tag-ranges on qinq ports or LAGs
- C-tag-ranges per s-tag on qinq ports or LAGs
- VC-ID ranges on SDPs

The following CLI examples show the vES configuration options:

```
config>service>system>bgp-evpn#
...
ethernet-segment vES-1 virtual create
 lag 1
 dot1q
   qtag-range 100 to 200
...
ethernet-segment vES-2 virtual create
 port 1/1/1
 qinq
   s-tag 1 c-tag-range 100 to 200
   s-tag-range 2 to 10
...
ethernet-segment vES-3 virtual create
 sdp 1
 vc-id-range 1000 to 2000
...
```

Where:

- The **virtual** keyword creates an ES as defined in *draft-sajassi-bess-evpn-virtual-eth-segment*. The configuration of the dot1q or qinq nodes is allowed only when the ES is created as **virtual**.
- On the vES, the user must first create a port, LAG, or SDP before configuring a VLAN or VC-ID association. When added, the port/LAG type and encap-type is checked as follows:
  - If the encap-type is dot1q, only the dot1q context configuration is allowed; the qinq context cannot be configured.
  - If the encap-type is qinq, only the qinq node is allowed; the dot1q context cannot be configured.
  - A dot1q, qinq, or vc-id range is required for the vES to become operationally active.
- The **dot1q qtag-range** <*qtag1*> [**to** *qtag1*] command determines which VIDs are associated with the vES on a specific dot1q port or LAG. The group of SAPs that match the configured port/LAG and VIDs will be part of the vES.
- The **qinq s-tag-range** <*qtag1*> [**to** *qtag1*] command determines which outer VIDs are associated with the vES on the qinq port or LAG.
- The **qinq s-tag** <*qtag1*> **c-tag-range** <*qtag2*> [**to** <*qtag2*] command determines which inner c-tags per s-tag is associated with the vES on the qinq port or LAG.
- The **vc-id range** <*vcid*> **[to** *vc-id*] command determines which VC ids are associated with the vES on the configured SDP.

Although qtag values 0, * and 1 to 4094 are allowed, the following considerations must be taken in to account when configuring a dot1q or qinq vES:

- Up to 8 dot1q or qinq ranges may be configured in the same vES.
- When configuring a qinq vES, a qtag included in a s-tag-range cannot be included in the s-tag qtag of the **s-tag** *qtag1* **c-tag-range** *qtag2* **[to** *qtag2*] command. For example, the following combination is not supported in the same vES:

```
s-tag-range 350 to 500
s-tag 500 c-tag-range 100 to 200
```

The following example shows a supported combination:

```
*A:PE75>config>service>system>bgp-evpn>eth-seg>qinq# info
----------------------------------------------
                s-tag-range 100 to 200
                s-tag-range 300 to 400
                s-tag 500 c-tag-range 100 to 200
                s-tag 600 c-tag-range 100 to 200
                s-tag 600 c-tag-range 150 to 200
```

- vES associations that contain qtags <0, *, null> are special and treated as follows:

– When a special qtag value is configured in the *from* value of the range, the *to* value must be the same.
– Qtag values <0, *> are only supported for the **qtag-range** and **c-tag-range**; they are not supported in the **s-tag-range**.
– The qtag “null” value is only supported in the **c-tag-range** if the **s-tag** is configured as “*”.

Table 82 lists examples of the supported qtag values between 1 to 4094.

Table 82 **Examples of Supported qtag Values**

vES Configuration for Port 1/1/1	SAP Association
dot1q qtag-range 100	1/1/1:100
dot1q qtag-range 100 to 102	1/1/1:100, 1/1/1:101, 1/1/1:102
qinq s-tag 100 c-tag-range 200	1/1/1:100.200
qinq s-tag-range 100	All the SAPs 1/1/1:100.*x* where: *x* is a value between 1 to 4094, 0, *
qinq s-tag-range 100 to 102	All SAPs 1/1/1:100.x, 1/1/1:101.x, 1/1/1:102.*x* where: *x* is a value between 1 to 4094, 0, *

Table 83 lists all the supported combinations that include qtag values <0, *, null>. Any other combination of these special values is not supported.

Table 83 **Examples of Supported Combinations**

vES Configuration for Port 1/1/1	SAP Association
dot1q qtag-range 0	1/1/1:0
dot1q qtag-range *	1/1/1:*
qinq s-tag 0 c-tag-range *	1/1/1:0.*
qinq s-tag * c-tag-range *	1/1/1:*.*
qinq s-tag * c-tag-range null	1/1/1:*.null
qinq s-tag x c-tag-range 0	1/1/1:x.0 where: *x* is a value between 1 to 4094
qinq s-tag x c-tag-range *	1/1/1:x.* where: *x* is a value between 1 to 4094

On vESs, the single-active and all-active modes are supported for EVPN-MPLS VPLS, Epipe, and PBB-EVPN services. Single-active multi-homing is supported on port and SDP-based vESs, and all-active multi-homing is only supported on LAG-based vESs.

The following considerations apply if the vES is used with PBB-EVPN services:

- BMAC allocation procedures are the same as the regular ES procedures.

  Note: Two all-active vESs must use different ES-BMACs, even if they are defined in the same LAG
- The vES implements CMAC flush procedures described in RFC 7623. Optionally, the ISID-based CMAC flush can be used for cases where the single-active vES does not use ES-BMAC allocation.

5.3.8 Preference-Based and Non-Revertive DF Election

In addition to the ES service-carving modes **auto** and **off**, the **manual** mode also supports the preference-based algorithm with the **non-revertive** option, as described in *draft-rabadan-bess-evpn-pref-df*.

When ES is configured to use the preference-based algorithm, the ES route is advertised with the Designated Forwarder (DF) election extended community (sub-type 0x06). Figure 181 shows the DF election extended community.

Figure 181 **DF Election Extended Community**

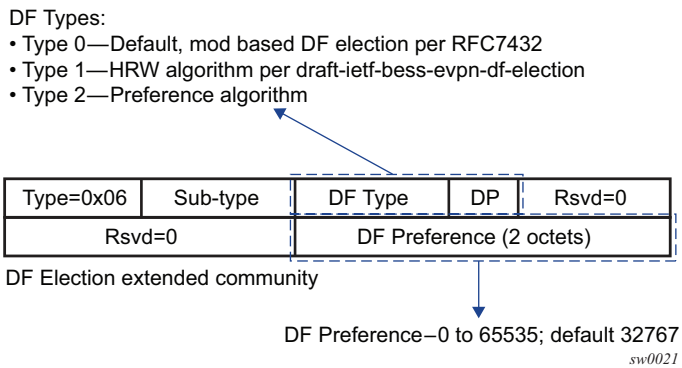

In the extended community, a DF type 2 preference algorithm is advertised with a 2-byte preference value (32767 by default) if the preference-based **manual** mode is configured. The Don't Preempt Me (DP) bit is set if the **non-revertive** option is enabled.

The following CLI excerpt shows the relevant commands to enable the preference-based DF election on a specific ES (regular or virtual):

```
config>service>system>bgp-evpn>ethernet-segment#
...
service-carving mode {manual|auto|off}
  service-carving manual
    [no] preference [create] [non-revertive]
      value <value>
    exit
  [no] evi <evi> [to <evi>]
  [no] isid <isid> [to <isid>]
# value 0..65535; default 32767
...
```

Where:

- The preference value can be changed on an active ES without shutting down the ES, and therefore, a new DF can be forced for maintenance or other reasons.
- The **service-carving** mode must be changed to **manual** mode to create the **preference** context.
- The **preference** command is supported on regular or virtual ES, regardless of the multi-homing mode (single-active or all-active) or the service type (VPLS, I-VPLS, or Epipe).
- By default, the highest-preference PE in the ES becomes the DF for an EVI or ISID, using the DP bit as the tiebreaker first (DP=1 wins over DP=0) and the lowest PE-IP as the last-resort tiebreaker. All the explicitly configured EVI or ISID ranges select the lowest preference PE as the DF (whereas the non-configured EVI or ISID values select the highest preference PE).

  This selection is displayed as **Cfg Range Type: lowest-pref** in the following **show** command example.

```
*A:PE-2# show service system bgp-evpn ethernet-segment name "vES-23"
===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : vES-23
Eth Seg Type            : Virtual
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:23:23:23:23:23:23:23:23:23
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 262141
Source BMAC LSB         : 00-23
ES BMac Tbl Size        : 8                  ES BMac Entries    : 0
Lag Id                  : 1
ES Activation Timer     : 3 secs (default)
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref
-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
```

```
Mode           Value                                    Value          Preempt
-------------------------------------------------------------------------------
non-revertive  100            12/21/2016 15:16:38      100            Enabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
===============================================================================
```

- The EVI and ISID ranges configured on the service-carving context are not required to be consistent with any ranges configured for vESs.
- If the **non-revertive** option is configured, when the former DF comes back up after a failure and checks existing ES routes, it will advertise an operational preference and DP bit, which does not cause a DF switchover for the ES EVI/ISID values.

The following configuration example shows the use of the preference-based algorithm and non-revertive option in an ES defined in PE1 and PE2.

```
*A:PE-1>config>service>system>bgp-evpn# info
----------------------------------------------
 ethernet-segment "ES1" create
   esi 01:00:00:00:00:12:00:00:00:01
   service-carving manual
       preference non-revertive create
         value 10000
       exit
       evi 2001 to 4000
   exit
   multi-homing single-active
   port 1/1/1
   no shutdown

/* example of vpls 1 - similar config exists for evis 2-4000 */
*A:PE-1>config>service>vpls# info
----------------------------------------------
 vxlan vni 1 create
 exit
 bgp-evpn
   evi 1
   mpls
     ecmp 2
     auto-bind-tunnel
       resolution any
     exit
 sap 1/1/1:1 create
 no shutdown
----------------------------------------------
*A:PE-2>config>service>system>bgp-evpn# info
----------------------------------------------
 ethernet-segment "ES1" create
   esi 01:00:00:00:00:12:00:00:00:01
   service-carving manual
       preference non-revertive create
         value 5000
       exit
       evi 2001 to 4000
```

```
    exit
    multi-homing single-active
    port 1/1/1
    no shutdown

*A:PE-2>config>service>vpls# info
----------------------------------------------
 vxlan vni 1 create
 exit
 bgp-evpn
   evi 1
   mpls
     ecmp 2
     auto-bind-tunnel
       resolution any
     exit
 sap 1/1/1:1 create
 no shutdown
----------------------------------------------
```

Based on the configuration in the preceding example, the PE behavior is as follows:

1. Assuming the ES is **no shutdown** on both PE1 and PE2, the PEs exchange ES routes, including the [Pref, DP-bit] in the DF election extended community.
2. For EVIs 1 to 2000, PE2 is immediately promoted to NDF, whereas PE1 becomes the DF, and (following the **es-activation-timer**) brings up its SAP in EVIs 1 to 2000.

   For EVIs 2001 to 4000, the result is the opposite and PE2 becomes the DF.
3. If port 1/1/1 on PE1 goes down, PE1 withdraws its ES route and PE2 becomes the DF for EVIs 1 to 2000.
4. When port 1/1/1 on PE1 comes back up, PE1 compares its ES1 preference with the preferences in the remote PEs in ES1. PE1 advertises the ES route with an “in-use operational” Pref = 5000 and DP=0. Because PE2's Pref is the same as PE1's operational value, but PE2's DP=1, PE2 continues to be the DF for EVIs 1 to 4000.

   Note: The DP bit is the tiebreaker in case of equal Pref and regardless of the choice of highest or lowest preference algorithm.
5. PE1's “in-use” Pref and DP will continue to be [5000,0] until one of the following conditions is true:
   a. PE2 withdraws its ES route, in which case PE1 will re-advertise its admin Pref and DP [10000,DP=1]
   b. The user changes PE1's Pref configuration

5.3.9 IGMP Snooping in EVPN-MPLS and PBB EVPN Services

IGMP snooping is supported in EVPN-MPLS VPLS and PBB-EVPN I-VPLS (where BGP EVPN is running in the associated B-VPLS service) services. It is required in scenarios where the operator does not want to flood all the IP multicast traffic to the access nodes or CEs, and only wants to deliver the IP multicast traffic for which IGMP reports have been received.

The following points apply when IGMP snooping is configured in EVPN-MPLS VPLS or PBB-EVPN I-VPLS services.

- IGMP snooping is enabled using the **config>service>vpls>igmp-snooping no shutdown** command.
- Queries and reports received on SAP or SDP bindings are snooped and properly handled; they are sent to SAP or SDP bindings as expected.
- Queries and reports on EVPN-MPLS or PBB-EVPN B-VPLS destinations are handled as follows.
  - If received from SAP or SDP bindings, the queries and reports are sent to all EVPN-MPLS and PBB-EVPN B-VPLS destinations, regardless of whether the service is using an ingress replication or mLDP provider tunnel.
  - If received on an EVPN-MPLS or PBB-EVPN B-VPLS destination, the queries and reports are processed and propagated to access SAP or SDP bindings, regardless of whether the service is using an ingress replication or mLDP provider tunnel.
  - EVPN-MPLS and PBB-EVPN B-VPLS destinations are is treated as a single IGMP snooping interface and is always added as an **mrouter**.
  - The debug trace output displays one copy of messages being sent to all EVPN-MPLS and PBB-EVPN B-VPLS destinations (the trace does not show a copy for each destination) and displays messages received from all EVPN-MPLS and PBB-EVPN B-VPLS destinations as coming from a single EVPN-MPLS interface.

Note: When IGMP snooping is enabled with P2MP LSPs, at least one EVPN-MPLS multicast destination must be established to enable the processing of IGMP messages by the system.

In the following show command output, the EVPN-MPLS destinations are shown as part of the MFIB (when **igmp-snooping** is in a **no shutdown** state), and the EVPN-MPLS logical interface is shown as an **mrouter**.

```
*A:PE-2# show service id 2000 mfib
```

```
===============================================================================
Multicast FIB, Service 2000
===============================================================================
Source Address  Group Address         SAP or SDP Id                   Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
*               *                     eMpls:192.0.2.3:262132       Local    Fwd
                                      eMpls:192.0.2.4:262136       Local    Fwd
                                      eMpls:192.0.2.5:262131       Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
*A:PE-2# show service id 2000 igmp-snooping base
===============================================================================
IGMP Snooping Base info for service 2000
===============================================================================
Admin State : Up
Querier     : 10.0.0.3 on evpn-mpls
-------------------------------------------------------------------------------
SAP or SDP                   Oper MRtr Pim  Send Max   Max  Max   MVR       Num
Id                         Stat Port Port Qrys Grps  Srcs Grp   From-VPLS Grps
                                                        Srcs
-------------------------------------------------------------------------------
sap:1/1/1:2000             Up   No   No   No   None  None None  Local     0
evpn-mpls                  Up   Yes  N/A  N/A  N/A   N/A  N/A   N/A       N/A
===============================================================================


*A:PE-4# show service id 2000 igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 2000
===============================================================================
MRouter          SAP or SDP Id                 Up Time         Expires   Version
-------------------------------------------------------------------------------
10.0.0.3         evpn-mpls                  0d 00:38:49    175s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
```

The equivalent output for PBB-EVPN services is similar to the output above for EVPN-MPLS services, with the exception that the EVPN destinations are named "b-EVPN-MPLS".

5.3.9.1 Data-driven IGMP Snooping Synchronization with EVPN Multihoming

When single-active multihoming is used, the IGMP snooping state is learned on the active multihoming object. If a failover occurs, the system with the newly active multihoming object must wait for IGMP messages to be received to instantiate the IGMP snooping state after the Ethernet Segment (ES) activation timer expires; this could result in an increased outage.

The outage can be reduced by using MCS synchronization, which is supported for IGMP snooping in both EVPN-MPLS and PBB-EVPN services (see Multi-Chassis Synchronization for Layer 2 Snooping States). However, MCS only supports synchronization between two PEs, whereas EVPN multihoming is supported between a maximum of four PEs. Also, IGMP snooping state can be synchronized only on a SAP.

An increased outage would also occur when using all-active EVPN multihoming. The IGMP snooping state on an ES LAG SAP or virtual ES to the attached CE must be synchronized between all the ES PEs, as the LAG link used by the DF PE might not be the same as that used by the attached CE. MCS synchronization is not applicable to all-active multihoming as MCS only supports active/standby synchronization.

To eliminate any additional outage on a multihoming failover, IGMP snooping messages can be synchronized between the PEs on an ES using data-driven IGMP snooping state synchronization, which is supported in both EVPN-MPLS and PBB-EVPN services. The IGMP messages received on an ES SAP or spoke SDP are sent to the peer ES PEs with an ESI label (for EVPN-MPLS) or ES BMAC (for PBB-EVPN) and these are used to synchronize the IGMP snooping state on the ES SAP or spoke SDP on the receiving PE.

Data-driven IGMP snooping state synchronization is supported for both all-active multihoming and single-active with an ESI label multihoming in EVPN-MPLS services and for all-active multihoming in PBB-EVPN services. All PEs participating in a multihomed ES must be running an SR OS version supporting this capability. PBB-EVPN with IGMP snooping using single-active multihoming is not supported.

Data-driven IGMP snooping state synchronization is also supported with P2MP mLDP LSPs in both EVPN-MPLS and PBB-EVPN services. When P2MP mLDP LSPs are used in EVPN-MPLS services, all PEs (including the PEs not connected to a multihomed ES) in the EVPN-MPLS service must be running an SR OS version supporting this capability with IGMP snooping enabled and all network interfaces must be configured on FP3 or higher-based line cards.

Figure 182 shows the processing of an IGMP message for EVPN-MPLS. In PBB-EVPN services, the ES BMAC is used instead of the ESI label to synchronize the state.

Figure 182 **Data-driven IGMP Snooping Synchronization with EVPN Multihoming**

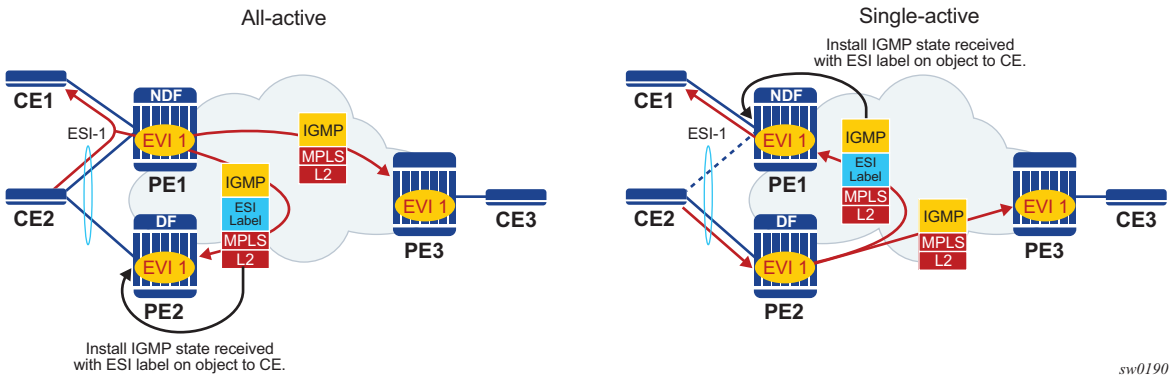

Data-driven synchronization is enabled by default when IGMP snooping is enabled within an EVPN-MPLS service using all-active multihoming or single-active with an ESI label multihoming, or in a PBB-EVPN service using all-active multihoming. If IGMP snooping MCS synchronization is enabled on an EVPN-MPLS or PBB-EVPN (I-VPLS) multihoming SAP then MCS synchronization takes precedence over the data-driven synchronization and the MCS information is used. Mixing data-driven and MCS IGMP synchronization within the same ES is not supported.

When using EVPN-MPLS, the ES should be configured as **non-revertive** to avoid an outage when a PE takes over the DF role. The Ethernet A-D per ESI route update is withdrawn when the ES is down which prevents state synchronization to the PE with the ES down, as it does not advertise an ESI label. The lack of state synchronization means that if the ES comes up and that PE becomes DF after the ES activation timer expires, it might not have any IGMP snooping state until the next IGMP messages are received, potentially resulting in an additional outage. Configuring the ES as **non-revertive** will avoid this potential outage. Configuring the ES to be **non-revertive** would also avoid an outage when PBB-EVPN is used, but there is no outage related to the lack of the ESI label as it is not used in PBB-EVPN.

The following steps can be used when enabling IGMP snooping in EVPN-MPLS and PBB-EVPN services.

- Upgrade SR OS on all ES PEs to a version supporting data-driven IGMP snooping synchronization with EVPN multihoming.
- Enable IGMP snooping in the required services on all ES PEs. Traffic loss will occur until all ES PEs have IGMP snooping enabled and the first set of join/query messages are processed by the ES PEs.
- There is no action required on the non-ES PEs.

If P2MP mLDP LSPs are also configured, the following steps can be used when enabling IGMP snooping in EVPN-MPLS and PBB-EVPN services.

- Upgrade SR OS on all PEs (both ES and non-ES) to a version supporting data-driven IGMP snooping synchronization with EVPN multihoming.
- For EVPN-MPLS:
  - Enable IGMP snooping on all non-ES PEs.Traffic loss will occur until the first set of join/query messages are processed by the non-ES PEs.
  - Then enable IGMP snooping on all ES PEs. Traffic loss will occur until all PEs have IGMP snooping enabled and the first set of join/query messages are processed by the ES PEs.
- For PBB-EVPN:
  - Enable IGMP snooping on all ES PEs. Traffic loss will occur until all PEs have IGMP snooping enabled and the first set of join/query messages are processed by the ES PEs.
  - There is no action required on the non-ES PEs.

To aid with troubleshooting, the debug packet output displays the IGMP packets used for the snooping state synchronization. An example of a join sent on ES esi-1 from one ES PE and the same join received on another ES PE follows.

```
6 2017/06/16 18:00:07.819 PDT MINOR: DEBUG #2001 Base IGMP
"IGMP: TX packet on svc 1
  from chaddr 5e:00:00:16:d8:2e
  send towards ES:esi-1
  Port  : evpn-mpls
  SrcIp : 0.0.0.0
  DstIp : 224.0.0.22
  Type  : V3 REPORT
    Num Group Records: 1
        Group Record Type: MODE_IS_EXCL (2), AuxDataLen 0, Num Sources 0
          Group Addr: 235.0.0.1


4 2017/06/16 18:00:07.820 PDT MINOR: DEBUG #2001 Base IGMP
"IGMP: RX packet on svc 1
  from chaddr d8:2e:ff:00:01:41
  received via evpn-mpls on ES:esi-1
  Port  : sap lag-1:1
  SrcIp : 0.0.0.0
  DstIp : 224.0.0.22
  Type  : V3 REPORT
    Num Group Records: 1
        Group Record Type: MODE_IS_EXCL (2), AuxDataLen 0, Num Sources 0
          Group Addr: 235.0.0.1
```

5.3.10 PIM Snooping for IPv4 in EVPN-MPLS and PBB-EVPN Services

PIM Snooping for VPLS allows a VPLS PE router to build multicast states by snooping PIM protocol packets that are sent over the VPLS. The VPLS PE then forwards multicast traffic based on the multicast states. When all receivers in a VPLS are IP multicast routers running PIM, multicast forwarding in the VPLS is efficient when PIM snooping for VPLS is enabled.

PIM snooping for IPv4 is supported in EVPN-MPLS and PBB-EVPN I-VPLS (where BGP EVPN is running in the associated B-VPLS service) services. It is enabled using the following command (as IPv4 multicast is enabled by default):

```
configure service vpls <service-id> pim-snooping
```

PIM snooping on SAPs and spoke-SDPs operates in the same way as in a plain VPLS service. However, EVPN-MPLS/PBB-EVPN B-VPLS destinations are treated as a single PIM interface, specifically:

- Hellos and join/prune messages from SAPs or SDPs are always sent to all EVPN-MPLS or PBB-EVPN B-VPLS destinations.
- As soon as a hello message is received from one PIM neighbor on an EVPN-MPLS or PBB-EVPN I-VPLS destination, then the single interface representing all EVPN-MPLS or PBB-EVPN I-VPLS destinations will have that PIM neighbor.
- The EVPN-MPLS or PBB-EVPN B-VPLS destination split horizon logic ensures that IP multicast traffic and PIM messages received on an EVPN-MPLS or PBB-EVPN B-VPLS destination are not forwarded back to other EVPN-MPLS or PBB-EVPN B-VPLS destinations.
- The debug trace output displays one copy of messages being sent to all EVPN-MPLS or PBB-EVPN B-VPLS destinations (the trace does not show a copy for each destination) and displays messages received from all EVPN-MPLS or PBB-EVPN B-VPLS destinations as coming from a single EVPN-MPLS interface.

PIM snooping for IPv4 is supported in EVPN-MPLS services using P2MP LSPs and PBB-EVPN I-VPLS services with P2MP LSPs in the associated B-VPLS service. When PIM snooping is enabled with P2MP LSPs, at least one EVPN-MPLS multicast destination is required to be established to enable the processing of PIM messages by the system.

Multi-chassis synchronization (MCS) of PIM snooping for IPv4 state is supported for both SAPs and spoke-SDPs which can be used with single-active multihoming. Care should be taken when using *.null to define the range for a QinQ virtual Ethernet Segment if the associated SAPs are also being synchronized by MCS, as there is no equivalent MCS sync-tag support to the *.null range.

PBB-EVPN services operate in a similar way to regular PBB services, specifically:

- The multicast flooding between the I-VPLS and the B-VPLS works in a similar way as for PIM snooping for IPv4 with an I-VPLS using a regular B-VPLS. The first PIM join message received over the local B-VPLS from a B-VPLS SAP or SDP or EVPN destination will add all of the B-VPLS SAP or SDP or EVPN components into the related multicast forwarding table associated with that I-VPLS context. The multicast packets will be forwarded throughout the B-VPLS on the per ISID single tree.
- When a PIM router is connected to a remote I-VPLS instance over the B-VPLS infrastructure, its location is identified by the B-VPLS SAP, SDP or by the set of all EVPN destinations on which its PIM hellos are received. The location is also identified by the source BMAC address used in the PBB header for the PIM hello message (this is the BMAC associated with the B-VPLS instance on the remote PBB PE).

In EVPN-MPLS services, the individual EVPN-MPLS destinations appear in the MFIB but the information for each EVPN-MPLS destination entry will always be identical, as shown below:

```
*A:PE# show service id 1 mfib
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
*               233.252.0.1           sap:1/1/9:1                  Local    Fwd
                                      eMpls:1.1.1.2:262141         Local    Fwd
                                      eMpls:1.1.1.3:262141         Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
*A:PE#
```

Similarly for the PIM neighbors:

```
*A:PE# show service id 1 pim-snooping neighbor
===============================================================================
PIM Snooping Neighbors ipv4
===============================================================================
Port Id                 Nbr DR Prty     Up Time       Expiry Time     Hold Time
  Nbr Address
-------------------------------------------------------------------------------
SAP:1/1/9:1             1               0d 00:08:17   0d 00:01:29     105
```

```
  10.0.0.1
EVPN-MPLS               1               0d 00:27:26   0d 00:01:19     105
  10.0.0.2
EVPN-MPLS               1               0d 00:27:26   0d 00:01:19     105
  10.0.0.3
-------------------------------------------------------------------------------
Neighbors : 3
===============================================================================
*A:PE#
```

A single EVPN-MPLS interface is shown in the outgoing interface, as can be seen in the output below:

```
*A:PE# show service id 1 pim-snooping group detail
===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 233.252.0.1
Source Address     : *
Up Time            : 0d 00:07:07
Up JP State        : Joined             Up JP Expiry       : 0d 00:00:37
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00
RPF Neighbor       : 10.0.0.1
Incoming Intf      : SAP:1/1/9:1
Outgoing Intf List : EVPN-MPLS, SAP:1/1/9:1
Forwarded Packets  : 0                  Forwarded Octets   : 0
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
*A:PE#
```

An example of the debug trace output for a join received on an EVPN-MPLS destination is shown below:

```
A:PE1# debug service id 1 pim-snooping packet jp
A:PE1#
32 2016/12/20 14:21:22.68 CET MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: Join/Prune
[000 02:16:02.460] PIM-RX ifId 1071394 ifName EVPN-MPLS 10.0.0.3 -
> 224.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0xd3eb
Upstream Nbr IP : 10.0.0.1 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 233.252.0.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Joined Srcs:
                10.0.0.1/32 Flag SWR  <*,G>
```

The equivalent output for PBB-EVPN services is similar to that above for EVPN-MPLS services, with the exception that the EVPN destinations are named "b-EVPN-MPLS".

5.3.10.1 Data-driven PIM Snooping for IPv4 Synchronization with EVPN Multihoming

When single-active multihoming is used, PIM snooping for IPv4 state is learned on the active multihoming object. If a failover occurs, the system with the newly active multihoming object must wait for IPv4 PIM messages to be received to instantiate the PIM snooping for IPv4 state after the ES activation timer expires, which could result in an increased outage.

This outage can be reduced by using MCS synchronization, which is supported for PIM snooping for IPv4 in both EVPN-MPLS and PBB-EVPN services (see Multi-Chassis Synchronization for Layer 2 Snooping States). However, MCS only supports synchronization between two PEs, whereas EVPN multihoming is supported between a maximum of four PEs.

An increased outage would also occur when using all-active EVPN multihoming. The PIM snooping for IPv4 state on an all-active ES LAG SAP or virtual ES to the attached CE must be synchronized between all the ES PEs, as the LAG link used by the DF PE might not be the same as that used by the attached CE. MCS synchronization is not applicable to all-active multihoming as MCS only supports active/standby synchronization.

To eliminate any additional outage on a multihoming failover, snooped IPv4 PIM messages should be synchronized between the PEs on an ES using data-driven PIM snooping for IPv4 state synchronization, which is supported in both EVPN-MPLS and PBB-EVPN services. The IPv4 PIM messages received on an ES SAP or spoke SDP are sent to the peer ES PEs with an ESI label (for EVPN-MPLS) or ES BMAC (for PBB-EVPN) and are used to synchronize the PIM snooping for IPv4 state on the ES SAP or spoke SDP on the receiving PE.

Data-driven PIM snooping state synchronization is supported for all-active multihoming and single-active with an ESI label multihoming in EVPN-MPLS services. All PEs participating in a multihomed ES must be running an SR OS version supporting this capability with PIM snooping for IPv4 enabled. It is also supported with P2MP mLDP LSPs in the EVPN-MPLS services, in which case all PEs (including the PEs not connected to a multihomed ES) must have PIM snooping for IPv4 enabled and all network interfaces must be configured on FP3 or higher-based line cards.

In addition, data-driven PIM snooping state synchronization is supported for all-active multihoming in PBB-EVPN services and with P2MP mLDP LSPs in PBB-EVPN services. All PEs participating in a multihomed ES, and all PEs using PIM proxy mode (including the PEs not connected to a multihomed ES) in the PBB-EVPN service must be running an SR OS version supporting this capability and must have PIM snooping for IPv4 enabled. PBB-EVPN with PIM snooping for IPv4 using single-active multihoming is not supported.

Figure 183 shows the processing of an IPv4 PIM message for EVPN-MPLS. In PBB-EVPN services, the ES BMAC is used instead of the ESI label to synchronize the state.

Figure 183 **Data-driven PIM Snooping for IPv4 Synchronization with EVPN Multihoming**

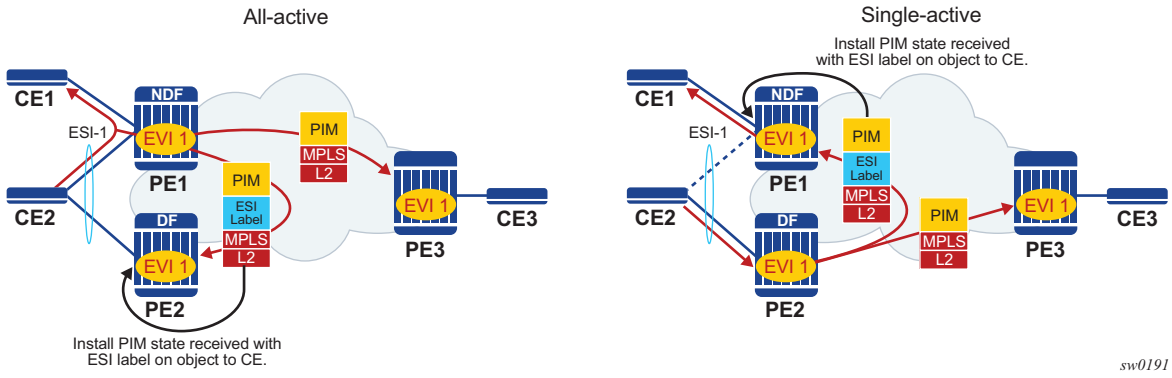

Data-driven synchronization is enabled by default when PIM snooping for IPv4 is enabled within an EVPN-MPLS service using all-active multihoming and single-active with an ESI label multihoming, or in a PBB-EVPN service using all-active multihoming. If PIM snooping for IPv4 MCS synchronization is enabled on an EVPN-MPLS or PBB-EVPN (I-VPLS) multihoming SAP or spoke SDP, then MCS synchronization takes preference over the data-driven synchronization and the MCS information is used. Mixing data-driven and MCS PIM synchronization within the same ES is not supported.

When using EVPN-MPLS, the ES should be configured as **non-revertive** to avoid an outage when a PE takes over the DF role. The Ethernet A-D per ESI route update is withdrawn when the ES is down, which prevents state synchronization to the PE with the ES down as it does not advertise an ESI label. The lack of state synchronization means that if the ES comes up and that PE becomes DF after the ES activation timer expires, it might not have any PIM snooping for IPv4 state until the next PIM messages are received, potentially resulting in an additional outage. Configuring the ES as **non-revertive** will avoid this potential outage. Configuring the ES to be **non-revertive** would also avoid an outage when PBB-EVPN is used, but there is no outage related to the lack of the ESI label as it is not used in PBB-EVPN.

The following steps can be used when enabling PIM snooping for IPv4 (using PIM snooping and PIM proxy modes) in EVPN-MPLS and PBB-EVPN services.

- PIM snooping mode
  - Upgrade SR OS on all ES PEs to a version supporting data-driven PIM snooping for IPv4 synchronization with EVPN multihoming.
  - Enable PIM snooping for IPv4 on all ES PEs. Traffic loss will occur until all PEs have PIM snooping for IPv4 enabled and the first set of join/hello messages are processed by the ES PEs.
  - There is no action required on the non-ES PEs.
- PIM proxy mode
  - EVPN-MPLS
    - Upgrade SR OS on all ES PEs to a version supporting data-driven PIM snooping for IPv4 synchronization with EVPN multihoming.
    - Enable PIM snooping for IPv4 on all ES PEs. Traffic loss will occur until all PEs have PIM snooping for IPv4 enabled and the first set of join/hello messages are processed by the ES PEs.
    - There is no action required on the non-ES PEs.
  - PBB-EVPN
    - Upgrade SR OS on all PEs (both ES and non-ES) to a version supporting data-driven PIM snooping for IPv4 synchronization with EVPN multihoming.
    - Then enable PIM snooping for IPv4 on all non-ES PEs. Traffic loss will occur until all PEs have PIM snooping for IPv4 enabled and the first set of join/hello messages are processed by each non-ES PE.
    - Then enable PIM snooping for IPv4 on all ES PEs. Traffic loss will occur until all PEs have PIM snooping for IPv4 enabled and the first set of join/hello messages are processed by the ES PEs.

If P2MP mLDP LSPs are also configured, the following steps can be used when enabling PIM snooping or IPv4 (using PIM snooping and PIM proxy modes) in EVPN-MPLS and PBB-EVPN services.

- PIM snooping mode
  - Upgrade SR OS on all PEs (both ES and non-ES) to a version supporting data-driven PIM snooping for IPv4 synchronization with EVPN multihoming.
  - Then enable PIM snooping for IPv4 on all ES PEs. Traffic loss will occur until all PEs have PIM snooping enabled and the first set of join/hello messages are processed by the ES PEs.
  - There is no action required on the non-ES PEs.
- PIM proxy mode

- Upgrade SR OS on all PEs (both ES and non-ES) to a version supporting data-driven PIM snooping for IPv4 synchronization with EVPN multihoming.
- Then enable PIM snooping for IPv4 on all non-ES PEs. Traffic loss will occur until all PEs have PIM snooping for IPv4 enabled and the first set of join/hello messages are processed by each non-ES PE.
- Then enable PIM snooping for IPv4 on all ES PEs. Traffic loss will occur until all PEs have PIM snooping enabled and the first set of join/hello messages are processed by the ES PEs.

In the above steps, when PIM snooping for IPv4 is enabled, the traffic loss can be reduced or eliminated by configuring a larger hold-time (up to 300 seconds), during which multicast traffic is flooded.

To aid with troubleshooting, the debug packet output displays the PIM packets used for the snooping state synchronization. An example of a join sent on ES esi-1 from one ES PE and the same join received on another ES PE follows:

```
6 2017/06/16 17:36:37.144 PDT MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: pimVplsFwdJPToEvpn
Forwarding to remote peer on bgp-evpn ethernet-segment esi-1"
7 2017/06/16 17:36:37.144 PDT MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: Join/Prune
[000 00:19:37.040] PIM-TX ifId 1071394 ifName EVPN-MPLS-ES:esi-1 10.0.0.10 -> 22
4.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0xd2de
Upstream Nbr IP : 10.0.0.1 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 235.0.0.10/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Joined Srcs:
                10.0.0.1/32 Flag SWR  <*,G>


4 2017/06/16 17:36:37.144 PDT MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: pimProcessPdu
Received from remote peer on bgp-evpn ethernet-segment esi-1, will be applied on
 lag-1:1
"
5 2017/06/16 17:36:37.144 PDT MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: Join/Prune
[000 00:19:30.740] PIM-RX ifId 1071394 ifName EVPN-MPLS-ES:esi-1 10.0.0.10 -> 22
4.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0xd2de
Upstream Nbr IP : 10.0.0.1 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 235.0.0.10/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Joined Srcs:
                10.0.0.1/32 Flag SWR  <*,G>
```

5.3.11 EVPN E-Tree

This section contains information about EVPN E-Tree.

5.3.11.1 BGP EVPN Control Plane for EVPN E-Tree

BGP EVPN control plane is extended and aligned with IETF Draft *draft-ietf-bess-evpn-etree* to support EVPN E-Tree services. Figure 184 shows the main EVPN extensions for the EVPN E-Tree information model.

Figure 184 **EVPN E-Tree BGP Routes**

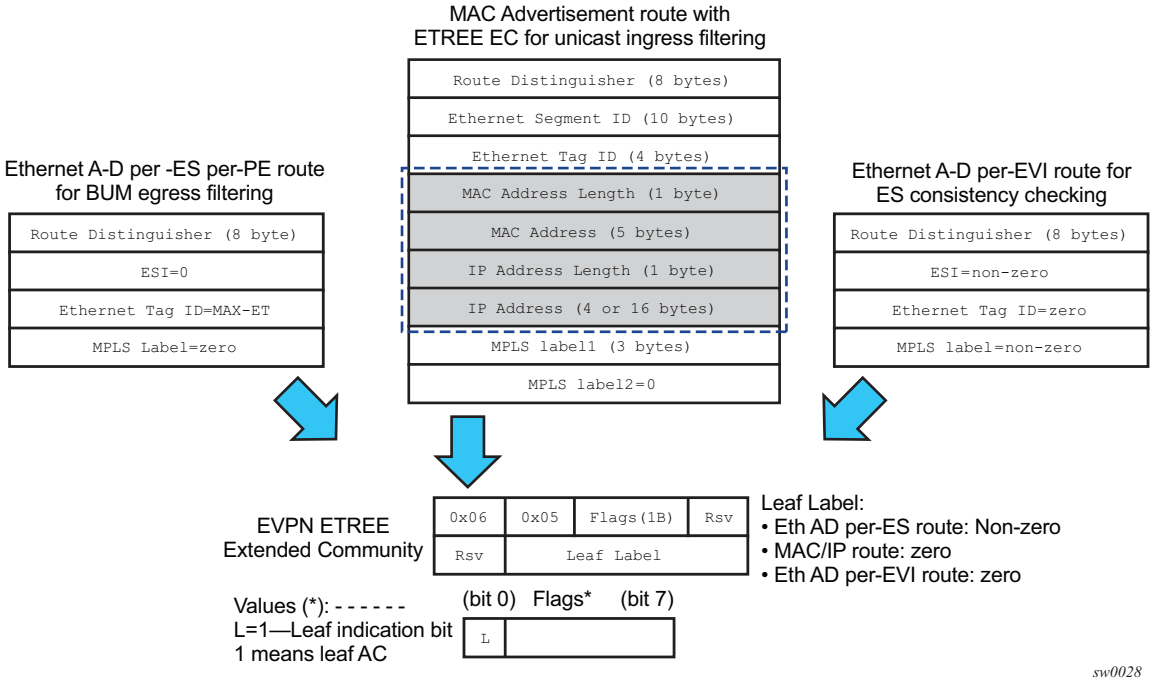

The following BGP extensions are implemented for EVPN E-Tree services.

- An EVPN E-Tree extended community (EC) sub-type 0x5 is defined. The following information is included.
  - The lower bit of the **Flags** field contains the L bit (where L=1 indicates leaf AC).
  - The Leaf label contains a 20-bit MPLS label in the high-order 20 bits of the label field.
- The new E-Tree EC is sent with the following routes.
  - AD per-ES per PE route for BUM egress filtering

    Each EVPN E-Tree capable PE advertises an AD per-ES route with the E-Tree EC, and the following information:
    - Service RD and route-target

If **ad-per-es-route-target evi-rt-set** is configured, then non-zero ESI AD per-ES routes (used for multi-homing) are sent per the **evi-rt-set** configuration, but E-Tree zero-ESI routes (used for E-Tree) are sent based on the default **evi-rt** configuration.

- ESI = 0

  Eth Tag = MAX-ET

  MPLS label = zero

– AD per-EVI route for root or leaf configuration consistency check as follows:
  - The E-Tree EC is sent with the AD per-EVI routes for a specific ES. In this case, no validation is performed by the implementation, and the leaf indication is only used for troubleshooting on the remote PEs.
  - The MPLS label value is zero.
  - All ACs in each EVI for a specific ES must be configured as either a root or leaf AC, but not a combination. In case of a configuration error, for example where the AC in PE1 is configured as root and in PE2 as leaf AC, the remote PE3 will receive the AD per-EVI routes with inconsistent leaf indication. However, the unicast filtering remains unaffected and is still handled by the FDB lookup information.

– MAC or IP routes for known unicast ingress filtering as follows:
  - An egress PE sends all MAC or IP routes learned over a leaf AC SAP or spoke-SDP with this E-Tree EC indicating that the MAC or IP belongs to a leaf AC.
  - The MPLS label value in the EC is 0.
  - Upon receiving a route with E-Tree EC, the ingress PE imports the route and installs the MAC in the FDB with a leaf flag (if there is a leaf indication in the route). Any frame coming from a leaf AC for which the MAC DA matches a leaf AC MAC is discarded at the ingress.
  - If two PEs send the same MAC with the same ESI but inconsistent root or leaf indication, the MAC is installed in the FDB as root.

5.3.11.2 EVPN for MPLS Tunnels in E-Tree Services

EVPN E-Tree services are modeled as VPLS services configured as E-Trees with **bgp-evpn mpls** enabled.

The following example CLI shows an excerpt of a VPLS E-Tree service with EVPN E-Tree service enabled.

```
*A:PE1>config>service>system>bgp-evpn#
  evpn-etree-leaf-label
*A:PE1>config>service# vpls 1 customer 1 etree create
```

```
*A:PE1>config>service>vpls(etree)# info
----------------------------------------------
  description "ETREE-enabled evpn-mpls-service"
  bgp-evpn
    evi 10
    mpls
      no shutdown
      ecmp 2
      auto-bind-tunnel resolution any
      ingress-replication-bum-label
  sap lag-1:1 leaf-ac create
  exit
  sap 2/1/1:1 leaf-ac create
  exit
  sap 2/2/2:1 create
  exit
  spoke-sdp 3:1 leaf-ac create
  exit
```

The following considerations apply to the configuration of the EVPN E-Tree service.

- The **config**>**service**>**system**>**bgp-evpn# evpn-etree-leaf-label** command is required prior to the configuration of any EVPN E-Tree service, and is relevant for EVPN E-Tree services only. The command allocates an E-Tree leaf label on the system and programs the ILM for this label. The ILM label ensures that in-flight traffic always has an ILM entry for lookup, therefore avoiding discards when the service is **shutdown** and subsequently **no shutdown** in a short period of time.
- The **config**>**service# vpls x etree create** command is compatible with the **bgp-evpn mpls** context.
- As in VPLS E-Tree services, an AC that is not configured as a leaf AC is treated as root AC.
- MAC addresses learned over a leaf AC SAP or SDP binding are advertised as leaf MAC addresses.
- Any PE with one or more **bgp-evpn** enabled VPLS E-Tree service advertises an AD per-ES per-PE route with the leaf indication and leaf label that will be used for BUM egress filtering.
- Any **leaf-ac** SAP or SDP binding defined in an ES triggers the advertisement of an AD per-EVI route with the leaf indication.
- EVPN E-Tree services use the following CLI commands:
  - **sap** *sap-id* **leaf-ac create**
  - **mesh-sdp** *sdp-id:vc-id* **leaf-ac create**
  - **spoke-sdp** *sdp-id:vc-id* **leaf-ac create**
- The **root-leaf-tag** command is blocked in VPLS E-Tree services where **bgp-evpn mpls** is enabled

5.3.11.3 EVPN E-Tree Operation

EVPN E-Tree supports all operations related to flows among local root AC and leaf AC objects in accordance with IETF Draft *draft-ietf-bess-evpn-etree*. This section describes the extensions required to forward to or from BGP-EVPN destinations.

5.3.11.3.1 EVPN E-Tree Known Unicast Ingress Filtering

Known unicast traffic forwarding is based on ingress PE filtering. Figure 185 shows an example of EVPN-E-Tree forwarding behavior for known unicast.

Figure 185 **EVPN E-Tree Known Unicast Ingress Filtering**

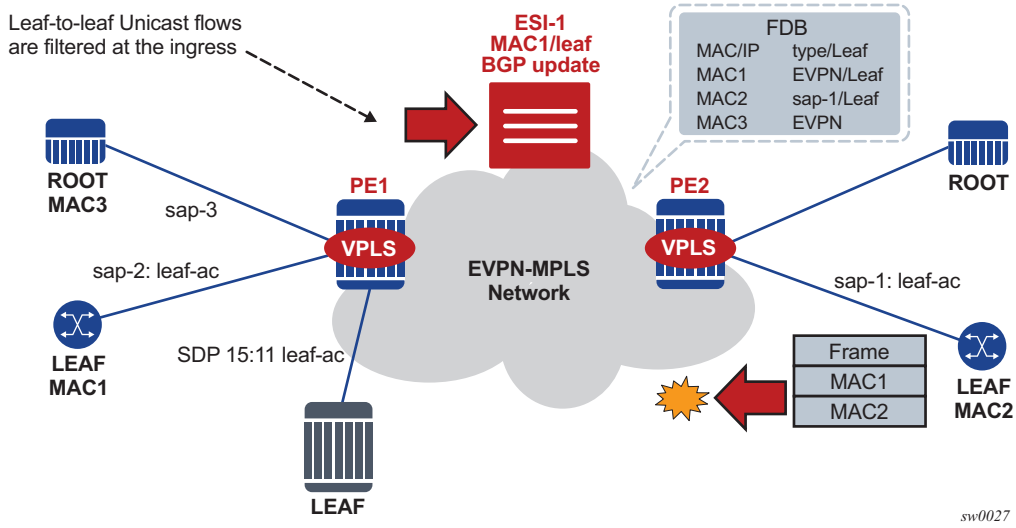

MAC addresses learned on *leaf-ac* objects are advertised in EVPN with their corresponding leaf indication.

In Figure 185, PE1 advertises MAC1 using the E-Tree EC and leaf indication, and PE2 installs MAC1 with a leaf flag in the FDB.

Assuming MAC DA is present in the local FDB (MAC1 in the FDB of PE2) when PE2 receives a frame, it is handled as follows.

- If the unicast frame enters a root AC, the frame follows regular data plane procedures; that is, it is sent to the owner of the MAC DA (local SAP or SDP binding or remote BGP EVPN PE) without any filtering.
- If the unicast frame enters a leaf AC, it is handled as follows.
  - A MAC DA lookup is performed on the FDB.

- If there is a hit and the MAC was learned as an EVPN leaf (or from a leaf AC), then the frame is dropped at ingress.
- The source MAC (MAC2) is learned and marked as a leaf-learned MAC. It is advertised by the EVPN with the corresponding leaf indication.

- A MAC received with a root and leaf indication from different PEs in the same ES is installed as root.

The ingress filtering for E-Tree leaf-to-leaf traffic requires the implementation of an extra leaf EVPN MPLS destination per remote PE (containing leaf objects) per E-Tree service. The ingress filtering for E-Tree leaf-to-leaf traffic is as follows.

- A separate EVPN MPLS bind is created for unicast leaf traffic in the service. The internal EVPN MPLS destination is created for each remote PE that contains a leaf and advertises at least one leaf MAC.
- The creation of the internal EVPN MPLS destination is triggered when a MAC route with L=1 in the E-Tree EC is received. Any EVPN E-Tree service can potentially use one extra EVPN MPLS destination for leaf unicast traffic per remote PE.

  The extra destination in the EVPN E-Tree service is for unicast only and it is not part of the flooding list. It is resource-accounted and displayed in the **tools dump service evpn usage** command, as shown in the following example output.

```
A:PE-4# tools dump service evpn usage
vxlan-evpn-mpls usage statistics at 01/23/2017 00:53:14:
MPLS-TEP                                          :              3
VXLAN-TEP                                         :              0
Total-TEP                                         :      3/ 16383
Mpls Dests (TEP, Egress Label + ES + ES-BMAC)     :             10
Mpls Etree Leaf Dests                             :              1
Vxlan Dests (TEP, Egress VNI)                     :              0
Total-Dest                                        :     10/196607
Sdp Bind +  Evpn Dests                            :     13/245759
ES L2/L3 PBR                                      :      0/ 32767
Evpn Etree Remote BUM Leaf Labels                 :              3
```

- MACs received with L=1 point to the EVPN MPLS destination, whereas root MACs point to the “root” destination.

5.3.11.3.2 EVPN E-Tree BUM Egress Filtering

BUM traffic forwarding is based on egress PE filtering. Figure 186 shows an example of EVPN E-Tree forwarding behavior for BUM traffic.

Figure 186 **EVPN E-Tree BUM Egress Filtering**

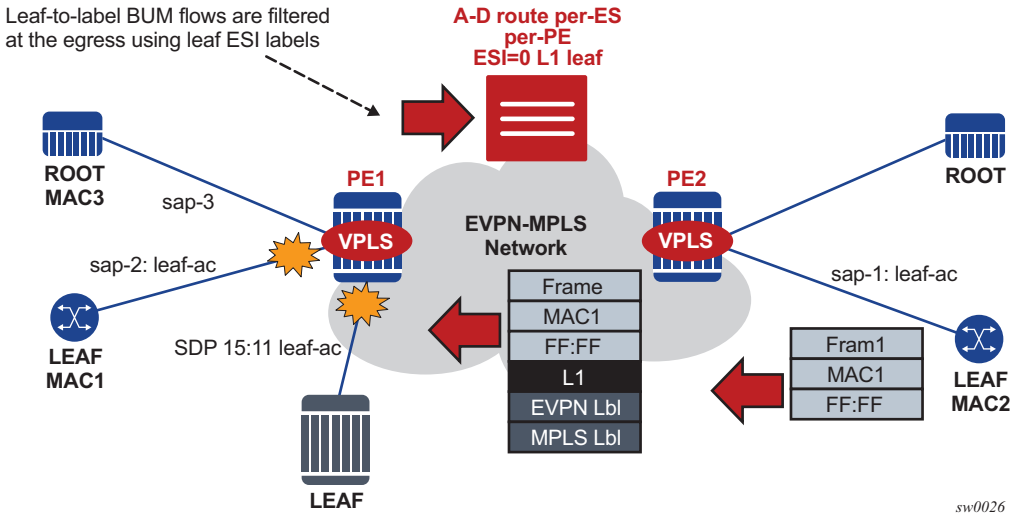

In Figure 186, BUM frames are handled as follows when they ingress PE or PE2.

- If the BUM frame enters a root AC, the frame follows regular EVPN data plane procedures.
- If the BUM frame enters a leaf AC, the frame handling is as follows:
  - The frame is marked as leaf and forwarded or replicated to the egress IOM.
  - At the egress IOM, the frame is flooded in the default multicast list subject to the following.
    - Leaf entries are skipped when BUM traffic is forwarded, but this excludes BUM traffic forwarded to local leaf AC objects.
    - Traffic to remote BGP EVPN PEs is encapsulated with the EVPN label stack. If a leaf ESI label present for the far-end PE (L1 in Figure 186), the leaf ESI label is added at the bottom of the stack; the remaining stack will follow (including EVI label). If there is no leaf ESI label for the far-end egress PE, no additional label is added to the stack. This means that the egress PE does not have any E-Tree enabled service, but it can still work with the VPLS E-Tree service available in PE2.

The BUM-encapsulated packet is received on the network ingress interface at the egress PE or PE1. The packet is processed as follows.

- A normal ILM lookup is performed for each label (including the EVI label) in the stack.
- Further label lookups are performed when the EVI label ILM lookup is complete. If the lookup yields a leaf label, all the leaf ACs are skipped when flooding to the default-multicast list at the egress PE.

5.3.11.3.3 EVPN E-Tree Egress Filtering Based on MAC Source Address

The egress PE checks the MAC Source Address (SA) for traffic received without the leaf MPLS label. This check covers corner cases where the ingress PE sends traffic originating from a leaf AC but without a leaf indication.

In Figure 186, PE2 receives a frame with MAC DA = MAC3 and MAC SA = MAC2. Because MAC3 is a root MAC, MAC lookup at PE2 allows the system to unicast the packet to PE1 without the leaf label. If MAC3 was no longer in PE1's FDB, PE1 would flood the frame to all the root and leaf ACs, despite the frame having originated from a leaf AC.

To minimize and prevent leaf traffic from leaking to other leaf ACs (as described in the preceding case), the egress PE always performs a MAC SA check for all types of traffic. The data path performs MAC SA-based egress filtering as follows.

- An Ethernet frame may be treated as originating from a leaf AC due to several reasons, which require the system to set a flag to indicate leaf traffic. The flag is set if one of the following conditions is true: if the frames arrive on a leaf SAP, if EVPN traffic arrives with a leaf label, or if a MAC SA is flagged as a leaf SA.
- After the flag is set, the action taken depends on the type of traffic.
  - Unicast traffic: An FDB lookup is performed, and if the MAC DA FDB entry is marked as a leaf type, the frame is dropped to prevent leaf-to-leaf forwarding.
  - BUM traffic: The flag is taken into account at the egress IOM and leaf-to-leaf forwarding is suppressed.

5.3.11.4 EVPN E-Tree and EVPN Multi-homing

EVPN E-Tree procedures support all-active and single-active EVPN multi-homing. Ingress filtering can handle MACs learned on ES leaf AC SAP or SDP-bindings. If a MAC associated with an ES leaf AC is advertised with a different E-Tree indication or if the AD per-EVI routes have inconsistent leaf indications, then the remote PEs performing the aliasing treat the MAC as root.

Figure 187 shows the expected behavior for multi-homing and egress BUM filtering.

Figure 187 **EVPN E-Tree BUM Egress Filtering and Multi-homing**

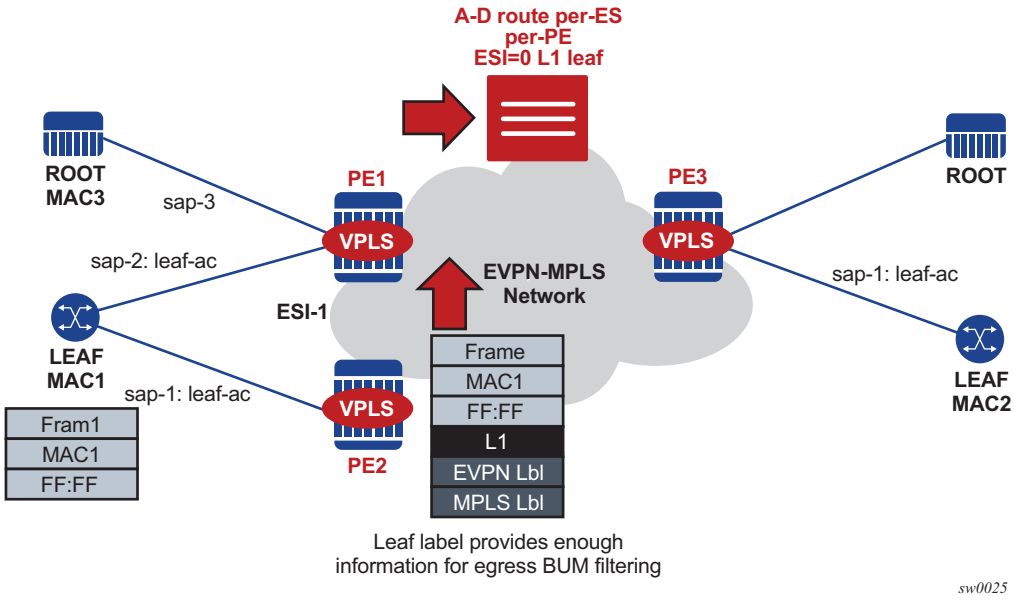

Multi-homing and egress BUM filtering in Figure 187 is handled as follows:

- BUM frames received on an ES leaf AC are flooded to the EVPN based on EVPN E-Tree procedures. The leaf ESI label is sent when flooding to other PEs in the same ES, and additional labels are not added to the stack.

  When flooding in the default multicast list, the egress PE skips all the leaf ACs (including the ES leaf ACs) on the assumption that all ACs in a specific ES for a specified EVI have a consistent E-Tree configuration, and they send an AD per-EVI route with a consistent E-Tree indication.
- BUM frames received on an ES root AC are flooded to the EVPN based on regular EVPN procedures. The regular ES label is sent for split-horizon when packets are sent to the DF or NDF PEs in the same ES. When flooding in the default multicast list, the egress PE skips the ES SAPs based on the ES label lookup.

If the PE receives an ES MAC from a peer that shares the ES and decides to install it against the local ES SAP that is **oper-up**, it checks the E-Tree configuration (root or leaf) of the local ES SAP against the received MAC route. The MAC route is processed as follows.

- If the E-Tree configuration does not match, then the MAC is not installed against any destination until the misconfiguration is resolved.
- If the SAP is **oper-down**, the MAC is installed against the EVPN destination to the peer.

5.3.11.5 PBB-EVPN E-Tree Services

SR OS supports PBB-EVPN E-Tree services in accordance with IETF Draft *draft-ietc-bess-evpn-etree*. PBB-EVPN E-Tree services are modeled as PBB-EVPN services where some I-VPLS services are configured as **etree** and some of their SAP or spoke-SDPs are configured as leaf ACs.

The procedures for the PBB-EVPN E-Tree are similar to those for the EVPN E-Tree, except that the egress leaf-to-leaf filtering for BUM traffic is based on the BMAC source address. Also, the leaf label and the EVPN AD routes are not used.

The PBB-EVPN E-Tree operation is as follows.

- When one or more I-VPLS E-Tree services are linked to a B-VPLS, the leaf backbone source MAC address (**leaf-source-bmac** parameter) is used for leaf-originated traffic in addition to the source B-VPLS MAC address (**source-bmac** parameter) that is used for sourcing root traffic.
- The leaf backbone source MAC address for PBB must be configured using the command **config**>**service**>**pbb**>**leaf-source-bmac** *ieee-address* prior to the configuration of any I-VPLS E-Tree service.
- The **leaf-source-bmac** address is advertised in a BMAC route with a leaf indication.
- Known unicast filtering occurs at the ingress PE. When a frame enters an I-VPLS leaf AC, a MAC lookup is performed. If the CMAC DA is associated with a leaf BMAC, the frame is dropped.
- Leaf-to-leaf BUM traffic filtering occurs at the egress PE. When flooding BUM traffic with the BMAC SA matching a leaf BMAC, the egress PE skips the I-VPLS leaf ACs.

The following CLI example shows an I-VPLS E-Tree service that uses PBB-EVPN E-Tree. The **leaf-source-bmac** address must be configured prior to the configuration of the I-VPLS E-Tree. As is the case in regular E-Tree services, SAP and spoke-SDPs that are not explicitly configured as leaf ACs are considered root AC objects.

```
A:PE-2>config>service# info
----------------------------------------------
        pbb
            leaf-source-bmac 00:00:00:00:00:22
        exit
        vpls 1000 customer 1 b-vpls create
            service-mtu 2000
            bgp
            exit
            bgp-evpn
                evi 1000
                vxlan
                    shutdown
```

```
            exit
            mpls
                ingress-replication-bum-label
                auto-bind-tunnel
                    resolution any
                exit
                no shutdown
            exit
        exit
        stp
            shutdown
        exit
        no shutdown
    exit
    vpls 1001 customer 1 i-vpls etree create
        pbb
            backbone-vpls 1000
            exit
        exit
        stp
            shutdown
        exit
        sap 1/1/1:1001 leaf-ac create
            no shutdown
        exit
        sap 1/1/1:1002 create
            no shutdown
        exit
        no shutdown
    exit
```

The following considerations apply to PBB-EVPN E-Trees and multi-homing.

- All-active multi-homing is not supported on leaf AC I-VPLS SAPs.
- Single-active multi-homing is supported on leaf AC I-VPLS SAPs and spoke-SDPs.
- ISID- and RFC 7623-based CMAC-flush are supported in addition to PBB-EVPN E-Tree services and single-active multi-homing.

5.3.12 MPLS Entropy Label and Hash Label

The router supports the MPLS entropy label (RFC 6790) and the Flow Aware Transport label (known as the hash label) (RFC 6391). These labels allow LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. The entropy label can be enabled on bgp-evpn services (VPLS and Epipe). See the *7450 ESS, 7750 SR, and 7950 XRS MPLS Guide* for further information.

5.3.13 Inter-AS Option B and Next-Hop-Self Route-Reflector for EVPN-MPLS

Inter-AS Option B and Next-Hop-Self Route-Reflector (VPN-NH-RR) functions are supported for the BGP-EVPN family in the same way both functions are supported for IP-VPN families.

A typical use-case for EVPN Inter-AS Option B or EVPN VPN-NH-RR is Data Center Interconnect (DCI) networks, where cloud and service providers are looking for efficient ways to extend their Layer 2 and Layer 3 tenant services beyond the data center and provide a tighter DC-WAN integration. While the instantiation of EVPN services in the DC GW to provide this DCI connectivity is a common model, some operators use Inter-AS Option B or VPN-NH-RR connectivity to allow the DC GW to function as an ASBR or ABR respectively, and the services are only instantiated on the edge devices.

Figure 188 shows a DCI example where the EVPN services in two DCs are interconnected without the need for instantiating services on the DC GWs.

Figure 188 **EVPN Inter-AS Option B or VPN-NH-RR Model**

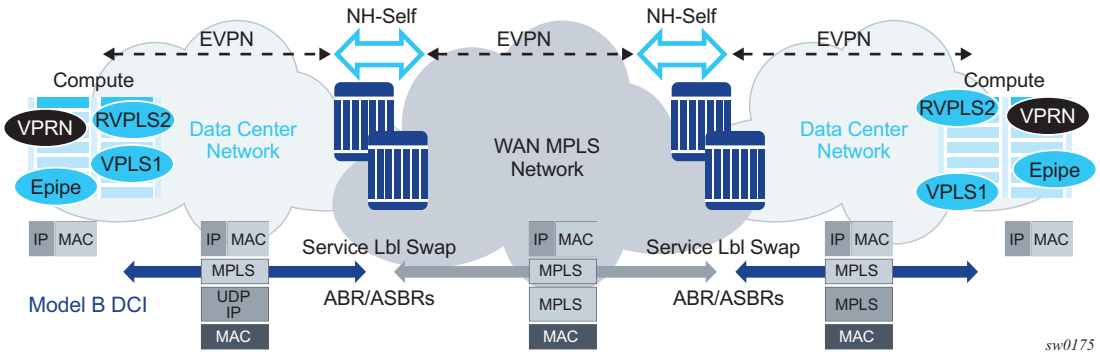

The ASBRs or ABRs connect the DC to the WAN at the control plane and data plane levels where the following considerations apply.

- From a control plane perspective, the ASBRs or ABRs perform the following tasks:
  - accept EVPN-MPLS routes from a BGP peer

    EVPN-VXLAN routes are not supported.
  - extract the MPLS label from the EVPN NLRI or attribute and program a label swap operation on the IOM
  - re-advertise the EVPN-MPLS route to the BGP peer in the other Autonomous Systems (ASs) or IGP domains

The re-advertised route will have a Next-Hop-Self and a new label encoded for those routes that came with a label.

- From a data plan perspective, the ASBRs and ABRs terminate the ingress transport tunnel, perform an EVPN label swap operation, and send the packets on to an interface (if E-BGP is used) or a new tunnel (if I-BGP is used).
- The ASBR or ABR resolves the EVPN routes based on the existing **bgp next-hop-resolution** command for **family vpn**, where **vpn** refers to EVPN, VPN-IPv4, and VPN-IPv6 families.

```
*A:ABR-1# configure router bgp next-hop-resolution labeled-routes transport-
tunnel family vpn resolution-filter
  - resolution-filter
 [no] bgp             - Use BGP tunnelling for next hop resolution
 [no] ldp             - Use LDP tunnelling for next hop resolution
 [no] rsvp            - Use RSVP tunnelling for next hop resolution
 [no] sr-isis         - Use sr-isis tunnelling for next hop resolution
 [no] sr-ospf         - Use sr-ospf for next hop resolution
 [no] sr-te           - Use sr-te for next hop resolution
 [no] udp             - Use udp for next hop resolution
```

Refer to the *7450 ESS, 7750 SR, and 7950 XRS Unicast Routing Protocols Guide* for more information about the next-hop resolution of BGP-labeled routes.

Inter-AS Option B for EVPN services on ABSRs and VPN-NH-RR on ABRs re-use the existing commands **enable-inter-as-vpn** and **enable-rr-vpn-forwarding** respectively. The two commands enable the ASBR or ABR function for both EVPN and IP-VPN routes. These two features can be used with the following EVPN services:

- EVPN-MPLS Epipe services (EVPN-VPWS)
- EVPN-MPLS VPLS services
- EVPN-MPLS R-VPLS services
- PBB-EVPN and PBB-EVPN E-Tree services
- EVPN-MPLS E-Tree services
- PE and ABR functions (EVPN services and **enable-rr-vpn-forwarding**), which are both supported on the same router
- PE and ASBR functions (EVPN services and **enable-inter-as-vpn**), which are both supported on the same router

The following sub-sections clarify some aspects of EVPN when used in an Inter-AS Option B or VPN-NH-RR network.

5.3.13.1 Inter-AS Option B and VPN-NH-RR Procedures on EVPN Routes

When **enable-rr-vpn-forwarding** or **enable-inter-as-vpn** is configured, only EVPN-MPLS routes are processed for label swap and the next hop is changed. EVPN-VXLAN routes are re-advertised without a change in the next hop.

The following shows how the router processes and re-advertises the different EVPN route types. Refer to BGP-EVPN Control Plane for MPLS Tunnels for detailed information about the route fields.

- Auto-discovery (AD) routes (Type 1)

  For AD per EVI routes, the MPLS label is extracted from the route NLRI. The route is re-advertised with Next-Hop-Self (NHS) and a new label. No modifications are made for the remaining attributes.

  For AD per ES routes, the MPLS label in the NLRI is zero. The route is re-advertised with NHS and the MPLS label remains zero. No modifications are made for the remaining attributes.
- MAC/IP routes (Type 2)

  The MPLS label (Label-1) is extracted from the NLRI. The route is re-advertised with NHS and a new Label-1. No modifications are made for the remaining attributes.
- Inclusive Multicast Ethernet Tag (IMET) routes (Type 3)

  Because there is no MPLS label present in the NLRI, the MPLS label is extracted from the PMSI Tunnel Attribute (PTA) if needed, and the route is then re-advertised with NHS, with the following considerations.
  - For IMET routes with tunnel-type Ingress Replication, the router extracts the IR label from the PTA. The router programs the label swap and re-advertises the route with a new label in the PTA.
  - For tunnel-type P2MP mLDP, the router re-advertises the route with NHS. No label is extracted; therefore, no swap operation occurs.
  - For tunnel-type Composite, the IR label is extracted from the PTA, the swap operation is programmed and the route re-advertised with NHS. A new label is encoded in the PTA's IR label with no other changes in the remaining fields.
  - For tunnel-type AR, the routes are always considered VXLAN routes and are re-advertised with the next-hop unchanged.
- Ethernet-Segment (ES) routes (Type 4)

Because ES routes do not contain an MPLS label, the route is re-advertised with NHS and no modifications to the remaining attributes. Although an ASBR or ABR will re-advertise ES routes, EVPN multi-homing for ES PEs located in different ASs or IGMP domains is not supported.

- IP-Prefix routes (Type 5)

  The MPLS label is extracted from the NLRI and the route is re-advertised with NHS and a new label. No modifications are made to the remaining attributes.

5.3.13.2 BUM Traffic in Inter-AS Option B and VPN-NH-RR Networks

Inter-AS Option B and VPN-NH-RR support the use of non-segmented trees for forwarding BUM traffic in EVPN.

For ingress replication and non-segmented trees, the ASBR or ABR performs an EVPN BUM label swap without any aggregation or further replication. This concept is shown in Figure 189.

Figure 189 **VPN-NH-RR and Ingress Replication for BUM Traffic**

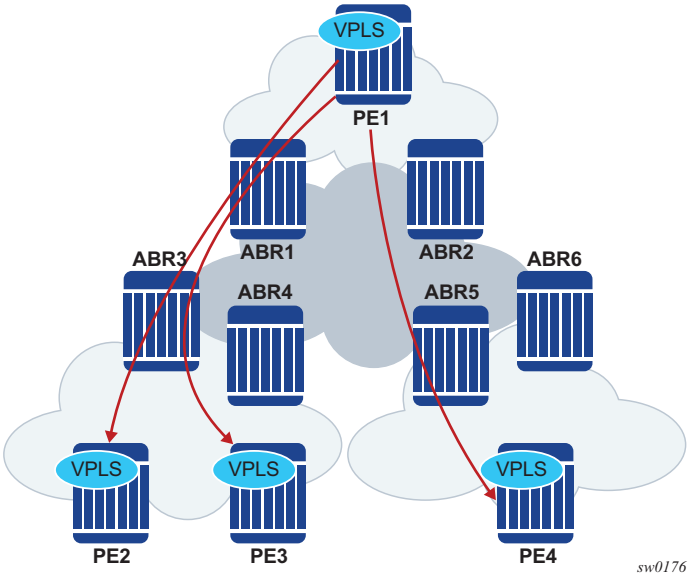

In Figure 189, when PE2, PE3, and PE4 advertise their IMET routes, the ABRs re-advertise the routes with NHS and a different label. However, IMET routes are not aggregated; therefore, PE1 sets up three different EVPN multicast destinations and sends three copies of every BUM packet, even if they are sent to the same ABR. This example is also applicable to ASBRs and Inter-AS Option B.

P2MP mLDP may also be used with VPN-NH-RR, but not with Inter-AS Option B. The ABRs, however, will not aggregate or change the mLDP root IP addresses in the IMET routes. The root IP addresses must be leaked across IGP domains. For example, if PE2 advertises an IMET route with mLDP or composite tunnel type, PE1 will be able to join the mLDP tree if the root IP is leaked into PE1's IGP domain.

5.3.13.3 EVPN Multi-Homing in Inter-AS Option B and VPN-NH-RR Networks

In general, EVPN multi-homing is supported in Inter-AS Option B or VPN-NH-RR networks with the following limitations.

- An ES PE can only process a remote ES route correctly if the received next hop and origination IP address match. EVPN multi-homing is not supported when the ES PEs are in different ASs or IGP domains, or if there is an NH-RR peering the ES PEs and overriding the ES route next hops.
- EVPN multi-homing Ethernet segments are not supported on EVPN PEs that are also ABRs or ASBRs.
- Mass-withdraw based on the AD per-ES routes is not supported for a PE that is in a different AS or IGP domain that the ES PEs. Figure 190 shows an EVPN multi-homing scenario where the ES PEs, PE2 and PE3, and the remote PE, PE1, are in different ASs or IGP domains.

Figure 190 **EVPN Multi-Homing with Inter-AS Option B or VPN-NH-RR**

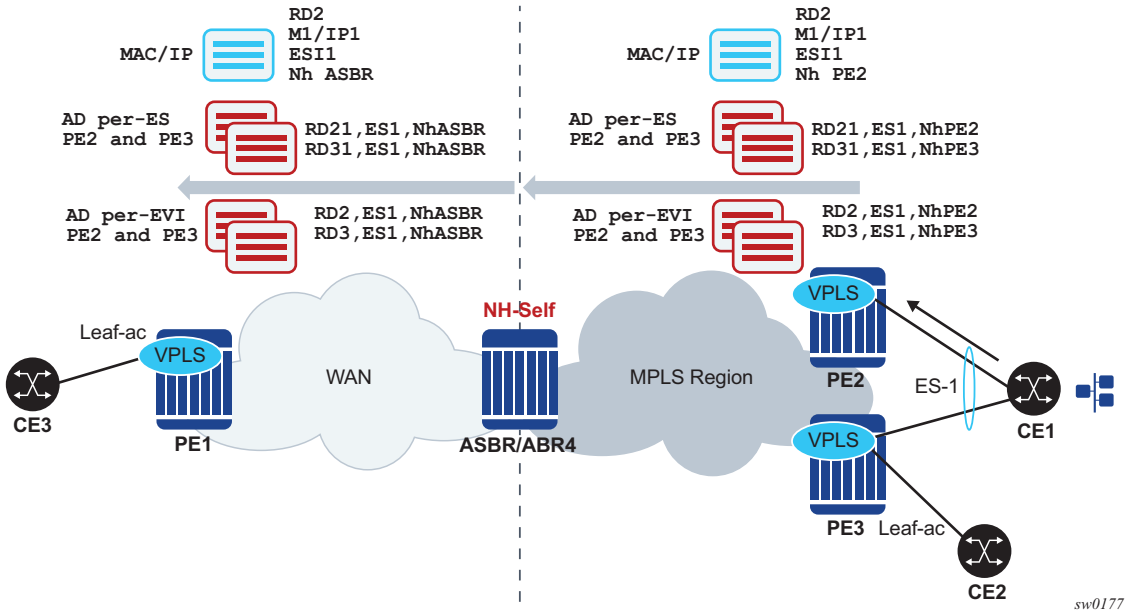

In Figure 190, PE1's aliasing and backup functions to the remote ES-1 are supported. However, PE1 cannot identify the originating PE for the received AD per-ES routes because they are both arriving with the same next hop (ASBR/ABR4) and RDs may not help to correlate each AD per-ES route to a given PE. Therefore, if there is a failure on PE2's ES link, PE1 cannot remove PE2 from the destinations list for ES-1 based on the AD per-ES route. PE1 must wait for the AD per-EVI route withdrawals to remove PE2 from the list. In summary, when the ES PEs and the remote PE are in different ASs or IGP domains, per-service withdrawal based on AD per-EVI routes is supported, but mass-withdrawal based on AD per-ES routes is not supported.

5.3.13.4 EVPN E-Tree in Inter-AS Option B and VPN-NH-RR Networks

Unicast procedures known to EVPN-MPLS E-Tree are supported in Inter-AS Option B or VPN-NH-RR scenarios, however, the BUM filtering procedures are affected.

As described in EVPN E-Tree, leaf-to-leaf BUM filtering is based on the Leaf Label identification at the egress PE. In a non-Inter-AS or non-VPN-NH-RR scenario, EVPN E-tree AD per-ES (ESI-0) routes carrying the Leaf Label are distinguished by the advertised next hop. In Inter-AS or VPN-NH-RR scenarios, all the AD per-ES routes are received with the ABR or ASBR next hop. Therefore, AD per-ES routes originating from different PEs would all have the same next hop, and the ingress PE would not be able to determine which leaf label to use for a given EVPN multicast destination.

A simplified EVPN E-Tree solution is supported, where an E-Tree Leaf Label is not installed in the IOM if the PE receives more than one E-Tree AD per-ES route, with different RDs, for the same next hop. In this case, leaf BUM traffic is transmitted without a Leaf Label and the leaf-to-leaf traffic filtering depends on the egress source MAC filtering on the egress PE. See EVPN E-Tree Egress Filtering Based on MAC Source Address.

PBB-EVPN E-tree services are not affected by Inter-AS or VPN-NH-RR scenarios, as AD per-ES routes are not used.

5.4 General EVPN Topics

This section provides information on general topics related to EVPN.

5.4.1 ARP/ND Snooping and Proxy Support

VPLS services support proxy-ARP (Address Resolution Protocol) and proxy-ND (Neighbor Discovery) functions that can be enabled or disabled independently per service. When enabled (proxy-ARP/proxy-ND no shutdown), the system will populate the corresponding proxy-ARP/proxy-ND table with IP->MAC entries learned from the following sources:

- EVPN-received IP->MAC entries
- User-configured static IP->MAC entries
- Snooped dynamic IP->MAC entries (learned from ARP/GARP/NA messages received on local SAPs/SDP-bindings)

In addition, any ingress ARP or ND frame on a SAP or SDP-binding will be intercepted and processed. ARP requests and Neighbor Solicitations will be answered by the system if the requested IP address is present in the proxy table.

Figure 191 shows an example of how proxy-ARP is used in an EVPN network. Proxy-ND would work in a similar way. The MAC address notation in the diagram is shortened for readability.

Figure 191 **Proxy-ARP Example Usage in an EVPN Network**

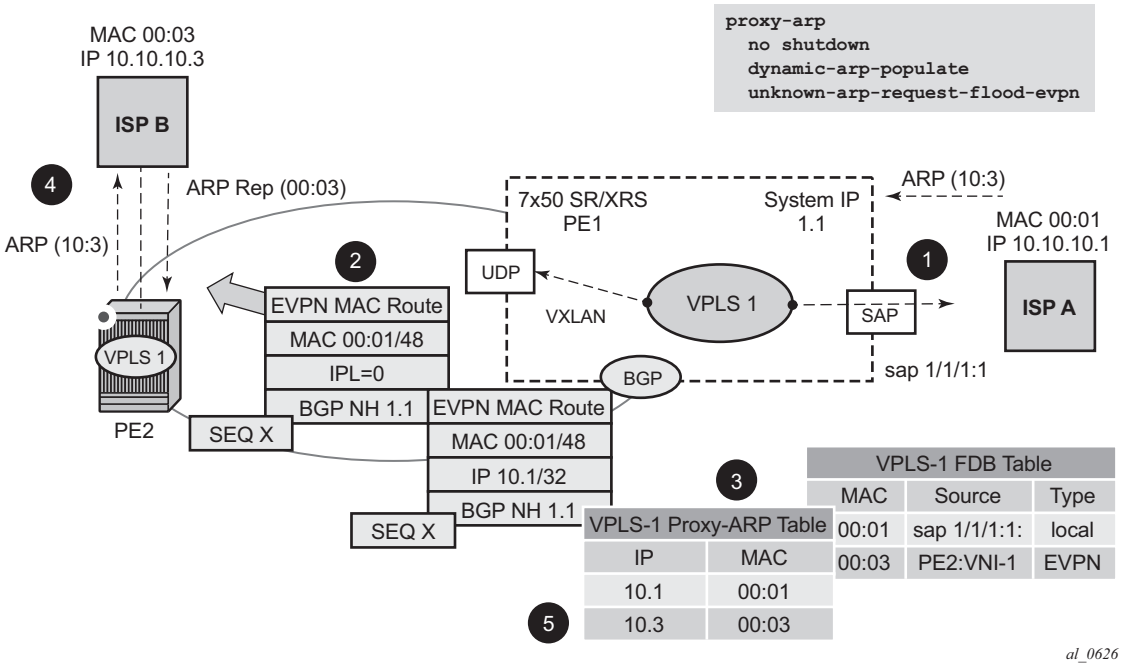

PE1 is configured as follows:

```
*A:PE1>config>service>vpls# info
----------------------------------------------
vxlan vni 600 create
        exit
        bgp
            route-distinguisher 192.0.2.71:600
            route-target export target:64500:600 import target:64500:600
    exit
    bgp-evpn
        vxlan
            no shutdown
        exit
    exit
    proxy-arp
        age-time 600
        send-refresh 200
        dup-detect window 3 num-moves 3 hold-down max anti-spoof-
mac 00:ca:ca:ca:ca:ca
        dynamic-arp-populate
          no shutdown
            exit
            sap 1/1/1:600 create
            exit
no shutdown
----------------------------------------------
```

Figure 191 shows the following steps, assuming proxy-ARP is no shutdown on PE1 and PE2, and the tables are empty:

1. ISP-A sends ARP-request for (10.10.)10.3.
2. PE1 learns the MAC 00:01 in the FDB as usual and advertises it in EVPN without any IP. Optionally, the MAC can be configured as a CStatic mac, in which case it will be advertised as protected. If the MAC is learned on a SAP or SDP-binding where **auto-learn-mac-protect** is enabled, the MAC will also be advertised as protected.
3. The ARP-request is sent to the CPM where:
   – An ARP entry (IP 10.1'MAC 00:01) is populated into the proxy-ARP table.
   – EVPN advertises MAC 00:01 and IP 10.1 in EVPN with the same SEQ number and Protected bit as the previous route-type 2 for MAC 00:01.
   – A GARP is also issued to other SAPs/SDP-bindings (assuming they are not in the same split horizon group as the source). If garp-flood-evpn is enabled, the GARP message is also sent to the EVPN network.
   – The original ARP-request can still be flooded to the EVPN or not based on the **unknown-arp-request-flood-evpn** command.
4. Assuming PE1 was configured with **unknown-arp-request-flood-evpn**, the ARP-request is flooded to PE2 and delivered to ISP-B. ISP-B replies with its MAC in the ARP-reply. The ARP-reply is finally delivered to ISP-A.
5. PE2 will learn MAC 00:01 in the FDB and the entry 10.1'00:01 in the proxy-ARP table, based on the EVPN advertisements.
6. When ISP-B replies with its MAC in the ARP-reply:
   – MAC 00:03 is learned in FDB at PE2 and advertised in EVPN.
   – MAC 00:03 and IP 10.3 are learned in the proxy-ARP table and advertised in EVPN with the same SEQ number as the previous MAC route.
   – ARP-reply is unicasted to MAC 00:01.
7. EVPN advertisements are used to populate PE1's FDB (MAC 00:03) and proxy-ARP (IP 10.3—>MAC 00:03) tables as mentioned in 5.

From this point onward, the PEs reply to any ARP-request for 00:01 or 00:03, without the need for flooding the message in the EVPN network. By replying to known ARP-requests / Neighbor Solicitations, the PEs help to significantly reduce the flooding in the network.

Use the following commands to customize proxy-ARP/proxy-ND behavior:

- **dynamic-arp-populate** and **dynamic-nd-populate**

Enables the addition of dynamic entries to the proxy-ARP or proxy-ND table (disabled by default). When executed, the system will populate proxy-ARP/ proxy-ND entries from snooped GARP/ARP/NA messages on SAPs/SDP-bindings in addition to the entries coming from EVPN (if EVPN is enabled). These entries will be shown as *dynamic*.

- **static <IPv4-address> <mac-address>** and **static <IPv4-address> <mac-address>** and **static <ipv6-address> <mac-address> {host|router}**

  Configures static entries to be added to the table.

Note: A static IP->MAC entry requires the addition of the MAC address to the FDB as either learned or CStatic (conditional static mac) in order to become active (*Status —> active*).

- **age-time <60..86400>** (seconds)

  Specifies the aging timer per proxy-ARP/proxy-ND entry. When the aging expires, the entry is flushed. The age is reset when a new ARP/GARP/NA for the same IP—>MAC is received.

- **send-refresh <120..86400>** (seconds)

  If enabled, the system will send ARP-request/Neighbor Solicitation messages at the configured time, so that the owner of the IP can reply and therefore refresh its IP—>MAC (proxy-ARP entry) and MAC (FDB entry).

- **table-size [1..16384]**

  Enables the user to limit the number of entries learned on a specified service. By default, the table-size limit is 250.

  The unknown ARP-requests, NS, or the unsolicited GARPs and NA messages can be configured to be flooded or not in an EVPN network with the following commands:

  – proxy-arp [no] unknown-arp-request-flood-evpn
  – proxy-arp [no] garp-flood-evpn
  – proxy-nd [no] unknown-ns-flood-evpn
  – proxy-nd [no] host-unsolicited-na-flood-evpn
  – proxy-nd [no] router-unsolicited-na-flood-evpn

- **dup-detect [anti-spoof-mac <mac-address>] window <minutes> num-moves <count> hold-down <minutes|max>**

  Enables a mechanism that detects duplicate IPs and ARP/ND spoofing attacks. The working of the **dup-detect** command can be summarized as follows:

  – Attempts (relevant to dynamic and EVPN entry types) to add the same IP (different MAC) are monitored for <*window*> minutes and when <count> is reached within that *window,* the proxy-ARP/proxy-ND entry for the IP is suspected and marked as *duplicate*. An alarm is also triggered.

- The condition is cleared when hold-down time expires (*max* does not expire) or a **clear** command is issued.
- If the **anti-spoof-mac** is configured, the proxy-ARP/proxy-ND offending entry's MAC is replaced by this <mac-address> and advertised in an unsolicited GARP/NA for local SAP or SDP-bindings and in EVPN to remote PEs.
- This mechanism assumes that the same **anti-spoof-mac** is configured in all the PEs for the same service and that traffic with destination **anti-spoof-mac** received on SAPs/SDP-bindings will be dropped. An ingress MAC filter has to be configured in order to drop traffic to the **anti-spoof-mac**.

Table 84 shows the combinations that will produce a **Status = Active** proxy-arp entry in the table. The system will only reply to proxy-ARP requests for active entries. Any other combination will result in a **Status = inActv** entry. If the service is not active, the proxy-arp entries will not be active either, regardless of the FDB entries

Note: A static entry is active in the FDB even when the service is down.

Table 84 **Proxy-arp Entry combinations**

Proxy-arp Entry Type	FDB Entry Type (for the same MAC)
Dynamic	learned
Static	learned
Dynamic	CStatic/Static
Static	CStatic/Static
EVPN	EVPN
Duplicate	—

When proxy-ARP/proxy-ND is enabled on services with all-active multi-homed ethernet-segments, a proxy-arp entry type 'EVPN' might be associated with a 'learned' FDB entry (because the CE can send traffic for the same MAC to all the multi-homed PEs in the ES). If that is the case, the entry will be inactive, as per Table 84.

5.4.1.1 Proxy-ARP/ND Periodic Refresh, Unsolicited Refresh and Confirm-Messages

When proxy-ARP/proxy-ND is enabled, the system starts populating the proxy table and responding to ARP-requests/NS messages. To keep the active IP->MAC entries alive and ensure that all the host/routers in the service update their ARP/ND caches, the system may generate the following three types of ARP/ND messages for a specified IP->MAC entry:

- Periodic refresh messages (ARP-requests or NS for a specified IP):

  These messages are activated by the *send-refresh* command and their objective is to keep the existing FDB and Proxy-ARP/ND entries alive, in order to minimize EVPN withdrawals and re-advertisements.
- Unsolicited refresh messages (unsolicited GARP or NA messages):

  These messages are sent by the system when a new entry is learned or updated. Their objective is to update the attached host/router caches.
- Confirm messages (unicast ARP-requests or unicast NS messages):

  These messages are sent by the system when a new MAC is learned for an existing IP. The objective of the confirm messages is to verify that a specified IP has really moved to a different part of the network and is associated with the new MAC. If the IP has not moved, it will force the owners of the duplicate IP to reply and cause *dup-detect* to kick in.

5.4.1.2 Proxy-ND and the Router Flag in Neighbor Advertisement messages

RFC 4861 describes the use of the (R) or "Router" flag in NA messages as follows:

- A node capable of routing IPv6 packets must reply to NS messages with NA messages where the R flag is set (R=1).
- Hosts must reply with NA messages where R=0.

The use of the "R" flag in NA messages impacts how the hosts select their default gateways when sending packets "off-link". Therefore, it is important that the proxy-ND function on the 7750 SR, 7450 ESS, or 7950 XRS must meet one of the following criteria:

a. Either provide the appropriate R flag information in proxy-ND NA replies
b. Flood the received NA messages if it cannot provide the appropriate R flag when replying

Due to the use of the "R" flag, the procedure for learning proxy-ND entries and replying to NS messages differs from the procedures for proxy-ARP in IPv4: the router or host flag will be added to each entry, and that will determine the flag to use when responding to a NS.

5.4.1.3 Procedure to Add the R Flag to a Specified Entry

The procedure to add the R flag to a specified entry is as follows:

- Dynamic entries are learned based on received NA messages. The R flag is also learned and added to the proxy-ND entry so that the appropriate R flag is used in response to NS requests for a specified IP.
- Static entries are configured as host or router as per the command **[no] static <ip-address> <ieee-address> {host | router}**.
- EVPN entries are learned from BGP and the command **evpn-nd-advertise {host | router}** determines the R flag added to them.
- In addition, the **evpn-nd-advertise {host | router}** command will indicate what static and dynamic IP->MAC entries the system will advertise in EVPN. If **evpn-nd-advertise router** is configured, the system should flood the received unsolicited NA messages for hosts. This is controlled by the **[no] host-unsolicited-na-flood-evpn** command. The opposite is also recommended so that the **evpn-nd-advertise host** is configured with the **router-unsolicited-na-flood-evpn**.

5.4.1.4 Proxy-ARP/ND Mac-List for Dynamic Entries

SR OS supports the association of configured MAC lists with a configured dynamic proxy-ARP or proxy-ND IP address. The actual proxy-ARP or proxy-ND entry is not created until an ARP or Neighbor Advertisement message is received for the IP and one of the MACs in the associated MAC-list. This is in accordance with IETF Draft *draft-ietf-bess-evpn-proxy-arp-nd.* which states that a proxy-ARP or proxy-ND IP entry can be associated to one MAC among a list of allowed MACs.

The following example shows the use of MAC lists for dynamic entries.

```
A:PE-2>config>service#
  proxy-arp-nd
    mac-list ISP-1 create
      mac 00:de:ad:be:ef:01
      mac 00:de:ad:be:ef:02
      mac 00:de:ad:be:ef:03

A:PE-2>config>service>vpls>proxy-arp#
```

```
  dynamic 1.1.1.1 create
    mac-list ISP-1
    resolve 30

A:PE-2>config>service>vpls>proxy-nd#
  dynamic 200::1 create
    mac-list ISP-1
    resolve 30
```

Where:

- A dynamic IP (**dynamic** *ip* **create**) is configured and associated to a MAC list (**mac-list** *name*).
- The MAC list is created in the **config**>**service** context and can be reused by multiple configured dynamic IPs as follows:
  - in different services
  - in the same service, for proxy-ARP and proxy-ND entries
- If the MAC list is empty, the proxy-ARP or proxy-ND entry is not created for the configured IP.
- The same MAC list can be applied to multiple configured dynamic entries even within the same service.
- The new proxy-ARP and proxy-ND entries behave as dynamic entries and are displayed as type **dyn** in the **show** commands.

  The following output sample displays the entry corresponding to the configured dynamic IP.

```
*A:PE-2# show service id 1 proxy-arp detail
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : 900 secs            Send Refresh      : 300 secs
Table Size        : 250                 Total             : 1
Static Count      : 0                   EVPN Count        : 0
Dynamic Count     : 1                   Duplicate Count   : 0
Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None
EVPN
-------------------------------------------------------------------------------
Garp Flood        : enabled             Req Flood         : enabled
Static Black Hole : disabled
-------------------------------------------------------------------------------
===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address       Type Status    Last Update
-------------------------------------------------------------------------------
1.1.1.1           00:de:ad:be:ef:01  dyn active    02/23/2016 09:05:49
```

```
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
*A:PE-2# show service proxy-arp-nd mac-list "ISP-1" associations
===============================================================================
MAC List Associations
===============================================================================
Service Id                    IP Addr
-------------------------------------------------------------------------------
1                             1.1.1.1
1                             200::1
-------------------------------------------------------------------------------
Number of Entries: 2
===============================================================================
```

Although no new proxy-ARP or proxy-ND entries are created when a dynamic IP is configured, the router triggers the following resolve procedure.

1. The router sends a resolve message with a configurable frequency of 1 to 60 minutes; the default value is 5 minutes.

   The resolve message is an ARP-request or NS message flooded to all the non-EVPN endpoints in the service.
2. The router sends resolve messages at the configured frequency until a dynamic entry for the IP is created.

Note: The dynamic entry is created only if an ARP, GARP, or NA message is received for the configured IP, and the associated MAC belongs to the configured MAC list of the IP. If the MAC list is empty, the proxy-ARP or proxy-ND entry is not created for the configured IP.

After a dynamic entry (with a MAC address included in the list) is successfully created, its behavior (for send-refresh, age-time, and other activities) is the same as a configured dynamic entry with the following exceptions.

- Regular dynamic entries may override configured dynamic entries, but static or EVPN entries cannot override configured dynamic entries.
- If the corresponding MAC is flushed from the FDB after the entry is successfully created, the entry becomes inactive in the proxy-ARP or proxy-ND table and the resolve process is restarted.
- If the MAC list is changed, all the IPs that point to the list delete the proxy entries and the resolve process is restarted.
- If there is an existing configured dynamic entry and the router receives a GARP, ARP, or NA for the IP with a MAC that is not contained in the MAC list, the message is discarded and the proxy-ARP or proxy-ND entry is deleted. The resolve process is restarted.

- If there is an existing configured dynamic entry and the router receives a GARP, ARP, or NA for the IP with a MAC contained in the MAC list, the existing entry is overridden by the IP and new MAC, assuming the confirm procedure passes.
- The dup-detect and confirm procedures work for the configured dynamic entries when the MAC changes are between MACs in the MAC list. Changes to an off-list MAC cause the entry to be deleted and the resolve process is restarted.

5.4.2 BGP-EVPN MAC-Mobility

EVPN defines a mechanism to allow the smooth mobility of MAC addresses from an NVE to another NVE. The 7750 SR, 7450 ESS, and 7950 XRS support this procedure as well as the MAC-mobility extended community in MAC advertisement routes as follows:

- The router honors and generates the SEQ (Sequence) number in the mac mobility extended community for mac moves.
- When a MAC is EVPN-learned and it is attempted to be learned locally, a BGP update is sent with SEQ number changed to "previous SEQ"+1 (exception: mac duplication num-moves value is reached).
- SEQ number = zero or no mac mobility **ext-community** are interpreted as sequence zero.
- In case of mobility, the following MAC selection procedure is followed:
  - If a PE has two or more active remote EVPN routes for the same MAC (VNI can be the same or different), the highest SEQ number is selected. The tie-breaker is the lowest IP (BGP NH IP).
  - If a PE has two or more active EVPN routes and it is the originator of one of them, the highest SEQ number is selected. The tie-breaker is the lowest IP (BGP NH IP of the remote route is compared to the local system address).

Note: When EVPN multi-homing is used in EVPN-MPLS, the ESI is compared to determine whether a MAC received from two different PEs has to be processed within the context of MAC mobility or multi-homing. Two MAC routes that are associated with the same remote or local ESI but different PEs are considered reachable through all those PEs. Mobility procedures are not triggered as long as the MAC route still belongs to the same ESI.

5.4.3 BGP-EVPN MAC-Duplication

EVPN defines a mechanism to protect the EVPN service from control plane churn as a result of loops or accidental duplicated MAC addresses. The 7750 SR, 7450 ESS, and 7950 XRS support an enhanced version of this procedure as described in this section.

A situation may arise where the same MAC address is learned by different PEs in the same VPLS because of two (or more hosts) being mis-configured with the same (duplicate) MAC address. In such situation, the traffic originating from these hosts would trigger continuous MAC moves among the PEs attached to these hosts. It is important to recognize such situation and avoid incrementing the sequence number (in the MAC Mobility attribute) to infinity.

To remedy such situation, a router that detects a MAC mobility event by way of local learning starts a **window <in-minutes>** timer (default value of window = 3) and if it detects **num-moves <num>** before the timer expires (default value of num-moves = 5), it concludes that a duplicate MAC situation has occurred. The router then alerts the operator with a trap message. The offending MAC address can be shown using the show service id x bgp-evpn command:

```
10 2014/01/14 01:00:22.91 UTC MINOR: SVCMGR #2331 Base
"VPLS Service 1 has MAC(s) detected as duplicates by EVPN mac-
duplication detection."
# show service id 1 bgp-evpn
===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
VXLAN Admin Status : Enabled            Creation Origin    : manual
MAC Dup Detn Moves : 5                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 9                  Number of Dup MACs : 1
-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses             Time Detected
-------------------------------------------------------------------------------
00:00:00:00:00:12                            01/14/2014 01:00:23
-------------------------------------------------------------------------------
===============================================================================
```

After detecting the duplicate, the router stops sending and processing any BGP MAC advertisement routes for that MAC address until one of the following occurs:

a. The MAC is flushed due to a local event (SAP or SDP-binding associated with the MAC fails) or the reception of a remote update with better SEQ number (due to a mac flush at the remote router).
b. The retry <in-minutes> timer expires, which will flush the MAC and restart the process.

Note: The other routers in the VPLS instance will forward the traffic for the duplicate MAC address to the router advertising the best route for the MAC.

The values of **num-moves** and window are configurable to allow for the required flexibility in different environments. In scenarios where BGP rapid-update evpn is configured, the operator might want to configure a shorter window timer than in scenarios where BGP updates are sent every (default) min-route-advertisement interval.

Mac-duplication is always enabled in EVPN-VXLAN VPLS services, and the preceding described mac duplication parameters can be configured per VPLS service under the **bgp-evpn mac-duplication** context:

```
*A:DGW1>config>service>vpls>bgp-evpn# info
----------------------------------------------
       mac-advertisement
       unknown-mac-route
       mac-duplication
         detect num-moves num window in_mins
         [no] retry in_mins
       vxlan
         no shutdown
       exit
```

5.4.4 Conditional Static MAC and Protection

RFC 7432 defines the use of the sticky bit in the mac-mobility extended community to signal static mac addresses. These addresses must be protected in case there is an attempt to dynamically learn them in a different place in the EVPN-VXLAN VPLS service.

In the 7750 SR, 7450 ESS, and 7950 XRS, any conditional static mac defined in an EVPN-VXLAN VPLS service will be advertised by BGP-EVPN as a static address, that is, with the sticky bit set. An example of the configuration of a conditional static mac is shown below:

```
*A:PE63>config>service>vpls# info
----------------------------------------------
            description "vxlan-service"
...
            sap 1/1/1:1000 create
            exit
            static-mac
                mac 00:ca:ca:ca:ca:00 create sap 1/1/1:1000 monitor fwd-status
            exit
            no shutdown
```

```
*A:PE64# show router bgp routes evpn mac hunt mac-address 00:ca:ca:ca:ca:00
...
===============================================================================
BGP EVPN Mac Routes
===============================================================================
Network        : 0.0.0.0/0
Nexthop        : 192.0.2.63
From           : 192.0.2.63
Res. Nexthop   : 192.168.19.1
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None
Connector      : None
Community      : target:65000:1000     mac-mobility:Seq: 0/Static
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.63
Flags          : Used  Valid  Best  IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : MAC
ESI            : 0:0:0:0:0:0:0:0:0:0    Tag            : 1063
IP Address     : ::                     RD             : 65063:1000
Mac Address    : 00:ca:ca:ca:ca:00      Mac Mobility   : Seq:0
Neighbor-AS    : N/A
Source Class   : 0                      Dest Class     : 0
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Local static MACs or remote MACs with sticky bit are considered as 'protected'. A packet entering a SAP / SDP-binding will be discarded if its source MAC address matches one of these 'protected' MACs.

5.4.5 Auto-Learn MAC Protect and Restricting Protected Source MACs

Auto-learn MAC protect, together with the ability to restrict where the protected source MACs are allowed to enter the service, can be enabled within an EVPN-MPLS and EVPN-VXLAN VPLS and routed VPLS services, but not in PBB-EVPN services. The protection, using the **auto-learn-mac-protect** command (described in Auto-Learn MAC Protect), and the restrictions, using the **restrict-protected-src [discard-frame]** command, operate in the same way as in a non-EVPN VPLS service.

- When **auto-learn-mac-protect** is enabled on an object, source MAC addresses learned on that object are marked as protected within the FDB.

- When **restrict-protected-src** is enabled on an object and a protected source MAC is received on that object, the object is automatically shutdown (requiring the operator to **shutdown** then **no shutdown** the object in order to make it operational again).
- When **restrict-protected-src discard-frame** is enabled on an object and a frame with a protected source MAC is received on that object, that frame is discarded.

In addition, the following behavioral differences are specific to EVPN services:

- An implicit **restrict-protected-src discard-frame** command is enabled by default on SAPs, mesh-SDPs and spoke-SDPs. As this is the default, it is not possible to configure this command in an EVPN service. This default state can be seen in the show output for these objects, for example on a SAP:

```
*A:PE# show service id 1 sap 1/1/9:1 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/9:1                  Encap             : q-tag
...
RestMacProtSrc Act : none (oper: Discard-frame)
```

- A **restrict-protected-src discard-frame** can be optionally enabled on EVPN-MPLS/VXLAN destinations within EVPN services. When enabled, frames that have a protected source MAC address are discarded if received on any EVPN-MPLS/VXLAN destination in this service, unless the MAC address is learned and protected on an EVPN-MPLS/VXLAN destination in this service. This is enabled as follows:

```
configure
   service
      vpls <service id>
           bgp-evpn
               mpls
                   [no] restrict-protected-src discard-frame
           vxlan vni <vni-id>
               [no] restrict-protected-src discard-frame
```

- Auto-learned protected MACs are advertised to remote PEs in an EVPN MAC/IP advertisement route with the sticky bit set.
- The source MAC protection action relating to the **restrict-protected-src [discard-frame]** commands also applies to MAC addresses learned by receiving an EVPN MAC/IP advertisement route with the sticky bit set from remote PEs. This causes remotely configured conditional static MACs and auto-learned protected MACs to be protected locally.

- In all-active multi-homing scenarios, if **auto-learn-mac-protect** is configured on all-active SAPs and **restrict-protected-src discard-frame** is enabled on EVPN-MPLS/VXLAN destinations, traffic from the CE that enters one multi-homing PE and needs to be switched through the other multi-homing PE will be discarded on the second multi-homing PE. Each multi-homing PE will protect the CE's MAC on its local all-active SAP, which results in any frames with the CE's MAC address as the source MAC being discarded as they are received on the EVPN-MPLS/VXLAN destination from the other multi-homing PE.

Conditional static MACs, EVPN static MACs and locally protected MACs are marked as protected within the FDB, as shown in the example output.

```
*A:PE# show service fdb-mac
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:00:00:00:00:01 sap:1/1/9:1              LP/30    01/05/16 11:58:22
1         00:00:00:00:00:02 vxlan:                   EvpnS:P  01/05/16 11:58:23
                            1.1.1.2:1
1         00:00:00:00:01:01 sap:1/1/9:1              CStatic: 01/04/16 20:05:02
                                                     P
1         00:00:00:00:01:02 vxlan:                   EvpnS:P  01/04/16 20:18:02
                            1.1.1.2:1
-------------------------------------------------------------------------------
No. of Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

In this output:

- the first MAC is locally protected using the **auto-learn-mac-protect** command
- the second MAC has been protected using the **auto-learn-mac-protect** command on a remote PE
- the third MAC is a locally configured conditional static MAC
- the fourth MAC is a remotely configured conditional static MAC

5.4.6 Black-hole MAC and its Application to Proxy-ARP/ Proxy-ND Duplicate Detection

A black-hole MAC is a local FDB record. It is similar to a conditional static MAC; it is associated with a “black-hole” (similar to a VPRN black-hole static-route in VPRNs) instead of a SAP or SDP-binding. A black-hole MAC can be added by using the following command:

```
config>service>vpls# static-mac mac
mac <ieee-address> [create] black-hole
```

The static black-hole MAC can have security applications (for example, replacement of MAC filters) for certain MACs. When used in combination with **restrict-protected-src**, the static black-hole MAC provides a simple and scalable way to filter MAC DA or SA in the data plane, regardless of how the frame arrived at the system (using SAP or SDP-bindings or EVPN endpoints).

For example, when a specified **static-mac mac 00:00:ca:fe:ca:fe create black-hole** is added to a service, the following behavior occurs:

- The configured MAC is created as a static MAC with a **black-hole** source identifier.

```
*A:PE1# show service id 1 fdb detail
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
-------------------------------------------------------------------------------
1         00:ca:ca:ba:ca:01 eES:                     Evpn     06/29/15 23:21:34
                            01:00:00:00:00:71:00:00:00:01
1         00:ca:ca:ba:ca:06 eES:                     Evpn     06/29/15 23:21:34
                            01:74:13:00:74:13:00:00:74:13
1         00:ca:00:00:00:00 sap:1/1/1:2              CStatic:P 06/29/15 23:20:58
1         00:ca:fe:ca:fe:00 black-hole               CStatic:P  06/29/15 23:20:00
1         00:ca:fe:ca:fe:69 eMpls:                   EvpnS:P    06/29/15 20:40:13
                            192.0.2.69:262133
-------------------------------------------------------------------------------
No. of MAC Entries: 5
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
```

- After it has been successfully added to the FDB, the black-hole MAC will be treated like any other protected MAC, as follows:
  - The black-hole MAC will be added as protected (CStatic:P) and advertised in EVPN as static.

  - SAP or SDP-bindings or EVPN endpoints where the **restrict-protected-src discard-frame** is enabled will discard frames where MAC SA is equal to black-hole MAC.
  - SAP or SDP-bindings where **restrict-protected-src** (**no discard-frame**) is enabled will go operationally down if a frame with MAC SA is equal to black-hole MAC is received.
- After the black-hole MAC has been successfully added to the FDB, any frame arriving at any SAP or SDP-binding or EVPN endpoint with MAC DA is equal to black-hole MAC will be discarded.

Black-hole MACs can also be used in services with **proxy-ARP**/**proxy-ND** enabled to filter traffic with destination to **anti-spoof-macs**. The **anti-spoof-mac** provides a way to attract traffic to a specified IP when a duplicate condition is detected for that IP address (see section ARP/ND Snooping and Proxy Support for more information); however, the system still needs to drop the traffic addressed to the **anti-spoof-mac** by using either a MAC filter or a black-hole MAC.

The user does not need to configure MAC filters when configuring a **static-black-hole** MAC address for the **anti-spoof-mac** function. To use a black-hole MAC entry for the **anti-spoof-mac** function in a proxy-ARP/proxy-ND service, the user needs to configure:

- the **static-black-hole** option for the **anti-spoof-mac**

```
*A:PE1# config>service>vpls>proxy-arp#
dup-detect window 3 num-moves 5 hold-down max anti-spoof-
mac 00:66:66:66:66:00 static-black-hole
```

- a static black-hole MAC using the same MAC address used for the **anti-spoof-mac**

```
*A:PE1# config>service>vpls#
static-mac mac 00:66:66:66:66:00 create black-hole
```

When this configuration is complete, the behavior of the **anti-spoof-mac** function changes as follows:

- In the EVPN, the MAC is advertised as Static. Locally, the MAC will be shown in the FDB as “CStatic” and associated with a **black-hole**.
- The combination of the **anti-spoof-mac** and the **static-black-hole** ensures that any frame that arrives at the system with MAC DA = **anti-spoof-mac** is discarded, regardless of the ingress endpoint type (SAP or SDP-binding or EVPN) and without the need for a filter.
- If, instead of discarding traffic, the user wants to redirect it using MAC DA as the **anti-spoof-mac**, then redirect filters should be configured on SAPs or SDP-bindings (instead of the **static-black-hole** option).

When the **static-black-hole** option is not configured with the **anti-spoof-mac**, the behavior of the **anti-spoof-mac** function, as described in ARP/ND Snooping and Proxy Support, remains unchanged. In particular:

- the **anti-spoof-mac** is not programmed in the FDB
- any attempt to add a Static MAC (or any other MAC) with the **anti-spoof-mac** value will be rejected by the system
- a MAC filter is needed to discard traffic with MAC DA = **anti-spoof-mac**.

5.4.7 Black-hole MAC for EVPN Loop Detection

SR OS can combine a black-hole MAC address concept and the EVPN MAC duplication procedures to provide loop protection in EVPN networks. The feature is compliant with the MAC mobility and multi-homing functionality in RFC 7432. The **config>service>vpls>bgp-evpn>mac-duplication>black-hole-dup-mac** CLI command enables the feature.

If enabled, there are no apparent changes in the MAC duplication; however, if a duplicated MAC is detected (for example, M1), then the router performs the following:

- adds M1 to the duplicate MAC list
- programs M1 in the FDB as a “Protected” MAC associated with a black-hole endpoint (where “type” is set to EvpnD:P and Source-Identifier is “black-hole”)

While the MAC type value remains EvpnD:P, the following additional operational details apply.

- Incoming frames with MAC DA = M1 are discarded by the ingress IOM, regardless of the ingress endpoint type (SAP, SDP, or EVPN), based on an FDB MAC lookup.
- Incoming frames with MAC SA = M1 are discarded by the ingress IOM or cause the router to bring down the SAP or SDP-binding, depending on the **restrict-protected-src** setting on the SAP, SDP, or EVPN endpoint.

The following sample CLI shows an example EVPN-MPLS service where **black-hole-dup-mac** is enabled and MAC duplication programs the duplicate MAC as a black-hole.

```
19 2016/12/20 19:45:59.69 UTC MINOR: SVCMGR #2331 Base
"VPLS Service 30 has MAC(s) detected as duplicates by EVPN mac-duplication detec
tion."
*A:PE-5# configure service vpls 30
*A:PE-5>config>service>vpls# info
----------------------------------------------
            bgp
```

```
            exit
            bgp-evpn
                evi 30
                mac-duplication
                    detect num-moves 3 window 3
                    retry 6
                    black-hole-dup-mac
                exit
                vxlan
                    shutdown
                exit
                mpls
                    ingress-replication-bum-label
                    auto-bind-tunnel
                        resolution any
                    exit
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            sap 1/1/1:30 create
                no shutdown
            exit
            spoke-sdp 56:30 leaf-ac create
                no shutdown
            exit
            no shutdown
-------------------------------------------
*A:PE-5# show service id 30 bgp-evpn
===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
CFM MAC Advertise  : Disabled
VXLAN Admin Status : Disabled           Creation Origin    : manual
MAC Dup Detn Moves : 3                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 6                  Number of Dup MACs : 1
MAC Dup Detn BH    : Enabled
IP Route Advert    : Disabled

EVI                : 30
Ing Rep Inc McastAd: Enabled
Accept IVPLS Flush : Disabled
Send EVPN Encap    : Enabled
-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses            Time Detected
-------------------------------------------------------------------------------
00:11:00:00:00:01                           12/20/2016 19:46:00
-------------------------------------------------------------------------------
<snip>
...
*A:PE-5# show service id 30 fdb detail
===============================================================================
Forwarding Database, Service 30
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
```

```
-------------------------------------------------------------------------------
30        00:11:00:00:00:01 black-hole              EvpnD:P  12/20/16 19:46:00
-------------------------------------------------------------------------------
No. of MAC Entries: 1
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

If the **retry** time expires, the MAC is flushed from the FDB and the process starts again. The **clear service id 30 evpn mac-dup-detect** {*ieee-address* | **all**} command clears the duplicate black-hole MAC address.

Note: The **clear service id 30 fdb** command clears learned MAC addresses; black-hole MAC addresses are not cleared.

Support for the **black-hole-dup-mac** command and the preceding associated loop detection procedures is as follows:

- not supported on B-VPLS, I-VPLS, M-VPLS, or R-VPLS services
- fully supported on EVPN-VXLAN and EVPN-MPLS VPLS services (including EVPN E-Tree)
- fully supported with EVPN MAC mobility and EVPN-MPLS multi-homing

5.4.8 CFM Interaction with EVPN Services

Ethernet Connectivity and Fault Management (ETH-CFM) allows the operator to validate and measure Ethernet Layer 2 services using standard IEEE 802.1ag and ITU-T Y.1731 protocols. Each tool performs a unique function and adheres to that tool's specific PDU and frame format and the associate rules governing the transmission, interception, and process of the PDU. Detailed information describing the ETH-CFM architecture, the tools, and various functions is located in the various OAM and Diagnostics guides and is not repeated here.

EVPN provides powerful solution architectures. ETH-CFM is supported in the various Layer 2 EVPN architectures. Since the destination Layer 2 MAC address, unicast or multicast, is ETH-CFM tool dependent (i.e. ETH-CC is sent as a L2 multicast and ETH-DM is sent as an L2 unicast), the ETH-CFM function is allowed to multicast and broadcast to the virtual EVPN connections. The Maintenance Endpoint (MEP) and Maintenance Intermediate Point (MIP) do not populate the local Layer 2 MAC Address forwarding database (FDB) with the MAC related to the MEP and MIP. This means that the 48-bit IEEE MAC address is not exchanged with peers and all ETH-CFM frames are broadcast across all virtual connections. To prevent the

flooding of unicast packets and allow the remote forwarding databases to learn the remote MEP and MIP Layer 2 MAC addresses, the command **cfm-mac-advertisement** must be configured under the **config>service>vpls>bgp-evpn** context. This allows the MEP and MIP Layer 2 IEEE MAC addresses to be exchanged with peers. This command will track configuration changes and send the required updates via the EVPN notification process related to a change.

Up MEP, Down MEP, and MIP creation is supported on the SAP, spoke, and mesh connections within the EVPN service. There is no support for the creation of ETH-CFM Management Points (MPs) on the virtual connection. VirtualMEP (vMEP) is supported with a VPLS context and the applicable EVPN Layer 2 VPLS solution architectures. The vMEP follows the same rules as the general MPs. When a vMEP is configured within the supported EVPN service, the ETH-CFM extraction routines are installed on the SAP, Binding, and EVPN connections within an EVPN VPLS Service. The vMEP extraction within the EVPN-PBB context requires the **vmep-extensions** parameter to install the extraction on the EVPN connections.

When MPs are used in combination with EVPN multi-homing, the following must be considered:

- Behavior of operationally down MEPs on SAPs/SDP-bindings with EVPN multi-homing:
  - All-active multi-homing: no ETH-CFM is expected to be used in this case, since the two (or more) SAPs/SDP-bindings on the PEs will be oper-up and active; however, the CE will have a single LAG and will respond as though it is connected to a single system. In addition to that, **cfm-mac-advertisement** can lead to traffic loops in all-active multi-homing.
  - Single-active multi-homing: operationally down MEPs defined on single-active ethernet-segment SAPs/SDP-bindings will not send any CCMs when the PE is non-DF for the ES and fault-propagation is configured. For single-active multi-homing, the behavior will be equivalent to MEPs defined on BGP-MH saps/binds.
- Behavior for operationally up MEPs on ES SAPs/SDP-bindings with EVPN multi-homing:
  - All-active multi-homing: operationally up MEPs defined on non-DF ES SAPs can send CFM packets. However, they cannot receive CCMs (the SAP is removed from the default multicast list) or unicast CFM packets (because the MEP MAC is not installed locally in the FDB; unicast CFM packets will be treated as unknown, and not sent to the non-DF SAP MEP).
  - Single-active multi-homing: operationally up MEPs should be able to send or receive CFM packets normally.
  - operationally up MEPs defined on LAG SAPs require the command process_cpm_traffic_on_sap_down so that they can process CFM when the LAG is down and act as regular Ethernet ports.

Due to the above considerations, the use of ETH-CFM in EVPN multi-homed SAPs/ SDP-bindings is only recommended on operationally down MEPs and single-active multi-homing. ETH-CFM is used in this case to notify the CE of the DF or non-DF status.

5.4.9 Configuring EVPN-VXLAN and EVPN-MPLS in the Same VPLS Service

When two BGP instances are added to a VPLS service, both BGP-EVPN MPLS and BGP-EVPN VXLAN can be configured at the same time in the service. A maximum of two BGP instances are supported in the same VPLS, such that BGP-EVPN MPLS can use BGP instance 1 or 2, and BGP-EVPN VXLAN can use BGP instance 1 only.

In a service where EVPN-VXLAN and EVPN-MPLS are configured together, the **config>service>vpls>bgp-evpn>mpls>bgp-instance 2** command allows the user to associate the BGP-EVPN MPLS to an instance different from BGP-EVPN VXLAN, and have both encapsulations simultaneously enabled in the same service. At the control plane level, MAC/IP routes received in one instance are consumed and re-advertised in the other instance as long as the route is the best route for a specific MAC or MAC/IP. Inclusive multicast routes are independently generated for each BGP instance. From a data plane perspective, the EVPN-MPLS and EVPN-VXLAN destinations are instantiated in different implicit Split Horizon Groups (SHGs) so that traffic can be forwarded between them.

The following example shows a VPLS service with two BGP instances and both VXLAN and MPLS encapsulations configured for the same BGP-EVPN service.

```
*A:PE-1>config>service>vpls# info
----------------------------------------------
  description "evpn-mpls and evpn-vxlan in the same service"
  vxlan vni 7000 create
  exit
  bgp
    route-distinguisher 10:2
    route-target target:64500:1
  exit
  bgp 2
    route-distinguisher 10:1
    route-target target:64500:1
  exit
  bgp-evpn
    evi 7000
    incl-mcast-orig-ip 12.12.12.12
    vxlan
      no shutdown
    mpls
      bgp-instance 2
      control-word
```

```
    auto-bind-tunnel
      resolution any
    exit
    force-vlan-vc-forwarding
    no shutdown
  exit
exit
no shutdown
```

The following list describe the preceding example:

- **bgp 1** or **bgp** is the default BGP instance
- **bgp 2** is the additional instance required when both **bgp-evpn vxlan** and **bgp-evpn mpls** are enabled in the service
- The commands supported in instance 1 are also available for the second instance with the following considerations.
  - **pw-template-binding**: the pw-template-binding can only exist in instance 1; it is not supported in instance 2.
  - **route-distinguisher**: the operating route-distinguisher in both bgp instances must be different
  - **route-target**: the route-target in both instances can be the same or different
  - **vsi-import** and **vsi-export**: import and export policies can also be defined for either bgp instance
- The **mpls bgp-instance 2** command assigns the second BGP instance to the MPLS; VXLAN always uses instance 1.

Note: The **bgp-evpn vxlan no shutdown** command is only allowed if **bgp-evpn shutdown** is configured, or if the **bgp-instance** associated with the MPLS has a different route distinguisher than the VXLAN instance (and vice versa).

The following features are not supported when two BGP instances are enabled on the same VPLS service:

- SDP-bindings
- R-VPLS, M-VPLS, I-VPLS, B-VPLS, or E-Tree VPLS
- Proxy-ARP and Proxy-ND
- BGP Multi-homing
- Assisted-Replication with leaf configuration
- IGMP and PIM snooping

5.4.9.1 BGP-EVPN Routes in Services Configured With Two BGP Instances

From a BGP perspective, the two BGP instances configured in the service are independent of each other. The redistribution of routes between the BGP instances is resolved at the EVPN application layer.

By default, if BGP-EVPN VXLAN and BGP-EVPN MPLS are both enabled in the same service, BGP will send the generated EVPN routes twice: once with the RFC 5512 BGP encapsulation extended community set to VXLAN and a second time with the encapsulation type set to MPLS.

Usually, a DC gateway will peer a pair of Route-Reflectors (RRs) in the DC and a pair of RRs in the WAN. For this reason, the user needs to add router policies so that EVPN-MPLS routes are only sent to the WAN RRs and EVPN-VXLAN routes are only sent to the DC RRs. The following examples show how you can configure router policies.

```
config>router>bgp#
 vpn-apply-import
 vpn-apply-export
 group "WAN"
  family evpn
  type internal
  export "allow only mpls"
  neighbor 192.0.2.6
 group "DC"
  family evpn
  type internal
  export "allow only vxlan"
  neighbor 192.0.2.2
config>router>policy-options# info
----------------------------------------------
         community "vxlan" members "bgp-tunnel-encap:VXLAN"
         community "mpls" members "bgp-tunnel-encap:MPLS"
            policy-statement "allow only mpls"
                entry 10
                  from
                    family evpn
                    community vxlan
                  action drop
                  exit
                exit
            exit
            policy-statement "allow only vxlan"
                entry 10
             from
              family evpn
              community mpls
                    action drop
                    exit
                exit
            exit
```

In a BGP instance, the EVPN routes are imported based on the route-targets and regular BGP selection procedures, regardless of their encapsulation.

The BGP-EVPN routes are generated and redistributed between BGP instances based on the following rules.

- Auto-discovery (AD) routes (type 1) are not generated by services with two BGP instances. However, AD routes are received from the EVPN-MPLS peers and processed for aliasing and backup functions as usual.
- MAC/IP routes (type 2) received in one of the two BGP instances are imported and the MACs added to the FDB according to the existing selection rules. If the MAC is installed in the FDB, it is re-advertised in the other BGP instance with the new BGP attributes corresponding to the BGP instance (route-target, route-distinguisher, and so on). The following considerations apply to these routes.
  - The **mac-advertisement** command governs the advertisement of any MACs (even those learned from BGP).
  - A MAC route is redistributed only if it is the best route based on the EVPN selection rules.
  - If a MAC route is the best route and has to be redistributed, the MAC/IP information, along with the MAC Mobility Extended Community, is propagated in the redistribution.
  - The router redistributes any MAC route update for which any attribute has changed. For example, a change in the SEQ or sticky bit in one instance is updated in the other instance for a route that is selected as the best MAC route.
- EVPN inclusive multicast routes are generated independently for each BGP instance with the corresponding BGP encapsulation extended community (VXLAN or MPLS). Also, the following considerations apply to these routes.
  - Ingress Replication (IR) and Assisted Replication (AR) routes are supported in the EVPN-VXLAN BGP instance. If AR is configured, the AR IP address must be a loopback address different from the **system-ip** and the configured **originating-ip** address.
  - The IR, P2MP mLDP, and composite inclusive multicast routes are supported in the EVPN-MPLS BGP instance.
  - The modification of the **incl-mcast-orig-ip** command is supported, subject to the following considerations.
    - The configured IP in the **incl-mcast-orig-ip** command is encoded in the **originating-ip** field of the inclusive multicast Routes for IR, P2MP, and composite routes.
    - The **originating-ip** field of the AR routes is still derived from the **service>system>vxlan>assisted-replication-ip** configured value.

– EVPN handles the inclusive multicast routes in a service based on the following rules.
  - For IR routes, the EVPN destination is set up based on the NLRI next hop.
  - For P2MP mLDP routes, the PMSI Tunnel Attribute **tunnel-id** is used to join the mLDP tree.
  - For composite P2MP-IR routes, the PMSI Tunnel Attribute **tunnel-id** is used to join the tree and create the P2MP bind. The NLRI next-hop is used to build the IR destination.
  - For AR routes, the NLRI next-hop is used to build the destination.
  - The following applies if a router receives two inclusive multicast routes in the same instance.

    - If the routes have the same **originating-ip** but different route-distinguishers and next-hops, the router processes both routes. In the case of IR routes, it sets up two destinations: one to each next-hop.

    - If the routes have the same **originating-ip**, different route distinguishers, but same next hops, the router sets up only one binding for IR routes.

    - The router ignores inclusive multicast routes received with its own **originating-ip**, regardless of the route-distinguisher.

- IP-Prefix routes (type 5) are not generated or imported by a service with two BGP instances.

5.4.9.2 Anycast Redundant Solution for Dual BGP Instance Services

Figure 192 shows the Anycast mechanism used to support gateway redundancy for dual BGP instance services. The example shows two redundant DC gateways (DC GWs) where the VPLS services contain two BGP instances: one each for EVPN-VXLAN and EVPN-MPLS.

Figure 192 **Multi-homed Anycast Solution**

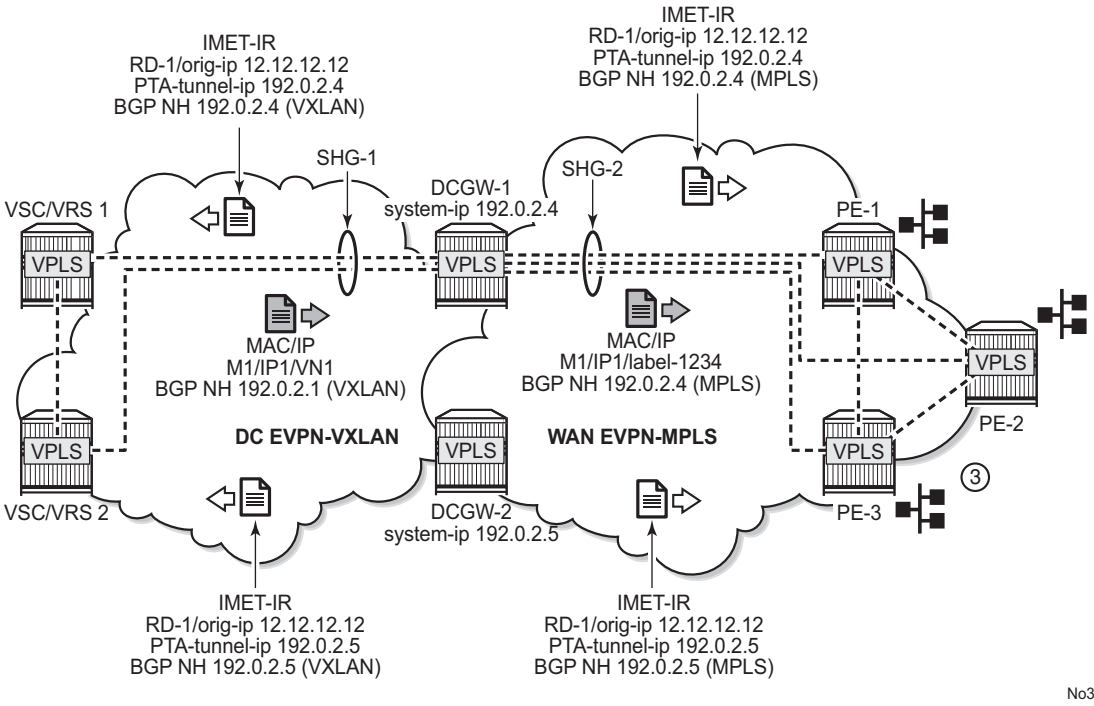

The example shown in Figure 192 depends on the ability of the two DC GWs to send the same inclusive multicast route to the remote PE or NVEs, such that:

- The remote PE or NVEs create a single BUM destination to one of the DC GWs (because the BGP selects only the best route to the DC GWs).
- The DC GWs do not create a destination between each other.

This solution avoids loops for BUM traffic, and known unicast traffic can use either DC GW router, depending on the BGP selection. The following CLI example output shows the configuration of each DC GW.

```
/* bgp configuration on DCGW1 and DCGW2 */
config>router>bgp#
 group "WAN"
  family evpn
  type internal
  neighbor 192.0.2.6
 group "DC"
  family evpn
  type internal
  neighbor 192.0.2.2
/* vpls service configuration */
DCGW-1# config>service>vpls(1)#
-----------------------
bgp
```

```
  route-distinguisher 64501:12
  route-target target:64500:1
exit
bgp 2
  route-distinguisher 64502:12
  route-target target:64500:1
exit
vxlan vni 1 create
exit
bgp-evpn
  evi 1
  incl-mcast-orig-ip 12.12.12.12
  vxlan
   no shutdown
  mpls
   bgp-instance 2
   no shutdown
   auto-bind-tunnel
     resolution any
   exit
<snip>
DCGW-2# config>service>vpls(1)#
-----------------------
bgp
  route-distinguisher 64501:12
  route-target target:64500:1
exit
bgp 2
  route-distinguisher 64502:12
  route-target target:64500:1
exit
vxlan vni 1 create
exit
bgp-evpn
  evi 1
  incl-mcast-orig-ip 12.12.12.12
  vxlan
   no shutdown
  mpls
   bgp-instance 2
   no shutdown
   auto-bind-tunnel
     resolution any
<snip>
```

Based on the preceding configuration example, the DC GWs behavior in this scenario is as follows:

- DCGW-1 and DCGW-2 send inclusive multicast routes to the DC RR and WAN RR with the same route key. For example:
  - DCGW-1 and DCGW-2 both send an IR route to DC RR with RD=64501:12, orig-ip=12.12.12.12, and a different next-hop and tunnel ID
  - DCGW-1 and DCGW-2 both send an IR route to WAN RR with RD=64502:12, orig-ip=12.12.12.12, and different next-hop and tunnel ID

- DCGW-1 and DCGW-2 both receive MAC/IP routes from DC and WAN that will be redistributed to the other BGP instances, assuming that the route is selected as best route and the MAC is installed in the FDB
  - As described in section BGP-EVPN Routes in Services Configured With Two BGP Instances, router peer policies are required so that only VXLAN or MPLS routes are sent or received for a specific peer
- Configuration of the same **incl-mcast-orig-ip** address in both DCGWs enables the Anycast solution for BUM traffic due to the following reasons.
  - The configured **originating-ip** is not required to be a reachable IP address and this forces the remote PE or NVEs to select only one of the two DC GWs.
  - The BGP next-hops are allowed to be the **system-ip** or even a loopback address. In both cases, the BGP next-hops are not required to be reachable in their respective networks.

In the example shown in Figure 192, PE-1 will pick up DC GW-1's inclusive multicast route (because of its lower BGP next-hop) and create a BUM destination to 192.0.2.4. When sending BUM traffic for VPLS-1, it will only send the traffic to DC GW-1. In the same way, the DC GWs will not set up BUM destinations between each other as they use the same **originating-ip** in their inclusive multicast routes.

The remote PE or NVEs perform a similar BGP selection for MAC or IP routes, as a specific MAC is sent by the two DC GWs with the same route-key. A PE or NVE will send known unicast traffic for a specific MAC to only one DC GW.

5.4.9.3 Using P2MP mLDP in Redundant Anycast DC GWs

Figure 193 shows an example of a common BGP EVPN service configured in redundant Anycast DC GWs and mLDP used in the MPLS instance.

Note: Packet duplication may occur if the service configuration is not performed carefully.

Figure 193 **Anycast Multi-homing and mLDP**

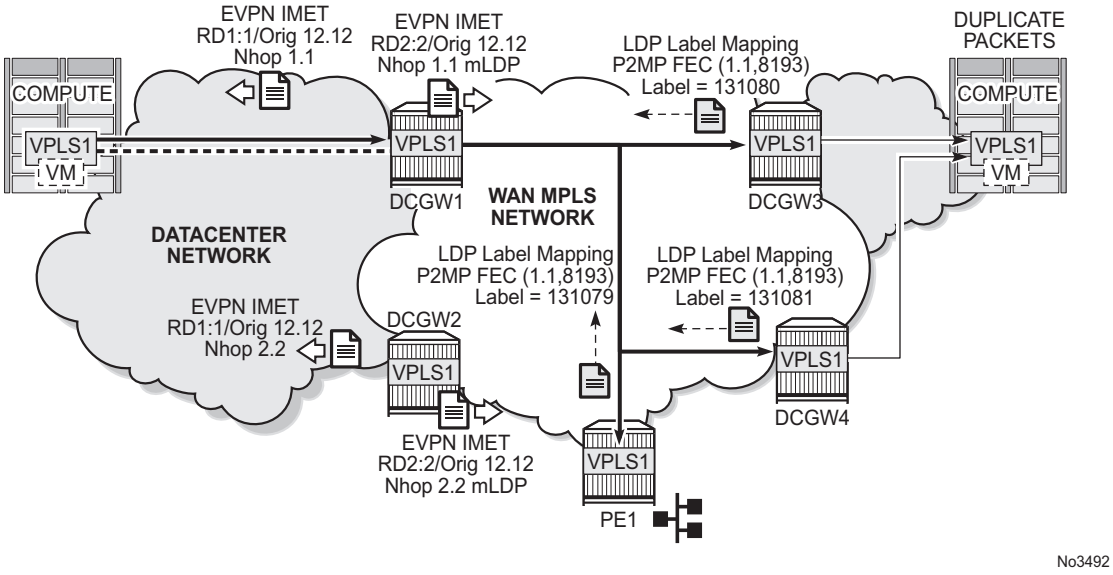

When mLDP is used with multiple Anycast multi-homing DC GWs, the same originating IP address must be used by all the DC GWs. Failure to do so may result in packet duplication.

In the example shown in Figure 193, each pair of DC GWs (DCGW1/DCGW2 and DCGW3/DCGW4) is configured with a different originating IP address, which causes the following behavior.

- DCGW3 and DCGW4 receive the inclusive multicast routes with the same route key from DCGW1 and DCGW2.
- Both DC GWs (DCGW3 and DCGW4) select only one route, which is generally the same, for example, DCGW1's inclusive multicast route.
- As a result, DCGW3 and DCGW4 join the mLDP tree with root in DCGW1, creating packet duplication when DCGW1 sends BUM traffic.
- Remote PE nodes with a single BGP-EVPN instance join the mLDP tree without any problem.

To avoid the packet duplication shown in Figure 193, Nokia recommends to configure the same originating IP address in all four DC GWs (DCGW1/DCGW2 and DCGW3/DCGW4). However, the route-distinguishers can be different per pair.

The following behavior occurs if the same originating IP address is configured on the DC GW pairs shown in Figure 193.

Note: This configuration allows the use of mLDP as long as BUM traffic is not required between the two DCs. Ingress Replication must be used if BUM traffic between the DCs is required.

- DCGW3 and DCGW4 do not join any mLDP tree sourced from DCGW1 or DCGW2, which prevents any packet duplication. This is because a router will ignore inclusive multicast routes received with its own **originating-ip**, regardless of the route-distinguisher.
- PE1 joins the mLDP trees from the two DCs.

5.4.9.4 Interconnect Ethernet-Segment Solution for Dual BGP Instance Services

SR OS supports Interconnect ESs (I-ES) for VXLAN as per IETF Draft *draft-ietf-bess-dci-evpn-overlay*. An I-ES is a virtual ES that allows DC GWs with two BGP instances to handle VXLAN access networks as any other type of ES. I-ESs support the RFC 7432 multi-homing functions, including single-active and all-active, ESI-based split-horizon filtering, DF election, and aliasing and backup on remote EVPN-MPLS PEs.

In addition to the EVPN multi-homing features, the main advantages of the I-ES redundant solution compared to the redundant solution described in Anycast Redundant Solution for Dual BGP Instance Services are as follows.

- The use of I-ES for redundancy in dual BGP-instance services allows local SAPs on the DCGWs.
- P2MP mLDP can be used to transport BUM traffic between DCs that use I-ES without any risk of packet duplication. As described in Using P2MP mLDP in Redundant Anycast DC GWs, packet duplication may occur in the Anycast DC GW solution when mLDP is used in the WAN.

Where EVPN-MPLS networks are interconnected to EVPN-VXLAN networks, the I-ES concept applies only to the access VXLAN network; the EVPN-MPLS network does not modify its existing behavior.

Figure 194 shows the use of I-ES for Layer 2 EVPN DCI between VXLAN and MPLS networks.

Figure 194 **The Interconnect ES Concept**

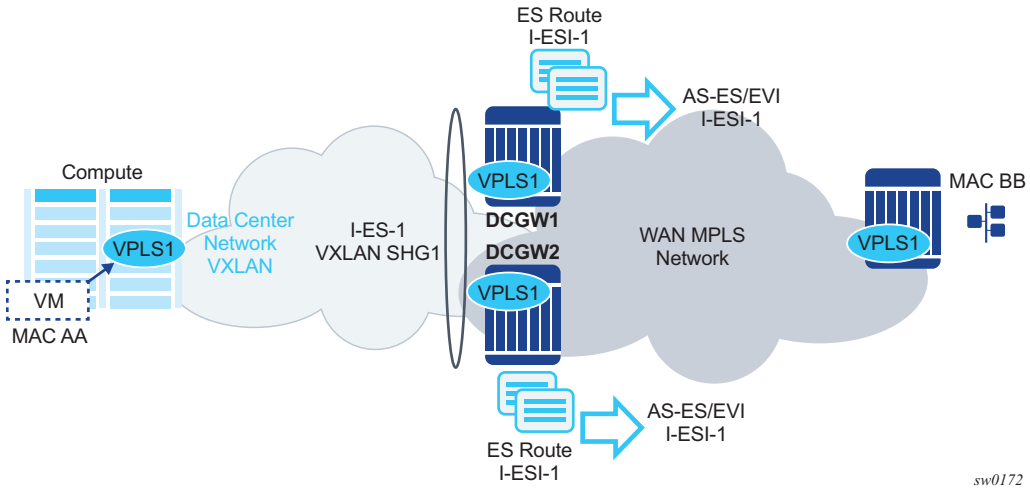

The following example shows how I-ES-1 would be provisioned on DC GW1 and the association between I-ES to a given VPLS service. A similar configuration would occur on DC GW2 in the I-ES.

New I-ES configuration:

```
DCGW1#config>service>system>bgp-evpn#
ethernet-segment I-ES-1 virtual create
 esi 01:00:00:00:12:12:12:12:12:00
 service-carving
  mode auto
 multi-homing all-active
 network-interconnect-vxlan 1
 service-id
   service-range 1 to 1000
 no shutdown
```

Service configuration:

```
DCGW1#config>service>vpls(1)#
 vxlan vni 1 instance 1 create
 exit
 bgp route-distinguisher 1:1
 bgp 2 route-distinguisher 2:2
 bgp-evpn
  evi 1
  vxlan
   no shutdown
  exit
  mpls
   bgp-instance 2
   auto-bind-tunnel resolution any
   no shutdown
```

```
...

DCGW1#config>service>vpls(2)#
 vxlan vni 2 instance 1 create
 exit
 bgp route-distinguisher 3:3
 bgp 2 route-distinguisher 4:4
 bgp-evpn
  evi 2
  vxlan
   no shutdown
  exit
  mpls
   bgp-instance 2
   auto-bind-tunnel resolution any
 sap 1/1/1:1 create
 exit
```

The above configuration associates I-ES-1 to the VXLAN instance in services VPLS1 and VPLS 2. The I-ES is modeled as a virtual ES, with the following considerations.

- The commands **network-interconnect-vxlan** and **service-id service-range** *svc-id* [**to** *svc-id*] are required within the Ethernet segment.
  - The **network-interconnect-vxlan** parameter identifies the VXLAN instance associated with the virtual ES. The value of the parameter must be set to 1. This command is rejected in a non-virtual ES.
  - The **service -range** parameter associates the specific service range to the ES. The ES must be configured as **network-interconnect-vxlan** before any service range can be added.
  - The ES parameters **port**, **lag**, **sdp**, **vc-id-range**, **dot1q**, and **qinq** cannot be configured in the Ethernet segment when a **network-interconnect-vxlan** instance is configured. The **source-bmac-lsb** option is blocked, as the I-ES cannot be associated with an I-VPLS or PBB-Epipe service. The remaining ES configuration options are supported.
  - All services with two BGP instances associate the VXLAN destinations and ingress VXLAN instances to the Ethernet segment.
- Multiple services can be associated with the same Ethernet segment, with the following considerations.
  - In a DC with two DC GWs (as in Figure 194), only two I-ESs are needed to load-balance, where one half of the dual BGP-instance services would be associated with one I-ES (for example, I-ES-1, in the above configuration) and one half to another I-ES.
  - Up to eight service ranges per VXLAN instance can be configured. Ranges may overlap within the same ES, but not between different ESs.
  - The service range can be configured before the service.

- After the I-ES is configured using **network-interconnect-vxlan**, the ES operational state depends exclusively on the ES administrative state. Because the I-ES is not associated with a physical port or SDP, when testing the non-revertive service carving manual mode, an Ethernet segment **shutdown** and **no shutdown** event results in the node sending its own administrative preference and DP bit and taking control if the preference and DP bit are higher than the current DF. This is because the peer ES routes are not present at the EVPN application layer when the ES is configured for **no shutdown**; therefore, the PE sends its own administrative preference and DP values. Therefore, for I-ESs, the non-revertive mode works only for node failures.
- A VXLAN instance may be placed in MhStandby under any of the following situations:
  - if the PE is single-active NDF for that I-ES
  - if the VXLAN service is added to the I-ES and either the ES or BGP-EVPN MPLS is shut down in all the services included in the ES

    The following example shows the change of the MhStandby flag from false to true when BGP-EVPN MPLS is shut down for all the services in the I-ES.

```
A:PE-4# show service id 500 vxlan instance 1 oper-flags
===============================================================================
VPLS VXLAN oper flags
===============================================================================
MhStandby                          : false
===============================================================================
A:PE-4# configure service vpls 500 bgp-evpn vxlan shutdown
*A:PE-4# show service id 500 vxlan instance 1 oper-flags
===============================================================================
VPLS VXLAN oper flags
===============================================================================
MhStandby                          : true
===============================================================================
```

5.4.9.4.1 BGP-EVPN Routes on Dual BGP-instance Services with I-ES

The configuration of an I-ES on DC GWs with two BGP-instances has the following impact on the advertisement and processing of BGP-EVPN routes.

- For EVPN MAC/IP routes, the following considerations apply.
  - MAC/IP routes received in the EVPN-MPLS BGP-instance are re-advertised in the EVPN-VXLAN BGP-instance with the ESI set to zero.
  - EVPN-VXLAN PEs and NVEs in the DC receive the same MAC from two or more different MAC/IP routes from the DC GWs, which perform regular EVPN MAC/IP route selection.

  - MAC/IP routes received in the EVPN-VXLAN BGP-instance are re-advertised in the EVPN-MPLS BGP-instance with the configured non-zero I-ESI value, assuming the VXLAN instance is not in an MhStandby operational state; otherwise the MAC/IP routes are dropped.
  - EVPN-MPLS PEs in the WAN receive the same MAC from two or more DC GWs, set with the same ESI. In this case, regular aliasing and backup functions occur as usual.
- ES routes are exchanged for the I-ES. The routes should be sent only to the MPLS network and not to the VXLAN network. This can be achieved by using router policies.
- AD per-ES and AD per-EVI are also advertised for the I-ES, and should be sent only to the MPLS network and not to the VXLAN. As for ES routes, router polices can be used to prevent AD routes from being sent to VXLAN peers.

In general, when I-ESs are used for redundancy, the use of router policies is needed to avoid control plane loops with MAC/IP routes. Consider the following to avoid control plane loops:

- Loops created by remote MACs

  Remote EVPN-MPLS MAC/IP routes are re-advertised into EVPN-VXLAN routes with an SOO (Site Of Origin) EC added by a BGP peer or VSI export policy identifying the DC GW pair. The other DC GW in the pair drops EVPN-VXLAN MAC routes tagged with the pair SOO. Router policies are needed to add SOO and drop routes received with self SOO.

  When remote EVPN-VXLAN MAC/IP routes are re-advertised into EVPN-MPLS, the DC GWs automatically drop EVPN-MPLS MAC/IP routes received with their own non-zero I-ESI.
- Loops created by local SAP MACs

  Local SAP MACs are learned and MAC/IP routes are advertised into both BGP instances. The MAC/IP routes advertised in the EVPN-VXLAN instance are dropped by the peer based on the SOO router policies as described above for loops created by remote MACs. The DC GW local MACs are always learned over the EVPN-MPLS destinations between the DC GWs.

The following outlines the considerations for BGP peer policies on DC GW1 to avoid control plane loops. Similar policies would be configured on DC GW2.

- Avoid sending service VXLAN routes to MPLS peers and service MPLS routes to VXLAN peers.
- Avoid sending AD and ES routes to VXLAN peers.
- Add SOO to VXLAN routes sent to the ES peer.
- Drop VXLAN routes received from the ES peer.

The following shows the CLI configuration.

```
A:DCGW1# configure router bgp
A:DCGW1>config>router>bgp# info
----------------------------------------------
            family vpn-ipv4 evpn
            vpn-apply-import
            vpn-apply-export
            rapid-withdrawal
            rapid-update vpn-ipv4 evpn
            group "wan"
                type internal
                export "allow only mpls"
                neighbor 192.0.2.4
                exit
                neighbor 192.0.2.5
                exit
            exit
            group "internal"
                type internal
                neighbor 192.0.2.1
                    export "allow only vxlan"
                exit
                neighbor 192.0.2.3
                    import "drop SOO-DCGW-23"
                    export "add SOO to vxlan routes"
                exit
            exit
            no shutdown
----------------------------------------------
A:DCGW1>config>router>bgp# /configure router policy-options
A:DCGW1>config>router>policy-options# info
----------------------------------------------
            community "mpls" members "bgp-tunnel-encap:MPLS"
            community "vxlan" members "bgp-tunnel-encap:VXLAN"
            community "SOO-DCGW-23" members "origin:64500:23"


//
 This policy prevents the router from sending service VXLAN routes to MPLS peers. //

            policy-statement "allow only mpls"
                entry 10
                    from
                        community "vxlan"
                        family evpn
                    exit
                    action drop
                    exit
                exit
            exit


//
 This policy ensures the router only exports routes that include the VXLAN encapsula
tion. //

            policy-statement "allow only vxlan"
                entry 10
```

```
                    from
                        community "vxlan"
                        family evpn
                    exit
                    action accept
                    exit
                exit
                default-action drop
                exit
            exit

// This import policy avoids importing routes with a self SOO. //

            policy-statement "drop SOO-DCGW-23"
                entry 10
                    from
                        community "SOO-DCGW-23"
                        family evpn
                    exit
                    action drop
                    exit
                exit
            exit

//
 This import policy adds SOO only to VXLAN routes. This allows the peer to drop rout
es based on the SOO, without affecting the MPLS routes. //

            policy-statement "add SOO to vxlan routes"
                entry 10
                    from
                        community "vxlan"
                        family evpn
                    exit
                    action accept
                        community add "SOO-DCGW-23"
                    exit
                exit
                default-action accept
                exit
            exit
-----------------------------------------------
```

5.4.9.4.2 Single-Active Multi-Homing on I-ES

When an I-ES is configured as single-active and is **no shutdown** with at least one associated service, the DC GWs send ES and AD routes as for any ES and runs DF election as normal, based on the ES routes, with the candidate list being pruned by the AD routes.

Figure 195 shows the expected behavior for a single-active I-ES.

Figure 195 **Interconnect ES — Single-Active**

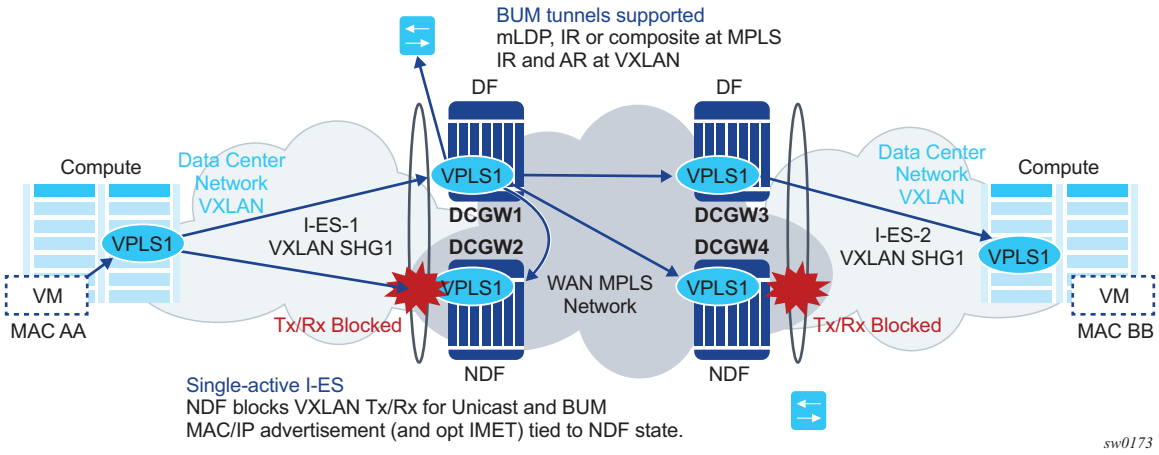

As shown in Figure 195, the Non-Designated-Forwarder (NDF) for a specified service carries out the following tasks.

- From a data path perspective, the VXLAN instance on the NDF goes into an MhStandby operational state and blocks ingress and egress traffic on the VXLAN destinations associated with the I-ES.
- MAC/IP routes and the FDB process
  - MAC/IP routes associated with the VXLAN instance and re-advertised to EVPN-MPLS peers are withdrawn.
  - MAC/IP routes corresponding to local SAP MACs or EVPN-MPLS binding MACs are withdrawn if they were advertised to the EVPN-VXLAN instance.
  - Received MAC/IP routes associated with the VXLAN instance are not installed in the FDB. MAC routes show as “used” in BGP, however, only the MAC route received from MPLS (from the ES peer) is programmed.
- Inclusive Multicast Ethernet Tag (IMET) routes process
  - IMET-AR-R routes (IMET-AR with replicator role) must be withdrawn if the VXLAN instance goes into an MhStandby operational state. Only the DF advertises the IMET-AR-R routes.
  - IMET-IR advertisements in the case of the NDF (or MhStandby) are controlled by the command **config>service>vpls>bgp-evpn>vxlan** [**no**] **send-imet-ir-on-ndf**.

    By default, the command is enabled and the router advertises IMET-IR routes, even if the PE is NDF (MhStandby). This attracts BUM traffic, but also speeds up convergence in the case of a DF switchover. The command is supported for single-active and all-active.

    If the command is disabled, the router withdraws the IMET-IR routes when the PE is NDF and will not attract BUM traffic.

The I-ES DF PE for the service continues advertising IMET and MAC/IP routes for the associated VXLAN instance as usual, as well as forwarding on the DF VXLAN bindings. When the DF DC GW receives BUM traffic, it sends the traffic with the egress ESI label if needed.

5.4.9.4.3 All-Active Multi-Homing on I-ES

The same considerations for ES and AD routes, and DF election apply for all-active multi-homing as for single-active multi-homing; the difference is in the behavior on the NDF DC GW. The NDF for a specified service performs the following tasks.

- From a data path perspective, the NDF blocks ingress and egress paths for broadcast and multicast traffic on the VXLAN instance bindings associated with the I-ES, while unknown and known unicast traffic is still allowed. The unknown unicast traffic is transmitted on the NDF if there is no risk of duplication. For example, unknown unicast packets are transmitted on the NDF if they do not have an ESI label, do not have an EVPN BUM label, and they pass a MAC SA suppression. In the example in Figure 196, the NDF transmits unknown unicast traffic. Regardless of whether DC GW1 is a DF or NDF, it accepts the unknown unicast packets and floods to local SAPs and EVPN destinations. When sending to DC GW2, the router sends the ESI-label identifying the I-ES. Due to the ESI-label suppression, DC GW2 does not send unknown traffic back to the DC.

Figure 196 **All-active Multi-Homing and Unknown Unicast on the NDF**

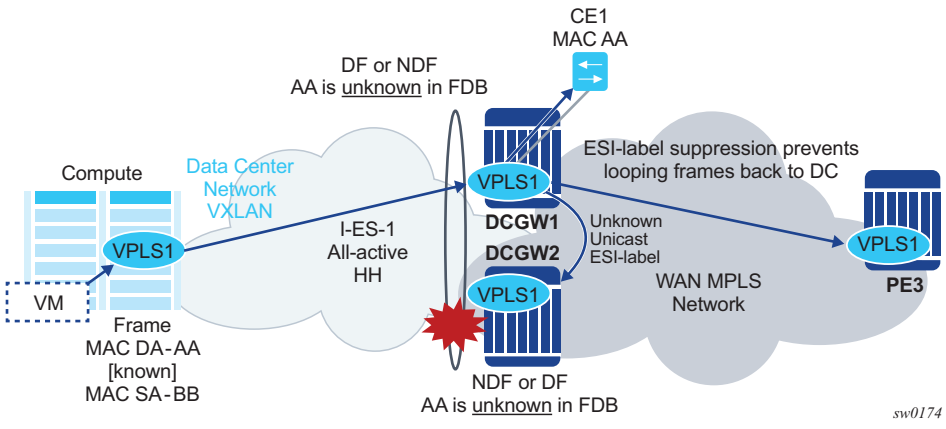

- MAC/IP routes and the FDB process
  - MAC/IP routes associated with the VXLAN instance are advertised normally.
  - MACs are installed as normal in the FDB for received MAC/IP routes associated with the VXLAN instance.

- IMET routes process
  - As is the case for single-active multi-homing, IMET-AR-R routes must be withdrawn on the NDF (MhStandby state). Only the DF advertises the IMET-AR-R routes.
  - The IMET-IR advertisements in the case of the NDF (or MhStandby) are controlled by the command **config**>**service**>**vpls**>**bgp-evpn**>**vxlan** [**no**] **send-imet-ir-on-ndf**, as in single-active multi-homing.

The I-ES DF PE for the service continues advertising IMET and MAC/IP routes for the associated VXLAN instance as usual. When the DF DC GW receives BUM traffic, it sends the traffic with the egress ESI label if needed.

5.4.10 BGP and EVPN Route Selection for EVPN Routes

When two or more EVPN routes are received at a PE, BGP route selection typically takes place when the route key or the routes are equal. When the route key is different, but the PE has to make a selection (for instance, the same MAC is advertised in two routes with different RDs), BGP will hand over the routes to EVPN and the EVPN application will perform the selection.

EVPN and BGP selection criteria are described below.

- EVPN route selection for MAC routes: when two or more routes with the same mac-length/mac but different route key are received, BGP will hand the routes over to EVPN. EVPN will select the route based on the following tie-break order:
  1. Conditional static MACs (local protected MACs)
  2. Auto-learned protected MACs (locally learned MACs on SAPs or mesh/spoke-SDPs due to the configuration of **auto-learn-mac-protect**)
  3. EVPN ES PBR MACs (see ES PBR MAC routes below)
  4. EVPN static MACs (remote protected MACs)
  5. Data plane learned MACs (regular learning on saps/sdp-bindings)
  6. EVPN MACs with higher SEQ number
  7. EVPN E-Tree root MACs
  8. Lowest IP (next-hop IP of the EVPN NLRI)
  9. Lowest eth-tag (that will be zero for MPLS and might be different from zero for VXLAN)
  10. Lowest RD
- ES PBR MAC routes: when a PBR filter with a forward action to an ESI and SF-IP (Service Function IP) exists, a MAC route is created by the system. This MAC route will be compared to other MAC routes received from BGP.

- When ARP resolves (it can be static, EVPN, or dynamic) for a SF-IP and the system has an AD EVI route for the ESI, a "MAC route" is created by ES PBR with the <MAC Address = ARPed MAC Address, VTEP = AD EVI VTEP, VNI = AD EVI VNI, RD = ES PBR RD (special RD), Static = 1> and installed in EVPN.
- This MAC route doesn't add anything (back) to ARP; however, it goes through the MAC route selection in EVPN and triggers the FDB addition if it is the best route.
- In terms of priority, this route's priority is lower than local static but higher than remote EVPN static (number 2 in the tie-break order above).
- If there are two competing ES PBR MAC routes, then the selection goes through the rest of checks (Lowest IP > Lowest RD).

- The BGP route selection for MAC routes with the same route-key follows the following priority order:
  1. EVPN static MACs (remote protected MACs).
  2. EVPN MACs with higher sequence number.
  3. Regular BGP selection (local-pref, aigp metric, shortest as-path, …, lowest IP).
- The BGP route selection for the rest of the EVPN routes: regular BGP selection is followed.

Note: In case BGP has to run an actual selection and a specified (otherwise valid) EVPN route 'loses' to another EVPN route, the non-selected route will be displayed by the **show router BGP routes evpn x detail** command with a 'tie-breaker' reason.

Note: Protected MACs do not overwrite EVPN static MACs; in other words, if a MAC is in the FDB and protected due to it being received with the sticky/static bit set in a BGP EVPN update and a frame is received with the source MAC on an object configured with **auto-learn-mac-protect**, that frame will be dropped due to the implicit **restrict-protected-src discard-frame**. The reverse is not true; when a MAC is learned and protected using **auto-learn-mac-protect**, its information is not overwritten with the contents of a BGP update containing the same MAC address.

5.4.11 Interaction of EVPN and Other Features

This section contains information about EVPN and how it interacts with other features.

5.4.11.1 Interaction of EVPN-VXLAN and EVPN-MPLS with Existing VPLS Features

When enabling existing VPLS features in an EVPN-VXLAN or an EVPN-MPLS enabled service, the following must be considered:

- EVPN-VXLAN services are not supported on I-VPLS/B-VPLS. VXLAN cannot be enabled on those services. EVPN-MPLS is only supported in regular VPLS and B-VPLS. Other VPLS types, such as **etree** or **m-vpls**, are not supported with either EVPN-VXLAN or EVPN-MPLS.
- In general, no router-generated control packets will be sent to the EVPN destination bindings, except for ARP, VRRP, ping, BFD and Eth-CFM for EVPN-VXLAN, and proxy-ARP/proxy-ND confirm messages and Eth-CFM for EVPN-MPLS.
- xSTP and M-VPLS services:
  - xSTP can be configured in bgp-evpn services. BPDUs will not be sent over the EVPN bindings.
  - **bgp-evpn** is blocked in m-vpls services; however, a different m-vpls service can manage a **SAP** or **spoke-sdp** in a **bgp-evpn** enabled service.
- **mac-move**—in **bgp-evpn** enabled VPLS services, **mac-move** can be used in **saps/sdp-bindings**; however, the MACs being learned through BGP-EVPN will not be considered.

Note: The mac duplication already provides a protection against mac-moves between EVPN and saps/sdp-bindings.

- **disable-learning** and other fdb-related tools—these will only work for data plane learned mac addresses.
- **mac-protec**t—**mac-protect** cannot be used in conjunction with EVPN.

Note: EVPN provides its own protection mechanism for static mac addresses.

- MAC OAM—MAC OAM tools are not supported for **bgp-evpn** services, that is: **mac-ping**, **mac-trace**, **mac-populate**, **mac-purge,** and **cpe-ping**.
- EVPN multi-homing and BGP-MH can be enabled in the same VPLS service, as long as they are not enabled in the same SAP-SDP or Spoke-SDP. There is no limitation on the number of BGP-MH sites supported per EVPN-MPLS service.

Note: The number of BGP-MH sites per EVPN-VXLAN service is limited to 1.

- SAPs/SDP-bindings that belong to a specified ES but are configured on non-bgp-evpn-mpls-enabled VPLS or Epipe services will be kept down with the **StandByForMHProtocol** flag.
- CPE-ping is not supported on EVPN services but it is in PBB-EVPN services (including I-VPLS and PBB-Epipe). CPE-ping packets will not be sent over EVPN destinations. CPE-ping will only work on local active SAP or SDP-bindings in I-VPLS and PBB-Epipe services.
- Other commands not supported in conjunction with **bgp-evpn**:
  - bgp-vpls
  - Endpoints and attributes
  - Subscriber management commands under service, SAP, and sdp-binding interfaces
  - MLD-snooping and attributes
  - BPDU translation
  - L2PT termination
  - MAC-pinning
- Other commands not supported in conjunction with **bgp-evpn mpls** are:
  - VSD-domains
  - SPB configuration and attributes

5.4.11.2 Interaction of PBB-EVPN with Existing VPLS Features

In addition to the B-VPLS considerations described in section Interaction of EVPN-VXLAN and EVPN-MPLS with Existing VPLS Features, the following specific interactions for PBB-EVPN should also be considered.

- When **bgp-evpn mpls** is enabled in a **b-vpls** service, an **i-vpls** service linked to that **b-vpls** cannot be an R-VPLS (the **allow-ip-int-bind** command is not supported).
- The ISID value of 0 is not allowed for PBB-EVPN services (I-VPLS and Epipes).
- The **ethernet-segments** can be associated with **b-vpls** SAPs/SDP-bindings and **i-vpls/epipe** SAPs/SDP-bindings,; however, the same ES cannot be associated with **b-vpls** and **i-vpls/epipe** SAP or SDP-bindings at the same time.

- When PBB-epipes are used with PBB-EVPN multi-homing, spoke-SDPs are not supported on **ethernet-segments**.
- When **bgp-evpn mpls** is enabled, eth-tunnels are not supported in the b-vpls instance.

5.4.11.3 Interaction of EVPN-VXLAN and EVPN-MPLS with Existing VPRN Features

When enabling existing VPRN features on interfaces linked to EVPN-VXLAN R-VPLS or EVPN-MPLS R-VPLS interfaces, consider that the following are not supported:

- the commands **arp-populate** and **authentication-policy**
- dynamic routing protocols such as IS-IS, RIP, and OSPF
- BFD on EVPN tunnel interfaces

5.4.11.4 Routing Policies for BGP EVPN IP Prefixes

BGP routing policies are supported for IP prefixes imported or exported through BGP-EVPN.

When applying routing policies to control the distribution of prefixes between EVPN and IP-VPN, the user must consider that both families are completely separate as far as BGP is concerned and that when prefixes are imported in the VPRN routing table, the BGP attributes are lost to the other family. The use of route tags allows the controlled distribution of prefixes across the two families.

Figure 197 shows an example of how VPN-IPv4 routes are imported into the RTM (Routing Table Manager), and then passed to EVPN for its own process.

Note: VPN-IPv4 routes can be tagged at ingress and that tag is preserved throughout the RTM and EVPN processing, so that the tag can be **matched** at the egress BGP routing policy.

Figure 197 **IP-VPN Import and EVPN Export BGP Workflow**

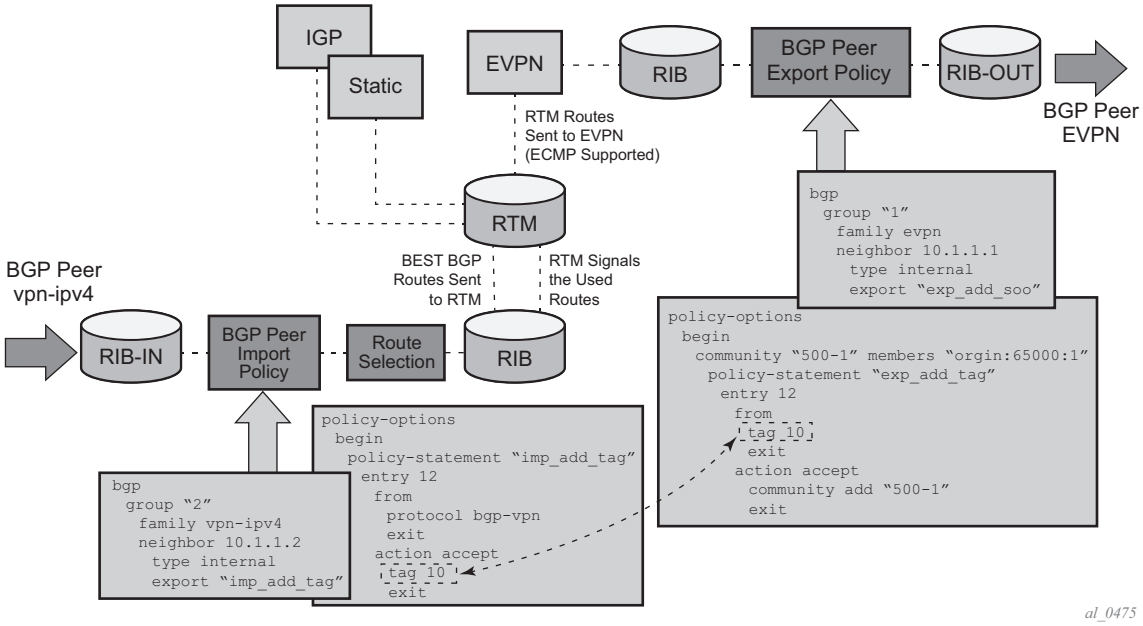

Policy tags can be used to match EVPN IP prefixes that were learned not only from BGP VPN-IPv4 but also from other routing protocols. The tag range supported for each protocol is different:

```
<tag>  : accepts in decimal or hex
         [0x1..0xFFFFFFFF]H (for OSPF and IS-IS)
         [0x1..0xFFFF]H (for RIP)
         [0x1..0xFF]H (for BGP)
```

Figure 198 shows an example of the reverse workflow: routes imported from EVPN and exported from RTM to BGP VPN-IPv4.

Figure 198 **EVPN Import and I-VPN Export BGP Workflow**

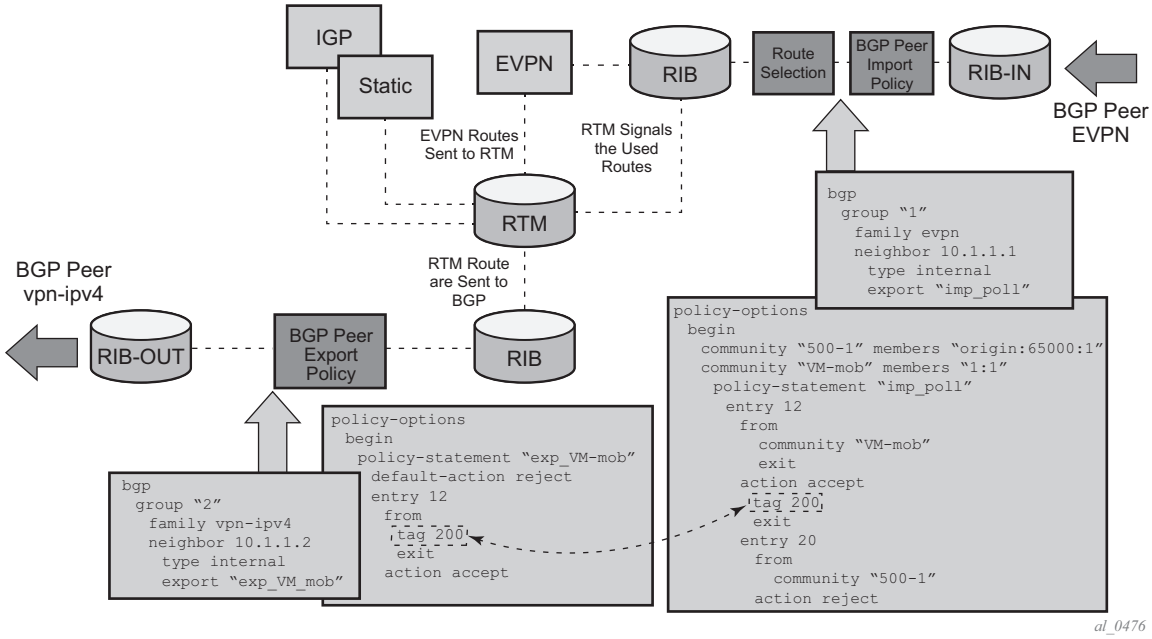

The preceding described behavior and the use of tags is also valid for vsi-import and vsi-export policies in the R-VPLS.

The policy behavior for EVPN ip-prefixes can then be summarized in the following statements.

- For EVPN prefix routes received and imported in RTM:
  - Policy entries can match on communities and add tags. This works at the peer level or at the vsi-import level.
  - Policy entries can match on *family evpn.*
- For exporting RTM to EVPN prefix routes:
  - Policy entries can match on tags and based on that, add communities, accept, or reject. This works at the peer level or the vsi-export level.

Policy entries can add tags for static-routes, RIP, OSPF, IS-IS, and BGP that can then be matched on the BGP peer export policy or vsi-export policy for EVPN prefix routes.

5.5 Configuring an EVPN Service with CLI

This section provides information to configure VPLS using the command line interface.

5.5.1 EVPN-VXLAN Configuration Examples

5.5.1.1 Layer 2 PE Example

This section shows a configuration example for three PEs in a Data Center, given the following assumptions:

- PE-1 is a Data Center Network Virtualization Edge device (NVE) where service VPLS 2000 is configured.
- PE-2 and PE-3 are redundant Data Center Gateways providing Layer 2 connectivity to the WAN for service VPLS 2000

DC PE-1 configuration for service VPLS 2000

DC PE-2 and PE-3 configuration with SAPs at the WAN side (advertisement of all macs and unknown-mac-route):

```
service vpls 2000 customer 1 create
    vxlan vni 2000 create
    bgp
        route-target 65000:2000
        route-distinguisher 65001:2000
    bgp-evpn
        mac-advertisement
        unknown-mac-route
        vxlan
            no shutdown
    site site-1 create
        sap 1/1/1:1
        no shutdown
        site-id 1
    sap 1/1/1:1 create
```

DC PE-2 and PE-3 configuration with BGP-AD spoke-SDPs at the WAN side (mac-advertisement disable, only unknown-mac-route advertised):

```
service vpls 2000 customer 1 create
    vxlan vni 2000 create
    bgp
        pw-template-binding 1 split-horizon-group "to-WAN" import-
```

```
rt target:65000:2500
        vsi-export "export-policy-1" #policy exporting the WAN and DC RTs
        vsi-import "import-policy-1" #policy importing the WAN and DC RTs
        route-distinguisher 65001:2000
    bgp-ad
        no shutdown
        vpls-id 65000:2000
    bgp-evpn
        mac-advertisement disable
        unknown-mac-route
        vxlan
            no shutdown
    site site-1 create
        split-horizon-group "to-WAN"
        no shutdown
        site-id 1
```

5.5.1.2 EVPN for VXLAN in R-VPLS Services Example

This section shows a configuration example for three 7750 SR, 7450 ESS, or 7950 XRS PEs in a Data Center, based on the following assumptions:

- PE-1 is a Data Center Network Virtualization Edge device (NVE) where the following services are configured:
  - R-VPLS 2001 and R-VPLS 2002 are subnets where Tenant Systems are connected
  - VPRN 500 is a VPRN instance providing inter-subnet forwarding between the local subnets and from local subnets to the WAN subnets
  - R-VPLS 501 is an IRB backhaul R-VPLS service that provides EVPN-VXLAN connectivity to the VPRNs in PE-2 and PE-3

```
*A:PE-1>config>service# info
        vprn 500 customer 1 create
            ecmp 4
            route-distinguisher 65071:500
            vrf-target target:65000:500
            interface "evi-501" create
              address 30.30.30.1/24
                vpls "evpn-vxlan-501"
                exit
            exit
            interface "subnet-2001" create
                address 10.10.10.1/24
                vpls "r-vpls 2001"
                exit
            exit
            interface "subnet-2002" create
                address 20.20.20.1/24
                vpls "r-vpls 2002"
                exit
            exit
            no shutdown
```

```
        exit
        vpls 501 customer 1 create
            allow-ip-int-bind
            vxlan vni 501 create
            exit
            bgp
                route-distinguisher 65071:501
                route-target export target:65000:501 import target:65000:501
          exit
            bgp-evpn
                ip-route-advertisement incl-host
                vxlan
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            service-name "evpn-vxlan-501"
            no shutdown
        exit
        vpls 2001 customer 1 create
            allow-ip-int-bind
            service-name "r-vpls 2001"
            sap 1/1/1:21 create
            exit
            sap 1/1/1:501 create
            exit
            no shutdown
        exit
        vpls 2002 customer 1 create
            allow-ip-int-bind
            service-name "r-vpls 2002"
            sap 1/1/1:22 create
            exit
            sap 1/1/1:502 create
            exit
            no shutdown
        exit
```

PE-2 and PE-3 are redundant Data Center Gateways providing Layer 3 connectivity to the WAN for subnets "subnet-2001" and "subnet-2002". The following configuration excerpt shows an example for PE-2. PE-3 would have an equivalent configuration.

```
*A:PE-2>config>service# info
        vprn 500 customer 1 create
            ecmp 4
            route-distinguisher 65072:500
            auto-bind mpls-gre
            vrf-target target:65000:500
            interface "evi-501" create
              address 30.30.30.2/24
                vpls "evpn-vxlan-501"
                exit
            exit
            no shutdown
```

```
        exit
        vpls 501 customer 1 create
            allow-ip-int-bind
            vxlan vni 501 create
            exit
            bgp
                route-distinguisher 65072:501
                route-target export target:65000:501 import target:65000:501
            exit
            bgp-evpn
                ip-route-advertisement incl-host
                vxlan
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            service-name "evpn-vxlan-501"
            no shutdown
        exit
```

5.5.1.3 EVPN for VXLAN in EVPN Tunnel R-VPLS Services Example

The example in EVPN for VXLAN in R-VPLS Services Example can be optimized by using EVPN tunnel R-VPLS services instead of regular IRB backhaul R-VPLS services. If EVPN tunnels are used, the corresponding R-VPLS services cannot contain SAPs or SDP-bindings and the VPRN interfaces will not need IP addresses.

The following excerpt shows the configuration in PE-1 for the VPRN 500. The R-VPLS 501, 2001 and 2002 can keep the same configuration as shown in the previous section.

```
*A:PE-1>config>service# info
        vprn 500 customer 1 create
            ecmp 4
            route-distinguisher 65071:500
            vrf-target target:65000:500
            interface "evi-501" create
                vpls "evpn-vxlan-501"
                    evpn-tunnel# no need to configure an IP address
                exit
            exit
            interface "subnet-2001" create
                address 10.10.10.1/24
                vpls "r-vpls 2001"
                exit
            exit
            interface "subnet-2002" create
                address 20.20.20.1/24
                vpls "r-vpls 2002"
                exit
```

```
            exit
            no shutdown
        exit
```

The VPRN 500 configuration in PE-2 and PE-3 would be changed in the same way by adding the evpn-tunnel and removing the IP address of the EVPN-tunnel R-VPLS interface. No other changes are required.

```
*A:PE-2>config>service# info
        vprn 500 customer 1 create
            ecmp 4
            route-distinguisher 65072:500
            auto-bind mpls-gre
            vrf-target target:65000:500
            interface "evi-501" create
                vpls "evpn-vxlan-501"
                    evpn-tunnel# no need to configure an IP address
                exit
            exit
            no shutdown
        exit
```

5.5.1.4 EVPN for VXLAN in R-VPLS Services with IPv6 interfaces and prefixes Example

In the following configuration example, PE1 is connected to CE1 in VPRN 30 through a dual-stack IP interface. VPRN 30 is connected to an EVPN-tunnel R-VPLS interface enabled for IPv6.

In the following excerpt configuration the PE1 will advertise, in BGP EVPN, the 172.16.0.0/24 and 200::/64 prefixes in two separate NLRIs. The NLRI for the IPv4 prefix will use GW IP = 0 and a non-zero GW MAC, whereas the NLRI for the IPv6 prefix will be sent with GW IP = Link-Local Address for interface "int-evi-301" and no GW MAC.

```
*A:PE1>config>service# info
        vprn 30 customer 1 create
            route-distinguisher 192.0.2.1:30
            vrf-target target:64500:30
            interface "int-PE-1-CE-1" create
                enable-ingress-stats
                address 172.16.0.254/24
                ipv6
                    address 200::1/64
                exit
                sap 1/1/1:30 create
                exit
            exit
            interface "int-evi-301" create
                ipv6
                exit
```

```
                vpls "evi-301"
                    evpn-tunnel
                exit
            exit
            no shutdown
-----------------------------------------------
```

5.5.2 EVPN-MPLS Configuration Examples

5.5.2.1 EVPN All-active Multi-homing Example

This section shows a configuration example for three 7750 SR, 7450 ESS, or 7950 XRS PEs, given the following assumptions:

- PE-1 and PE-2 are multi-homed to CE-12 that uses a LAG to get connected to the network. CE-12 is connected to LAG SAPs configured in an all-active multi-homing ethernet-segment.
- PE-3 is a remote PE that performs aliasing for traffic destined for the CE-12

The following configuration excerpt applies to a VPLS-1 on PE-1 and PE-2, as well as the corresponding ethernet-segment and LAG commands.

```
A:PE1# configure lag 1
A:PE1>config>lag# info
-----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/2
        lacp active administrative-key 1 system-id 00:00:00:00:69:72
        no shutdown
-----------------------------------------------
A:PE1>config>lag# /configure service system bgp-evpn
A:PE1>config>service>system>bgp-evpn# info
-----------------------------------------------
                route-distinguisher 192.0.2.69:0
                ethernet-segment "ESI-71" create
                    esi 0x01000000007100000001
                    es-activation-timer 10
                    service-carving
                        mode auto
                    exit
                    multi-homing all-active
                    lag 1
                    no shutdown
                exit
-----------------------------------------------
A:PE1>config>service>system>bgp-evpn# /configure service vpls 1
A:PE1>config>service>vpls# info
-----------------------------------------------
```

```
            bgp
            exit
            bgp-evpn
                cfm-mac-advertisement
                evi 1
                vxlan
                    shutdown
                exit
                mpls
                    ingress-replication-bum-label
                    auto-bind-tunnel
                        resolution any
                    exit
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            sap lag-1:1 create

            exit
            no shutdown
-----------------------------------------------

A:PE2# configure lag 1
A:PE2>config>lag# info
-----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/3
        lacp active administrative-key 1 system-id 00:00:00:00:69:72
        no shutdown
-----------------------------------------------
A:PE2>config>lag# /configure service system bgp-evpn
A:PE2>config>service>system>bgp-evpn# info
-----------------------------------------------
                route-distinguisher 192.0.2.72:0
                ethernet-segment "ESI-71" create
                    esi 0x01000000007100000001
                    es-activation-timer 10
                    service-carving
                        mode auto
                    exit
                    multi-homing all-active
                    lag 1
                    no shutdown
                exit
-----------------------------------------------
A:PE2>config>service>system>bgp-evpn# /configure service vpls 1
A:PE2>config>service>vpls# info
-----------------------------------------------
            bgp
            exit
            bgp-evpn
                cfm-mac-advertisement
                evi 1
                vxlan
                    shutdown
```

```
                exit
                mpls
                    ingress-replication-bum-label
                    auto-bind-tunnel
                        resolution any
                    exit
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            sap lag-1:1 create
            exit
            no shutdown
-----------------------------------------------
```

The configuration on the remote PE (i.e. PE-3), which supports aliasing to PE-1 and PE-2 is shown below. PE-3 does not have any ethernet-segment configured. It only requires the VPLS-1 configuration and ecmp>1 in order to perform aliasing.

```
*A:PE3>config>service>vpls# info
-----------------------------------------------
            bgp
            exit
            bgp-evpn
                cfm-mac-advertisement
                evi 1
                vxlan
                    shutdown
                exit
                mpls
                    ingress-replication-bum-label
                    ecmp 4
                    auto-bind-tunnel
                        resolution any
                    exit
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            sap 1/1/1:1 create
            exit
            spoke-sdp 4:13 create
                no shutdown
            exit
            no shutdown
-----------------------------------------------
```

5.5.2.2 EVPN Single-active Multi-homing Example

If we wanted to use **single-active** multi-homing on PE-1 and PE-2 instead of **all-active** multi-homing, we would only need to modify the following:

- change the LAG configuration to **single-active**

  The CE-12 will be now configured with two different LAGs, hence the key/system-id/system-priority must be different on PE-1 and PE-2
- change the ethernet-segment configuration to **single-active**

No changes are needed at service level on any of the three PEs.

The differences between single-active versus all-active multi-homing are highlighted in **bold** in the following example excerpts:

```
A:PE1# configure lag 1
A:PE1>config>lag# info
----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/2
        lacp active administrative-key 1 system-id 00:00:00:00:69:69
        no shutdown
----------------------------------------------
A:PE1>config>lag# /configure service system bgp-evpn
A:PE1>config>service>system>bgp-evpn# info
----------------------------------------------
                route-distinguisher 192.0.2.69:0
                ethernet-segment "ESI-71" create
                    esi 0x01000000007100000001
                    es-activation-timer 10
                    service-carving
                        mode auto
                    exit
                    multi-homing single-active
                    lag 1
                    no shutdown
                exit
----------------------------------------------

A:PE2# configure lag 1
A:PE2>config>lag# info
----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/3
        lacp active administrative-key 1 system-id 00:00:00:00:72:72
        no shutdown
----------------------------------------------
A:PE2>config>lag# /configure service system bgp-evpn
A:PE2>config>service>system>bgp-evpn# info
----------------------------------------------
                route-distinguisher 192.0.2.72:0
                ethernet-segment "ESI-71" create
```

```
                    esi 0x01000000007100000001
                    es-activation-timer 10
                    service-carving
                        mode auto
                    exit
                    multi-homing single-active
                    lag 1
                    no shutdown
                exit
-----------------------------------------------
```

5.5.3 PBB-EVPN Configuration Examples

5.5.3.1 PBB-EVPN All-active Multi-homing Example

As in the EVPN All-active Multi-homing Example, this section also shows a configuration example for three 7750 SR, 7450 ESS, or 7950 XRS PEs, however, PBB-EVPN is used in this excerpt, as follows:

- PE-1 and PE-2 are multi-homed to CE-12 that uses a LAG to get connected to I-VPLS 20001. CE-12 is connected to LAG SAPs configured in an **all-active** multi-homing ethernet-segment.
- PE-3 is a remote PE that performs aliasing for traffic destined for the CE-12.
- The three PEs are connected through B-VPLS 20000, a Backbone VPLS where EVPN is enabled.

The following excerpt shows the example configuration for I-VPLS 20001 and B-VPLS 20000 on PE-1 and PE-2, as well as the corresponding ethernet-segment and LAG commands:

```
*A:PE1# configure lag 1
*A:PE1>config>lag# info
----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/2
        lacp active administrative-key 1 system-id 00:00:00:00:69:72
        no shutdown
----------------------------------------------
*A:PE1>config>lag# /configure service system bgp-evpn
*A:PE1>config>service>system>bgp-evpn# info
----------------------------------------------
                route-distinguisher 192.0.2.69:0
                ethernet-segment "ESI-71" create
                    esi 01:00:00:00:00:71:00:00:00:01
                    source-bmac-lsb 71-71 es-bmac-table-size 8
                    es-activation-timer 5
                    service-carving
```

```
                        mode auto
                    exit
                    multi-homing all-active
                    lag 1
                    no shutdown
                exit
-----------------------------------------------
*A:PE1>config>service>system>bgp-evpn# /configure service vpls 20001
*A:PE1>config>service>vpls# info
-----------------------------------------------
            pbb
                backbone-vpls 20000
                exit
            exit
            stp
                shutdown
            exit
            sap lag-1:71 create
            exit
            no shutdown
-----------------------------------------------
*A:PE1>config>service>vpls# /configure service vpls 20000
*A:PE1>config>service>vpls# info
-----------------------------------------------
            service-mtu 2000
            pbb
                source-bmac 00:00:00:00:00:69
                use-es-bmac
            exit
            bgp-evpn
                evi 20000
                vxlan
                    shutdown
                exit
                mpls
                    auto-bind-tunnel
                        resolution any
                    exit
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            no shutdown
-----------------------------------------------

*A:PE2# configure lag 1
*A:PE2>config>lag# info
-----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/3
        lacp active administrative-key 1 system-id 00:00:00:00:69:72
        no shutdown
-----------------------------------------------
*A:PE2>config>lag# /configure service system bgp-evpn
*A:PE2>config>service>system>bgp-evpn# info
-----------------------------------------------
```

```
                route-distinguisher 192.0.2.72:0
                ethernet-segment "ESI-71" create
                    esi 01:00:00:00:00:71:00:00:00:01
                    source-bmac-lsb 71-71 es-bmac-table-size 8
                    es-activation-timer 5
                    service-carving
                        mode auto
                    exit
                    multi-homing all-active
                    lag 1
                    no shutdown
                exit
-------------------------------------------
*A:PE2>config>service>system>bgp-evpn# /configure service vpls 20001
*A:PE2>config>service>vpls# info
-------------------------------------------
            pbb
                backbone-vpls 20000
                exit
            exit
            stp
                shutdown
            exit
            sap lag-1:71 create
            exit
            no shutdown
-------------------------------------------
*A:PE2>config>service>vpls# /configure service vpls 20000
*A:PE2>config>service>vpls# info
-------------------------------------------
            service-mtu 2000
            pbb
                source-bmac 00:00:00:00:00:72
                use-es-bmac
            exit
            bgp-evpn
                evi 20000
                vxlan
                    shutdown
                exit
                mpls
                    auto-bind-tunnel
                        resolution any
                    exit
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            no shutdown
-------------------------------------------
*A:PE2>config>service>vpls#
```

The combination of the pbb s**ource-bmac** and the ethernet-segment **source-bmac-lsb** create the same BMAC for all the packets sourced from both PE-1 and PE-2 for ethernet-segment "ESI-71".

5.5.3.2 PBB-EVPN Single-Active Multi-Homing Example

In the following configuration example, PE-70 and PE-73 are part of the same single-active multi-homing, ethernet-segment ESI-7413. In this case, the CE is connected to PE-70 and PE-73 through spoke-sdps 4:74 and 34:74, respectively.

In this example PE-70 and PE-73 use a different source-bmac for packets coming from ESI-7413 and it is not an **es-bmac** as shown in the PBB-EVPN All-active Multi-homing Example.

```
*A:PE70# configure service system bgp-evpn
*A:PE70>config>service>system>bgp-evpn# info
----------------------------------------------
                route-distinguisher 192.0.2.70:0
                ethernet-segment "ESI-7413" create
                    esi 01:74:13:00:74:13:00:00:74:13
                    es-activation-timer 0
                    service-carving
                        mode auto
                    exit
                    multi-homing single-active
                    sdp 4
                    no shutdown
                exit
----------------------------------------------
*A:PE70>config>service>system>bgp-evpn# /configure service vpls 20001
*A:PE70>config>service>vpls# info
----------------------------------------------
            pbb
                backbone-vpls 20000
                exit
            exit
            stp
                shutdown
            exit
            spoke-sdp 4:74 create
                no shutdown
            exit
            no shutdown
----------------------------------------------
*A:PE70>config>service>vpls# /configure service vpls 20000
*A:PE70>config>service>vpls# info
----------------------------------------------
            service-mtu 2000
            pbb
                source-bmac 00:00:00:00:00:70
            exit
            bgp-evpn
                evi 20000
                vxlan
                    shutdown
                exit
                mpls
                    ecmp 2
                    auto-bind-tunnel
                        resolution any
```

```
                    exit
                    no shutdown
                exit
            exit
            stp
                shutdown
            exit
            no shutdown
-----------------------------------------------
*A:PE70>config>service>vpls#


A:PE73>config>service>system>bgp-evpn# info
-----------------------------------------------
                route-distinguisher 192.0.2.73:0
                ethernet-segment "ESI-7413" create
                    esi 01:74:13:00:74:13:00:00:74:13
                    es-activation-timer 0
                    service-carving
                        mode auto
                    exit
                    multi-homing single-active
                    sdp 34
                    no shutdown
                exit
-----------------------------------------------
A:PE73>config>service>system>bgp-evpn# /configure service vpls 20001
A:PE73>config>service>vpls# info
-----------------------------------------------
            pbb
                backbone-vpls 20000
                exit
            exit
            stp
                shutdown
            exit
            spoke-sdp 34:74 create
                no shutdown
            exit
            no shutdown
-----------------------------------------------
A:PE73>config>service>vpls# /configure service vpls 20000
A:PE73>config>service>vpls# info
-----------------------------------------------
            service-mtu 2000
            pbb
                source-bmac 00:00:00:00:00:73
            exit
            bgp-evpn
                evi 20000
                vxlan
                    shutdown
                exit
                mpls
                    auto-bind-tunnel
                        resolution any
                    exit
                    no shutdown
                exit
```

```
            exit
            stp
                shutdown
            exit
            no shutdown
-----------------------------------------------
A:PE73>config>service>vpls#
```

5.6 EVPN Command Reference

This chapter describes the EVPN command reference.

5.6.1 Command Hierarchies

5.6.1.1 EVPN Configuration Commands

```
config
    — service
        — epipe service-id [customer customer-id] [test] [create] [vpn vpn-id] [vc-switching]
        — no epipe service-id
            — bgp
                — route-distinguisher [ip-addr:comm-val | as-number:ext-comm-val |
                    auto-rd]
                — no route-distinguisher
                — route-target {ext-community | export ext-community] | import ext-
                    community}
                — no route-target
            — [no] bgp-evpn
                — [no] evi value
                — [no] local-ac-name ac-name
                    — [no] eth-tag value
                — [no] remote-ac-name ac-name
                    — [no] eth-tag value
                — mpls
                    — auto-bind-tunnel
                        — resolution {disabled | any | filter}
                        — resolution-filter
                            — [no] bgp
                            — [no] ldp
                            — [no] rsvp
                            — [no] sr-isis
                            — [no] sr-ospf
                            — [no] sr-te
                            — [no] udp
                    — [no] control-word
                    — ecmp max-ecmp-routes
                    — [no] entropy-label
                    — [no] force-vlan-vc-forwarding
                    — send-evpn-encap [mpls] [mplsoudp]
                    — no send-evpn-encap
                    — [no] shutdown
            — vxlan vni vni-id [create] [instance instance-id]
            — no vxlan vni vni-id
                — egr-vtep {ip-address | ipv6-address}
```

```
            — no egr-vtep
                — oper-group name
                — no oper-group
            — vxlan-src-vtep {ip-address | ipv6-address}
            — no vxlan-src-vtep
        — proxy-arp-nd
            — mac-list name [create]
            — no mac-list name
                — [no] mac ieee-address
        — vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [{b-vpls | i-vpls}] [etree]
          [create]
        — no vpls service-id
            — bgp [bgp-instance]
            — no bgp [bgp-instance]
                — route-distinguisher [{ip-addr:comm-val | as-number:ext-comm-val |
                  auto-rd}]
                — no route-distinguisher
                — route-target ext-community
                — route-target export ext-community [import ext-community]
                — route-target import ext-community
                — no route-target
                — vsi-export policy-name [policy-name...(up to 5 max)]
                — no vsi-export
                — vsi-import policy-name [policy-name...(up to 5 max)]
                — no vsi-import
            — [no] bgp-evpn
                — accept-ivpls-evpn-flush
                — no accept-ivpls-evpn-flush
                — [no] cfm-mac-advertisement
                — [no] evi value
                — incl-mcast-orig-ip ip-address
                — no incl-mcast-orig-ip
                — [no] ingress-repl-inc-mcast-advertisement
                — [no] ip-route-advertisement [incl-host]
                — [no] ip-route-advertisement [incl-host]
                — isid-route-target
                    — isid-range from [to to] {auto-rt | route-target rt}
                    — [no] isid-range from
                — [no] mac-advertisement
                — mac-duplication
                    — detect num-moves num-moves window minutes
                    — [no] retry minutes
                    — [no] black-hole-dup-mac
                — mpls
                    — auto-bind-tunnel
                        — resolution {disabled | any | filter}
                        — resolution-filter
                            — [no] bgp
                            — [no] ldp
                            — [no] rsvp
                            — [no] sr-isis
                            — [no] sr-ospf
                            — [no] sr-te
                            — [no] udp
```

```
                    — bgp-instance [bgp-instance]
                    — [no] control-word
                    — ecmp max-ecmp-routes
                    — [no] entropy-label
                    — [no] force-vlan-vc-forwarding
                    — [no] ingress-replication-bum-label
                    — [no] send-evpn-encap [mpls] [mplsoudp]
                    — [no] shutdown
                    — [no] split-horizon-group
                — [no] unknown-mac-route
                — vxlan
                    — [no] send-evpn-encap
                    — [no] send-imet-ir-on-ndf
                    — [no] shutdown
            — pbb
                — leaf-source-bmac ieee-address
                — no leaf-source-bmac
                — send-bvpls-evpn-flush
                — no send-bvpls-evpn-flush
                — [no] use-es-bmac
            — provider-tunnel
                — inclusive
                    — data-delay-interval seconds
                    — no data-delay-interval
                    — [no] mldp
                    — [no] root-and-leaf
                    — [no] owner {bgp-ad | bgp-vpls | bgp-evpn-mpls}
                    — [no] shutdown
            — [no] proxy-arp
                — [no] age-time seconds
                — dup-detect [anti-spoof-mac mac-address] [static-black-hole]
                    window minutes num-moves count hold-down minutes | max
                — dynamic ip-address [create]
                — no dynamic ip-address
                    — mac-list name
                    — [no] mac-list name
                    — resolve mins
                — [no] dynamic-arp-populate
                — evpn-route-tag tag
                — no evpn-route-tag
                — [no] garp-flood-evpn
                — [no] send-refresh seconds
                — [no] static ip-address ieee-address
                — table-size table-size
                — [no] unknown-arp-request-flood-evpn
                — [no] shutdown
            — [no] proxy-nd
                — [no] age-time seconds
                — dup-detect [anti-spoof-mac mac-address] [static-black-hole]
                    window minutes num-moves count hold-down {minutes | max}
                — dynamic ipv6-address [create]
                — no dynamic ipv6-address
                    — mac-list name
                    — [no] mac-list
```

```
                                    — resolve mins
                                — [no] dynamic-nd-populate
                                — evpn-nd-advertise {host | router}
                                — evpn-route-tag tag
                                — no evpn-route-tag
                                — [no] host-unsolicited-na-flood-evpn
                                — [no] router-unsolicited-na-flood-evpn
                                — [no] send-refresh seconds
                                — [no] static ip-address ieee-address {host | router}
                                — table-size table-size
                                — [no] unknown-ns-flood-evpn
                                — [no] shutdown
                            — sap
                                — [no] disable-send-bvpls-evpn-flush
                            — spoke-sdp
                                — [no] disable-send-bvpls-evpn-flush
                            — static-mac
                                — mac ieee-address [create] sap sap-id monitor fwd-status
                                — mac ieee-address [create] spoke-sdp sdp-id:vc-id monitor fwd-status
                                — mac ieee-address [create] black-hole
                                — no mac ieee-address
                            — vsd-domain name
                            — no vsd-domain vni
                            — vxlan
                            — vxlan vni vni-id [create] [instance instance-id]
                            — no vxlan vni vni-id
                                — assisted-replication {replicator | leaf] [replicator-activation-time
                                    seconds}
                                — no assisted-replication
                            — vxlan-src-vtep {ip-address | ipv6-address}
                            — no vxlan-src-vtep
                        — vprn
                            — interface
                                — vpls
                                    — [no] evpn-tunnel
                            — vxlan
                                — tunnel-termination {ip-address | ipv6-address} fpe fpe-id [create]
                                — no tunnel-termination {ip-address | ipv6-address}
                        — vsd
                            — domain name [type {l2-domain | vrf-gre | vrf-vxlan | l2-domain-irb}] [create]
                            — no domain name
                                — description description-string
                                — no description
                                — [no] shutdown
                            — service-range svc-id to svc-id

config
    — service
        — system
            — [no] bgp-auto-rd-range ip-address comm-val [1 to 65535] to [1 to 65535]
            — [no] bgp-evpn
                — ad-per-es-route-target [evi-rt | evi-rt-set route-distinguisher ip-
                    address]
                — ethernet-segment name [virtual] [create]
```

```
                        — no ethernet-segment name
                            — dot1q
                                — q-tag-range qtag1 [to qtag1]
                                — no q-tag-range qtag1
                            — es-activation-timer seconds
                            — no es-activation-timer
                            — esi esi
                            — no esi
                            — lag lag-id
                            — no lag
                            — multi-homing single-active [no-esi-label]
                            — multi-homing all-active
                            — no multi-homing
                            — network-interconnect-vxlan instance
                            — no network-interconnect-vxlan
                            — port port-id
                            — no port
                            — qinq
                                — s-tag qtag1 c-tag-range qtag2 [to qtag1]
                                — no s-tag qtag1 c-tag-range qtag2
                                — s-tag-range qtag1 [to qtag1]
                                — no s-tag-range qtag1
                            — sdp sdp-id
                            — no sdp
                            — service-carving
                                — manual
                                    — evi start [to to]
                                    — no evi start
                                    — isid start [to to]
                                    — no isid start
                                    — preference [create] [non-revertive]
                                    — no preference
                                        — value value
                                — mode {manual | auto | off}
                            — service-id
                                — service-range svc-id [to svc-id]
                                — no service-range svc-id
                            — [no] shutdown
                            — [no] source-bmac-lsb MAC-lsb [ex-bmac-table-size size]
                            — vc-id-range from [to to]
                            — no vc-id-range vc-id
                        — [no] evpn-etree-leaf-label
                        — route-distinguisher rd
                    — vxlan
                        — assisted-replication-ip ip-address
                        — no assisted-replication-ip
                        — tunnel-termination ip-address fpe fpe-id [create]
                        — no tunnel-termination ip-address

config
    — redundancy
        — bgp-evpn-multi-homing
                        — boot-timer seconds
                        — es-activation-timer seconds
```

```
config
    — system
        — vsd
            — system-id name
            — no system-id
        — xmpp
            — server xmpp-server-name [domain-name fqdn] [username user-name]
                [password password] [create] [service-name service-name]
            — server xmpp-server-name [domain-name fqdn] [username user-name]
                [password password] [create] [router router-instance]
            — no server xmpp-server-name
                — [no] shutdown
        — security
            — cli-script
                — authorization
                    — vsd
                        — [no] cli-user user-name
            — password
                — vsd-password

config
    — router
        — bgp
            — group
                — neighbor ip-address
                    — def-recv-evpn-encap [mpls | vxlan]

config
    — python
        — python-policy name
            — vsd script script

<root>
    — enable-vsd-config name
```

5.6.1.2 Show Commands

```
show
    — service
        — evpn-mpls [tep-ip-address]
        — id service-id
            — bgp bgp-instance
            — bgp-evpn
                — isid-route-target
            — proxy-arp ip-address [detail]
            — proxy-arp ip-address dynamic
            — es-pbr
            — evpn-mpls
                — esi esi
                — es-bmac ieee-address
            — vxlan
```

```
                    — vxlan assisted-replication replicator
                    — vxlan instance instance oper-flags
                — provider-tunnel-using leaf-only [bgp-ad | bgp-vpls | bgp-evpn-mpls]
                — provider-tunnel-using root-and-leaf [bgp-ad | bgp-vpls | bgp-evpn-mpls]
                — proxy-arp-nd
                    — mac-list name
                    — mac-list name associations
                    — mac-list name
                — service-using [vsd] [origin creation-origin]
                — system
                    — bgp-evpn
                        — ethernet-segment name name [all]
                        — ethernet-segment name name evi [evi]
                        — ethernet-segment name name isid [isid]
                        — ethernet-segment name name virtual-ranges
                — vsd [vsd] [origin creation-origin]
                    — domain [domain-name] [association]
                    — root-objects
                    — script
                        — snippets snippet-name [instance snippet-instance] [detail]
                        — statistics
                    — summary
                — vxlan [ip-address]
                — vxlan-instance-using ethernet-segment [name]
            — system
                — xmpp
                    — server [name]
                    — vsd [entry]
                — vsd
                    — domain [domain-name] [association]
                — vxlan

        show
            — redundancy
                — bgp-evpn-multi-homing
```

5.6.1.3 Clear Commands

```
        clear
            — service
                — id
                    — evpn
                        — mac-dup-detect {ieee-address | all}
                — statistics
                    — vsd
                        — domain name
                        — scripts name

        clear
            — system
```

```
— statistics
    — xmpp
        — server xmpp-server-name
```

5.6.1.4 Debug Commands

```
debug
— system
    — xmpp [connection] [gateway] [message] [vsd] [iq] [all]
— [no] xmpp
— vsd
    — scripts
        — [no] event
        — [no] cli
        — [no] errors
— [no] executed-cmd
        — [no] state-change
        — [no] warnings
    — instance instance
        — [no] event
        — [no] cli
        — [no] errors
— [no] executed-cmd
        — [no] state-change
        — [no] warnings
```

5.6.1.5 Tools Commands

```
tools
— dump
    — service
        — domain-to-vsd-mapping
            — domain name name
        — evpn
            — usage
        — id service-id
            — evpn-mpls [clear] [default-multicast-list]
            — vxlan [clear]
            — evpn
                — usage
        — system
            — bgp-evpn
                — ethernet-segment name evi evi df
                — ethernet-segment name isid isid df
        — vsd-services
            — command-list
        — vxlan [vtep]
            — dup-vtep-egrvni [clear]
```

```
tools
    — perform
        — service
            — id
                — proxy-arp
                    — dynamic-resolve all [force]
                    — dynamic-resolve ip-address [force]
                — proxy-nd
                    — dynamic-resolve all [force]
                    — dynamic-resolve ipv6-address [force]
            — vsd
                — domain name [name] refresh-config
                — fd-domain-sync {full | diff}
                — evaluate-script domain-name [domain-name] type [type] action
                    script-action [vni vni-id] [rt-i ext-community] [rt-e ext-community]
                    [metadata metadata] policy python-policy

tools
    — perform
        — system
            — xmpp
                — vsd-refresh
```

5.6.2 Command Descriptions

5.6.2.1 EVPN Configuration Commands

epipe

Syntax **epipe** *service-id* **customer** *customer-id* [**vpn** *vpn-id*] [**vc-switching**] [**create**]
epipe *service-id* [**test**] [**create**]
no epipe *service-id*

Context config>service

Description This command configures an Epipe service instance. This command is used to configure a point-to-point epipe service. An Epipe connects two endpoints defined as Service Access Points (SAPs). Both SAPs may be defined in one or they may be defined in separate devices connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far end SAP is generalized into a Service Distribution Point (SDP). This SDP describes a destination and the encapsulation method used to reach it. In addition to the SDPs, endpoint SAPs can also be connected by EVPN destinations.

No MAC learning or filtering is provided on an Epipe.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. After a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

After a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

By default, no Epipe services exist until they are explicitly created with this command.

The **no** form of this command deletes the Epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every node on which this service is defined.

Values *service-id*: 1 to 2147483648
svc-name: 64 characters maximum

customer *customer-id* — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

Default null (0)

vc-switching — Specifies if the pseudowire switching signaling is used for the spoke-SDPs configured in this service. When an Epipe is configured with **vc-switching**, no bgp-evpn is allowed.

test — Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs. This parameter applies to the 7450 ESS and 7750 SR only.

create — Keyword used to create the service instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

proxy-arp-nd

Syntax **proxy-arp-nd**

Context config>service

Description This command enables the context to configure the service-level **proxy-arp-nd** commands.

mac-list

Syntax **mac-list** *name* [**create**]
no mac-list *name*

Context config>service>proxy-arp-nd

Description This command creates a list of MAC addresses that can be pointed at from the service for a specified IP. The list may contain up to 10 MAC addresses; an empty list is also allowed.

The MAC list allows on-the-fly changes, but a change in the list deletes the proxy entries for all the IPs using that list.

The **no** form of the command deletes the entire MAC-list. Deleting a MAC list is only possible if it is not referenced in the configuration.

Parameters *name* — Specifies the name of the MAC address list, which can be up to 32 characters.

create — Mandatory keyword to create a MAC list.

mac

Syntax **mac** *ieee-address*
no mac *ieee-address*

Context config>service>proxy-arp-nd>mac-list

Description This command configures the proxy ARP or ND MAC address information.

The **no** form of the command deletes the MAC address.

Parameters *ieee-address* — Specifies the MAC address added to the list. The MAC list can be empty or contain up to 10 addresses.

Values xx:xx:xx:xx:xx:xx
xx-xx-xx-xx-xx-xx

vpls

Syntax **vpls** *service-id* **customer** *customer-id* **vpn** *vpn-id* [**m-vpls**] [{**bvpls** | **i-vpls**}] [**etree**] [**create**]
no vpls *service-id*

Context config>service

Description This command creates or edits a Virtual Private LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. After a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

To create a management VPLS on a 7450 ESS, the m-vpls keyword must be specified. See section **Hierarchical VPLS Redundancy** for an introduction to the concept of management VPLS.

After a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 2147483648
svc-name: 64 characters maximum

customer *customer-id* — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default null (0)

m-vpls — Specifies a management VPLS.

b-vpls | i-vpls — Creates a backbone-vpls or ISID-vpls.

etree — Creates an Ethernet-Tree (E-Tree) service.

bgp

Syntax **bgp** *bgp-instance*

Context config>service>epipe

Description This command enables the context to configure the BGP related parameters for BGP EVPN.

bgp

Syntax **bgp** *bgp-instance*
no bgp *bgp-instance*

Context config>service>vpls

Description This command enables the context to configure the BGP related parameters for BGP VPLS.

A maximum of two BGP instances can be configured in a VPLS service. The *bgp-instance* parameter value can be configured as 1 or 2. If it is not specified, the parameter value is configured as 1 by default.

The **route-distinguisher** configured in BGP instance 1 and 2 must be different. However, the route-target value may be configured the same or different for the two instances.

Only bgp-evpn mpls is allowed to be assigned to instance 2. Instance 1 must be used for the VXLAN and L2VPN address families.

BGP-EVPN VXLAN and BGP-EVPN MPLS can only be configured as **no shutdown** in the same service if they are associated with different instances (When the two BGP instances are created, the **bgp-instance** command must be configured in the **bgp-evpn mpls** context).

The **evi** value in **bgp-evpn** can be used to auto-derive the route distinguisher in instance 1 only. However, the **evi** value can be used to auto-derive the **route-target** in both instances.

The **no** version of the command removes the BGP instance.

Parameters *bgp-instance* — Specifies the value associated with the BGP instance.

Values 1 to 2

route-target

Syntax **route-target** *ext-community*

route-target export *ext-community* [**import** *ext-community*]
route-target import *ext-community*
no route-target

Context config>service>vpls>bgp-ad
config>service>vpls>bgp
config>service>epipe>bgp

Description This command configures the route target (RT) component that will be signaled in the related MP- BGP attribute to be used for BGP auto-discovery, BGP VPLS, BGP multi-homing and EVPN if these features are configured in this VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The ext-comm can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community. For BGP EVPN enabled VPLS and Epipe services, the route target can also be auto-derived from the **evi** value (**config>service>vpls>bgp-evpn>evi** or **config>service>epipe>bgp-evpn>evi**) if this command is not configured. See the **evi** command description for more information.

Parameters **export** *ext-community* — Specifies communities allowed to be sent to remote PE neighbors.

import *ext-community* — Specifies communities allowed to be accepted from remote PE neighbors.

vsi-export

Syntax **vsi-export** *policy-name* [*policy-name* ... (up to 5 max)]
no vsi-export

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command specifies the name of the VSI export policies to be used for BGP auto-discovery, BGP VPLS and BGP multi-homing if these features are configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

Parameters *policy-name* — Specifies a VSI export policy. 32 characters max.

vsi-import

Syntax **vsi-import** *policy-name* [*policy-name* ... (up to 5 max)]
no vsi-import

Context config>service>vpls>bgp-ad>vsi-id

config>service>vpls>bgp

Description This command specifies the name of the VSI import policies to be used for BGP auto-discovery, BGP VPLS and BGP multi-homing if these features are configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

Parameters *policy-name* — Specifies a VSI import policy. 32 characters max.

route-distinguisher

Syntax **route-distinguisher** [{*ip-addr:comm-val* | *as-number:ext-comm-val*}]
route-distinguisher auto-rd
no route-distinguisher

Context config>service>vpls>bgp
config>service>epipe>bgp

Description This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for L2VPN and EVPN families. This value will be used for BGP-AD, BGP VPLS and BGP multi-homing NLRI if these features are configured.

If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:

- if BGP AD VPLS-id is configured and no RD is configured under BGP node - RD=VPLS-ID
- if BGP AD VPLS-id is not configured then an RD value must be configured under BGP node (this is the case when only BGP VPLS is configured)
- if BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails

Values and format (6 bytes, other 2 bytes of type will be automatically generated)

Alternatively, the **auto-rd** option allows the system to automatically generate an RD based on the **bgp-auto-rd-range** command configured at the service level. For **bgp-evpn** enabled VPLS and Epipe services, the **route-distinguisher** value can also be auto-derived from the **evi** value (**config>service>vpls>bgp-evpn>evi** or **config>service>epipe>bgp-evpn>evi**) if this command is not configured. See the **evi** command description for more information.

Parameters *ip-addr:comm-val* — Specifies the IP address.

Values *ip-addr*: a.b.c.d
comm-val: 0 to 65535

as-number:ext-comm-val — Specifies the AS number.

Values *as-number*: 1 to 65535

ext-comm-val: 0 to 4294967295

auto-rd — The system will generate an RD for the service according to the IP address and range configured in the **bgp-auto-rd-range** command.

bgp-auto-rd-range

Syntax **bgp-auto-rd-range** *ip-address* **comm-val** *comm-val* **to** *comm-val*
no bgp-auto-rd-range

Context config>service>system

Description This command defines the type-1 route-distinguisher IPv4 address and community value range within which the system will select a route-distinguisher for the **bgp-enabled** services using **auto-rd**.

Interactions:

This command is used along with the **route-distinguisher auto-rd** command supported in VPLS, VPRN and Epipe services. The system forces the user to create a **bgp-auto-range** before the **auto-rd** option can be used in the services.

The system will keep allocating values for services configured with **route-distinguisher auto-rd** as long as there are available community values within the configured range. After the command is added, the following changes are allowed:

- The ip-address can be changed without modifying the *comm-val* range, even if services using **auto-rd** are present. The affected routes will be withdrawn and re-advertised with the new route-distinguishers.
- The *comm-val* range can be modified as long as no conflicting values are present in the new range. For example, the user may expand the range as long as the new range does not overlap with existing manual route-distinguishers. The user may also reduce the range as long as the new range can accommodate the already allocated auto-RDs.

Parameters *ip-address* — Specifies the IPv4 address used in the first 4 octets of all the type-1 auto route-distinguishers selected by the system.

comm-val — Specifies the community value of the type-1 auto route-distinguisher.

Values 0 to 65535

bgp-evpn

Syntax [**no**] **bgp-evpn**

Context config>service>vpls
config>service>system
config>service>epipe

Description This command enables the context to configure the BGP EVPN parameters in the base instance.

ad-per-es-route-target

Syntax **ad-per-es-route-target** {**evi-rt** | **evi-rt-set route-distinguisher** *ip-address*}
no ad-per-es-route-target

Context config>service>system>bgp-evpn

Description This command controls how Ethernet AD per-ES routes are generated.

The system can either send a separate Ethernet AD per-ES route per service, or an Ethernet AD per-ES routes aggregating the route-targets for multiple services. While both alternatives will inter-operate, RFC 7432 states that the EVPN Auto-Discovery per-ES route must be sent with a set of route-targets corresponding to all the EVIs defined on the Ethernet segment. The command supports both options.

The default option **ad-per-es-route-target evi-rt** configures the system to send a separate AD per-ES route per service.

When enabled, the **evi-rt-set** option allows the aggregation of routes: A single AD per-ES route with the associated RD (ip-address:1) and a set of EVI route-targets will be advertised (to a maximum of 128). When a significant number of EVIs are defined in the Ethernet segment (hence the number of route-targets), the system will send more than one route. For example:

- AD per-ES route for evi-rt-set 1 will be sent with RD ip-address:1
- AD per-ES route for evi-rt-set 2 will be sent with RD ip-address:2

Default ad-per-es-route-target evi-rt

Parameters **evi-rt** — Specifies the option to advertise a separate AD per-ES route per service.

evi-rt-set — Specifies the option to advertise a set of AD per-ES routes aggregating the route-targets for all the services in the Ethernet segment.

ip-address — Specifies the ip-address part of the route-distinguisher being used in the evi-rt-set option.

route-distinguisher

Syntax **route-distinguisher** [*ip-addr:comm-val* | *as-number:ext-comm-va*l]
no route-distinguisher

Context config>service>system>bgp-evpn

Description This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for EVPN corresponding to the base EVPN instance (Ethernet Segment routes). If the route-distinguisher component is not configured, the system will use system:ip-address as the default route-distinguisher

Default no route-distinguisher

Parameters *ip-addr:comm-val* — Specifies the IP address.

Values *ip-addr*: a.b.c.d
comm-val: 0 to 65535

as-number:ext-comm-val — Specifies the AS number.

Values *as-number*: 1 to 65535
ext-comm-val: 0 to 4294967295

evpn-etree-leaf-label

Syntax **evpn-etree-leaf-label**
no evpn-etree-leaf-label

Context config>service>system>bgp-evpn

Description This command enables EVPN Ethernet-Tree (E-Tree) VPLS services on the router (not B-VPLS). It allocates an E-Tree leaf label for the PE and programs the ILM entry.

The command ensures that in-flight traffic can perform an ILM entry lookup at any time, and therefore avoid the discards during shutdown or no shutdown services (or at least reduce the timing window so that it does not occur during normal operation or configuration).

Note: The **evpn-etree-leaf-label** command must be configured to execute **bgp-evpn mpls no shutdown**.

ethernet-segment

Syntax **ethernet-segment** *name* [**virtual**] [**create**]
no ethernet-segment

Context config>service>system>bgp-evpn

Description This command configures an Ethernet segment instance and its corresponding name. The configuration of the dot1q or qinq nodes is only allowed if the Ethernet segment (ES) is created as **virtual**.

For a virtual ES, a port, LAG, or SDP must be created for the ES before configuring a VLAN or vc-id association.

When a port or LAG is added, the **type** and **encap-type** values are checked. If the **encap-type** is **dot1q**, then only the dot1q node can be configured; the qinq context is not allowed. In the same way, if the **encap-type** is **qinq**, then only the qinq node is allowed. A dot1q, qinq, or vc-id range is required for a virtual ES to be operationally active.

Parameters *name* — Specifies the 28-character ES name.

virtual — This keyword specifies that the ES is virtual and is associated to logical interfaces, in addition to ports, LAGs, or SDPs.

create — Mandatory keyword for creating an ES.

dot1q

Syntax **dot1q**

Context config>service>system>bgp-evpn>ethernet-segment

Description This command creates the dot1q context for q-tag additions to the port or LAG virtual ES.

q-tag-range

Syntax **q-tag-range** *qtag1* [**to** *qtag1*]
no q-tag-range *qtag1*

Context config>service>system>bgp-evpn>ethernet-segment>dot1q

Description This command determines the VIDs associated with the virtual Ethernet segment on a specific dot1q port or LAG based on the following considerations:

- Values *, 0 to 4094 are allowed.
- Any SAP for which the service-delimiting qtag matches the range is associated with the virtual ES, and only those, for example, sap 1/1/1:0 will not match port 1/1/1, qtag-range 100.
- Maximum 8 ranges are allowed in the dot1q context.
- A range can be comprised of a single qtag.
- Shutting down the ES is not required prior to changing the q-tag-range.

The **no** form of the command removes the configured range. Only the first qtag1 value is required to remove the range.

Parameters *qtag1* — Specifies the VID. When configuring a range of qtags (and not a single value), the second qtag1 value must be greater than the first qtag1.

Values *, 0 to 4094

qinq

Syntax **qinq**

Context config>service>system>bgp-evpn>ethernet-segment

Description This command creates the qinq context for q-tag and s-tag additions to the port or LAG virtual Ethernet segments.

s-tag-range

Syntax **s-tag-range** *qtag1* [**to** *qtag1*]
no s-tag-range *qtag1*

Context config>service>system>bgp-evpn>ethernet-segment>qinq

Description This command determines the VIDs associated with the virtual Ethernet segment on a specific qinq port or LAG based on the following considerations:

- Values *, 0 to 4094 are allowed.
- Any SAP for which the service-delimiting qtag matches the range is associated with the virtual ES, and only those, for example, sap 1/1/1:0.* will not match port 1/1/1, s-tag-range 100.
- Maximum 8 ranges are allowed in the qinq context.
- A range can be comprised of a single qtag.
- Shutting down the ES is not required prior to making changes in the q-tag-range.

The **no** form of the command removes the configured range. Only the first qtag1 value is required to remove the range.

Parameters *qtag1* — Specifies the outer VID. When configuring a range of qtags (and not a single value), the second qtag1 value must be greater than the first qtag1.

Values *, 0 to 4094

s-tag

Syntax **s-tag** *qtag1* **c-tag-range** *qtag2* [**to** *qtag2*]
no s-tag *qtag1* **c-tag-range** *qtag2*

Context config>service>system>bgp-evpn>ethernet-segment>qinq

Description This command determines the inner VIDs (for a specified outer VID) associated with the virtual Ethernet segment on a specific qinq port or LAG based on the following:

- Values *, 0 to 4094 are allowed.

- Any SAP for which the outer and inner service-delimiting qtags match the range is associated with the virtual ES, and only those, for example, sap 1/1/1:10.* will not match port 1/1/1, s-tag 10 c-tag-range 10 to 100.
- A maximum of 8 ranges (including the s-tag ranges) are allowed in the qinq context.
- A c-tag range can be comprised of a single qtag.
- Shutting down the ES is not required prior to making changes.
- A qtag included in the **s-tag-range** command cannot be included in the s-tag qtag of this command.

Note: Not all qtag1 and qtag2 combinations are valid for values 0, *, and null. The following combinations are allowed:

- s-tag 0 c-tag-range *
- s-tag * c-tag-range *
- s-tag * c-tag-range null
- s-tag X c-tag-range 0 (where: X=1 to 4094)
- s-tag X c-tag-range * (where: X=1 to 4094)

The **no** form of the command removes the configured range. Only the first qtag1 value is required to remove the range.

Parameters *qtag1* — Specifies the outer VID for the c-tag range.

Values *, 0 to 4094

qtag2 — Specifies the inner VID for the c-tag range. When configuring a range of qtags (and not a single value), the second qtag1 value must be greater than the value of the first qtag1.

Values *, null, 0 to 4094

vc-id-range

Syntax **vc-id-range** *vc-id* [**to** *vc-id*]
no vc-id-range *vc-id*

Context config>service>system>bgp-evpn>ethernet-segment

Description This command determines the VC-IDs associated with the virtual Ethernet segment on a specific SDP based on the following considerations:

- VC-IDs for manual spoke-sdp and bgp-ad are included in the range.
- Th mesh-sdp VC-IDs are not allowed on a SDP used by a virtual ES.
- A maximum of 8 ranges are allowed.
- A range can be comprised of a single VC-ID.
- A **vc-id-range** can be comprised of a single VC-ID.

- Shutting down the ES is not required prior to making changes.

The **no** form of the command removes the configured range. Only the first VC-ID value is required to remove the range.

Parameters *vc-id* — Specifies the VC-ID. When configuring a range of VC-IDs (and not a single value), the value of the second VC-ID must be greater than the first VC-ID.

Values 1 to 4294967295

es-activation-timer

Syntax **es-activation-timer** *seconds*
no es-activation-timer

Context config>service>system>bgp-evpn>ethernet-segment

Description This command configures the Ethernet segment activation timer for a specified Ethernet segment. The **es-activation-timer** delays the activation of a specified **ethernet-segment** on a specified PE that has been elected as DF (Designated Forwarder). Only when the **es-activation-timer** has expired, the SAP/SDP-binding associated to an **ethernet-segment** can be activated (in case of single-active multi-homing) or added to the default-multicast-list (in case of all-active multi-homing).

If **no es-activation-timer** is configured, the system uses the value configured in the **config>redundancy>bgp-evpn-multi-homing>es-activation-timer** context, if configured. Otherwise the system uses a default value of 3 seconds.

Default no es-activation-timer

Parameters *seconds* — Specifies the number of seconds for the **es-activation-timer**.

Values 0 to 100

Default 3

esi

Syntax **esi** *value*
no esi

Context config>service>system>bgp-evpn>ethernet-segment

Description This command configures the 10-byte Ethernet segment identifier (ESI) associated to the ethernet-segment that will be signaled in the BGP-EVPN routes. The ESI value cannot be changed unless the ethernet-segment is shutdown. Reserved esi values (0 and MAX-ESI) are not allowed.

Parameters *value* — Specifies the 10-byte esi.

Values 00-11-22-33-44-55-66-77-88-99
Using any of these separators ('-',':')

lag

Syntax **lag** *lag-id*
no lag

Context config>service>system>bgp-evpn>ethernet-segment

Description This command configures a lag-id associated to the ethernet-segment. When the ethernet-segment is configured as **all-active.** only a lag can be associated to the ethernet-segment. When the ethernet-segment is configured as **single-active**, then a lag, port or sdp can be associated to the ethernet-segment. In either case, only one of the three objects can be configured in the ethernet-segment. A specified lag can be part of only one ethernet-segment.

Default no lag

Parameters *lag-id* — Specifies the lag-id associated with the ethernet-segment.

Values 1 to 800

multi-homing

Syntax **multi-homing single-active** [**no-esi-label**]
multi-homing all-active
no multi-homing

Context config>service>system>bgp-evpn>ethernet-segment

Description This command configures the multi-homing mode for the ethernet-segment as s**ingle-active** or all-active multi-homing, as defined in RFC7432.

By default, the use of **esi-labe**l is enabled for **all-active** and **single-active** as defined in RFC7432 (for **single-active multi-homing,** the esi-label is used to avoid transient loops).

When **single-active no-esi-label** is specified, the system will not allocate a label for the esi and hence advertise esi label 0 to peers. Even if the esi is configured to not send the esi-label, upon reception of an esi-label from a peer, the PE will always send traffic to that peer using the received esi-label.

Default no multi-homing

Parameters **single-active** — Configures single-active mode for the ethernet-segment.

all-active — Configures the system to not send an esi-label for **single-active** mode.

no-esi-label — Configures single-active mode for the ethernet-segment.

network-interconnect-vxlan

Syntax **network-interconnect-vxlan** *instance*
no network-interconnect-vxlan

Context config>service>system>bgp-evpn>ethernet-segment

Description This command associates the VXLAN instance with the virtual Ethernet segment. The association of the virtual ES is based on the VXLAN instance and range of services where the VXLAN instance is configured.

The **no** form of this command removes the VXLAN instance from the Ethernet segment association.

Parameters *instance* — Specifies the VXLAN instance that is to be associated with the virtual ES.

Values 1

port

Syntax **port** *port-id*
no port

Context config>service>system>bgp-evpn>ethernet-segment

Description This command configures a port-id associated with the ethernet-segment. If the ethernet-segment is configured as **all-active** only a lag can be associated to the ethernet-segment. If the ethernet-segment is configured as **single-active**, then a lag, port or sdp can be associated to the ethernet-segment. In any case, only one of the three objects can be configured in the ethernet-segment. A specified port can be part of only one ethernet-segment. Only Ethernet ports can be added to an ethernet-segment.

Default no port

Parameters *port-id* — Specifies the port ID associated to the ethernet-segment.

port-id	*slot*/*mda*/*port* [.*channel*]		
	eth-sat-id	esat-*id*/*slot*/*port*	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-*id*.*sub-port*	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

sdp

Syntax **sdp** *sdp-id*
no sdp

Context config>service>system>bgp-evpn>ethernet-segment

Description This command configures an *sdp-id* associated to the ethernet-segment. If the ethernet-segment is configured as **all-active** only a lag can be associated to the ethernet-segment. If the ethernet-segment is configured as **single-active**, then a lag, port or sdp can be associated to the ethernet-segment. In any case, only one of the three objects can be configured in the ethernet-segment. A specified SDP can be part of only one ethernet-segment. Only user-configured sdps can be added to an ethernet-segment.

Default no sdp

Parameters *sdp-id* — Specifies the IP address.

Values 1 to 17407

service-carving

Syntax **service-carving**

Context config>service>system>bgp-evpn>ethernet-segment

Description This command enables the context to configure service-carving in the ethernet-segment. The service-carving algorithm determines which PE is the Designated Forwarder (DF) in a specified Ethernet segment and for a specific service.

mode

Syntax **mode** {**manual** | **auto** | **off**}

Context config>service>system>bgp-evpn>ethernet-segment>service-carving

Description This command configures the service-carving mode. This determines how the DF is elected for a specified ethernet-segment and service.

Default mode auto

Parameters **auto** — This mode is the service-carving algorithm defined in RFC 7432. The DF for the service is calculated based on the modulo function of the service (identified by either the evi or the isid) and the number of PEs.

manual — In this mode the DF is elected based on the manual configuration added in the **service-carving>manual** context.

off — In this mode all the services elect the same DF PE (assuming the same PEs are active for all the configured services). The PE with the lowest IP is elected as DF for the ethernet-segment.

manual

Syntax **manual**

Context config>service>system>bgp-evpn>ethernet-segment>service-carving

Description This command enables the context to manually configure the service-carving algorithm, that is, configure the EVIs or ISIDs for which the PE is DF.

evi

Syntax **evi** *start* [**to** *to*]
no evi

Context config>service>system>bgp-evpn>ethernet-segment>service-carving>manual

Description This command configures the evi ranges for which the PE is primary, or uses the lowest preference algorithm.

Note: Multiple individual evi values and ranges are allowed.

There are two service-carving manual algorithms for DF election:

- Manual non-preference

  A **preference** command is not configured for this algorithm. The primary PE for the configured EVIs is determined by the EVI range. The manual non-preference algorithm only supports two PEs in the Ethernet segment
- Manual preference-based

  If a **preference** command is configured, the algorithm uses the configured value to determine the DF election. For EVIs not defined in the range, the highest-preference algorithm is used. For configured EVIs, the lowest-preference algorithm is used.

Parameters *start* — Specifies the initial evi value of the range.

Values 1 to 65535

to — Specifies the end evi value of the range. If not configured, only the individual start value is considered.

Values 1 to 65535

isid

Syntax **isid** *start* [**to** *to*]
no isid

Context config>service>system>bgp-evpn>ethernet-segment>service-carving>manual

Description This command configures the ISID ranges for which the PE is primary, or uses the lowest preference algorithm.

Note: Multiple individual ISID values and ranges are allowed.

The following service-carving manual algorithms are supported for DF election:

- Manual non-preference

  A **preference** command is not configured for this algorithm. The primary PE for the configured ISIDs is determined by the ISID range. The manual non-preference algorithm only supports two PEs in the Ethernet segment
- Manual preference-based

  If a **preference** command is configured, the algorithm uses the configured value to determine the DF election. For ISIDs not defined in the range, the highest-preference algorithm is used. For configured ISIDs, the lowest-preference algorithm is used.

Parameters *start* — Specifies the initial **isid** value of the range.

Values 1 to 16777215

to — Specifies the end **isid** value of the range. If not configured, only the individual start value is considered.

Values 1 to 16777215

preference

Syntax **preference** [**create**] [**non-revertive**]
no preference

Context config>service>system>bgp-evpn>ethernet-segment>service-carving>manual

Description This command creates the preference context for the Ethernet segment (ES) and determines whether the DF election for the ES is revertive or not. Creation of the **preference** context ensures that the PE will run the preference-based DF election algorithm.

Parameters **create** — Mandatory keyword required to create the preference context in an ES.

non-revertive — Configures a non-revertive ES, which ensures that when the Ethernet segment comes back after a failure, it does not take over an existing active DF PE.

value

Syntax **value** *value*

Context config>service>system>bgp-evpn>ethernet-segment>service-carving>manual>preference

Description This command modifies the default preference value used for the PE in the ES. An ES shut down is not required to modify this value during maintenance operations.

Default 32767

Parameters *value* — Determines the preference value used in the preference-based DF election algorithm.

Values 0 to 65535

service-id

Syntax **service-id**

Context config>service>system>bgp-evpn>ethernet-segment

Description This command enables the **service-id** context within the virtual **ethernet-segment** configuration.

service-range

Syntax **service-range** *svc-id* [**to** *svc-id*]
no service-range *svc-id*

Context config>service>system>bgp-evpn>ethernet-segment>service-id

Description This command associates a specified service range to a virtual ES, along with the **network-interconnect-vxlan** command. Up to eight service ranges per VXLAN instance can be configured, where the ranges may overlap. The service range may be configured before the service.

The **no** form of this command removes the association of the service range to the virtual ES for the configured VXLAN instance.

Parameters *svc-id* — Specifies which service range will be associated with the virtual Ethernet segment.

Values 1 to 2147483647

shutdown

Syntax **no shutdown**

Context config>service>system>bgp-evpn>ethernet-segment

Description This command changes the administrative status of the ethernet-segment.

The user can do **no shutdown** only when esi, multi-homing and lag/port/sdp are configured. If the ethernet-segment or the corresponding lag/port/sdp shutdown, the ethernet-segment route and the AD per-ES routes will be withdrawn. No changes are allowed when the ethernet-segment is **no shutdown**

Default shutdown

source-bmac-lsb

Syntax **source-bmac-lsb** *MAC-Lsb* [**es-bmac-table-size** *size*]
no source-bmac-lsb

Context config>service>system>bgp-evpn>ethernet-segment

Description This command configures the least significant two bytes of the BMAC used as the source BMAC for packets generated from the ethernet-segment in PBB-EVPN.

When the multi-homing mode is **all-active**, this value and the first high order four bytes must match on all the PEs that are part of the same ethernet-segment.

The **es-bmac-table-size** parameter modifies the default value (8) for the maximum number of virtual bmacs that can be associated to the ethernet-segment, that is, the es-bmacs. When the **source-bmac-lsb** is configured, the associated **es-bmac-table-size** is reserved out of the total FDB. The es-bmac will consume a separate BMAC per B-VPLS that is linked to an ethernet-segment

Parameters *MAC-Lsb* — Specifies the two least significant bytes of the es-bmac.

Values 1 to 65535, or xx-xx or xx:xx

size — Specifies the reserved space in the FDB for a specified es-bmac. By default the system reserves 8 entries for a specified ethernet-segment BMAC.

Values 1 to 511999

Default 8

vxlan

Syntax **vxlan**

Context config>service>system

Description This command enables the context where the vxlan global parameters are configured.

assisted-replication-ip

Syntax **assisted-replication-ip** *ip-address*
no assisted-replication-ip

Context config>service>system>vxlan

Description The assisted-replication-ip (AR-IP) command defines the IP address that supports the AR-R function in the router. The AR-IP address must also be defined as a loopback address in the base router and advertised in the IGP/BGP so that it is accessible to the remote NVE/PEs in the Overlay network.

If the AR-R function is enabled in a service, the Broadcast and Multicast frames encapsulated in VXLAN packets arriving at the router are replicated to the other VXLAN destinations within the service (except the destination pointing at the originator of the packet).

The **no** version of this command removes the AR IP address.

Default no assisted-replication-ip

Parameters *ip-address* — Specifies the assisted replication IP address.

tunnel-termination

Syntax **tunnel-termination** *ip-address* **fpe** *fpe-id* [**create**]
no tunnel-termination *ip-address*

Context config>service>system>vxlan

Description This command instructs the system to redirect traffic to the corresponding PXC interface associated with the configured Forwarding Path Extension (FPE) when the destination IP address matches the configured tunnel-termination IP address. The IP address is also registered, which enables the system to respond to ICMP packets directed to it.

Parameters *ip-address* — Specifies the non-system IPv4 or IPv6 address that terminates the VXLAN.

fpe *fpe-id* — Specifies the FPE identifier associated with the PXC port or LAG that will process and terminate the VXLAN.

Values 1 to 64

create — Creates the FPE.

redundancy

Syntax **redundancy**

Context config

Description This command enables the context to configure the global redundancy parameters.

bgp-evpn-multi-homing

Syntax **bgp-evpn-multi-homing**

Context config>redundancy
config>redundancy

Description This command enables the context to configure the bgp-evpn global timers

boot-timer

Syntax **boot-timer** *seconds*

Context config>redundancy>bgp-evpn-multi-homing

Description When the PE boots up, the **boot-timer** will allow the necessary time for the control plane protocols to come up before bringing up the ethernet-segments and running the DF algorithm.

The following considerations apply to the functionality:

- The boot-timer is configured at the system level **config>redundancy>bgp-evpn-multi-homing# boot-timer**. The configured value must provide enough time to allow the IOMs and BGP sessions to come up before exchanging ES routes and running the DF election for each EVI/ISID.
- The boot-timer is synchronized across CPMs and is relative to the System UP-time; hence it is not subject to change or reset upon CPM switchover.
- The boot-timer is never interrupted (the e**s-activation-timer**, however, can be interrupted if there is a new event triggering the DF election).
- The boot-timer runs per EVI/ISID on the ES's in the system. While **system-up-time** < **boot-timer** is true, the system does not run the DF election for any EVI/ISID. When the boot-timer expires, the DF election for the EVI/ISID is run and if the system is elected DF for the EVI/ISID, the **es-activation-timer** will kick-in.
- The system will **not** advertise ES routes until the boot timer has expired. This guarantees that the peer ES PEs do not run the DF election until the PE is ready to become the DF, if required.

Default boot-timer 10

Parameters *seconds* — Specifies the number of seconds for the boot-timer.

Values 0 to 600

es-activation-timer

Syntax **es-activation-timer** *seconds*

Context config>redundancy>bgp-evpn-multi-homing

Description This command configures the global ethernet-segment activation timer. The **es-activation-timer** delays the activation of a specified ethernet-segment on a specified PE that has been elected as DF (Designated Forwarder). Only when the **es-activation-timer** has expired, the SAP/SDP-binding associated to an ethernet-segment can be activated (in case of single-active multi-homing) or added to the default-multicast-list (in case of all-active multi-homing).

The **es-activation-timer** configured at the ethernet-segment level supersedes this global **es-activation-timer**.

Default es-activation-timer 3

Parameters *seconds* — Specifies the number of seconds for the **es-activation-timer**.

Values 0 to 100

accept-ivpls-evpn-flush

Syntax **accept-ivpls-evpn-flush**
no accept-ivpls-evpn-flush

Context config>service>vpls>bgp-evpn

Description This command enables the system to accept non-zero Ethernet tag MAC routes and process them only for CMAC flushing. This command can be changed on the fly without shutting down bgp-evpn mpls.

The **no** version of the command prevents the router from processing BMAC/ISID routes for cmac-flush.

Default no accept-ivpls-evpn-flush

cfm-mac-advertisement

Syntax [**no**] **cfm-mac-advertisement**

Context config>service>vpls>bgp-evpn

Description This command enables the advertisement and withdrawal, as appropriate, of the IEEE MAC address associated with the MP (MEP and MIP) created on a SAP, Spoke or Mesh, in an EVPN service.

The up-date occurs each time an MP is added or deleted, or an IEEE MAC address is changed for an MP on a SAP, Spoke or Mesh within the service. The size of the update depends on the number of MPs in the service affected by the modification.

Only enable this functionality, as required, for services that require a resident MAC address to properly forward unicast traffic and that do not perform layer two MAC learning as part of the dataplane.

Local MP IEEE MAC addresses are not stored in the local FDB and, as such, cannot be advertised through a control plane to a peer without this command.

The **no** version of the command disables the functionality and withdraws all previously advertised MP IEEE MAC addresses.

evi

Syntax **evi** *value*
[**no**] **evi**

Context config>service>vpls>bgp-evpn
config>service>epipe>bgp-evpn

Description This command allows you to specify a 2-byte EVPN instance unique in the system. It is used for the service-carving algorithm for multi-homing and auto-deriving route-target and route-distinguishers.

If not specified, the value will be zero and no route-distinguisher or route-targets will be auto-derived from it. If the *evi* value is specified and no other **route-distinguisher**/**route-target** are configured in the service, then the following rules apply:

- the route distinguisher is derived from <system_ip>:evi
- the route-target is derived from <autonomous-system>:evi

If vsi-import and export policies are configured, the route-target must be configured in the policies and those values take preference over the auto-derived route-targets. If **bgp-ad**>**vpls-id** and **bgp-evpn**>**evi** are both configured on the same service, the vpls-id auto-derived route-target/route-distinguisher takes precedence over the evi auto-derived ones The operational route-target for a service will be shown in the **show service id bgp** command.

The **no** version of the command will set the evi value back to zero.

Parameters *value* — Specifies the EVPN instance.

Values 1 to 65535

incl-mcast-orig-ip

Syntax **incl-mcast-orig-ip** *ip-address*

no incl-mcast-orig-ip

Context config>service>vpls>bgp-evpn>mpls

Description The IP address configured by the user in the **incl-mcast-orig-ip** command is encoded in the **originating-ip** field of EVPN Inclusive Multicast Routes with tunnel type Ingress Replication (value 6), mLDP (2), and Composite IR and mLDP (130).

The configured address does not need to be reachable in the base router or have an interface in the base router. The originating-ip address is used solely for BGP route-key selection.

The originating-ip is never changed for Inclusive Multicast Routes with tunnel type AR (Assisted Replication, value 10).

The **no** version of the command withdraws the affected Inclusive Multicast Routes and re-advertises it with the default system-ip address in the originating-ip field.

Default 1

Parameters *ip-address* — Specifies the IPv4 address value.

Values a.b.c.d

ingress-repl-inc-mcast-advertisement

Syntax [**no**] **ingress-repl-inc-mcast-advertisement**

Context config>service>vpls>bgp-evpn

Description This command enables and disables the advertisement of the Inclusive Multicast Ethernet Tag route (IMET route) with tunnel-type Ingress-Replication in the PMSI Tunnel Attribute, or with the tunnel-type Composite Point-to-Multipoint and Ingress-Replication (P2MP+IR) in the root-and-leaf nodes. The following considerations must be taken into account:

- When **no ingress-repl-inc-mcast-advertisement** is configured, no IMET routes will be sent for the service unless the **provider-tunnel** is configured with **owner bgp-evpn-mpls** and **root-and-leaf**, in which case, an IMET-P2MP route is sent.
- When **ingress-repl-inc-mcast-advertisement** and **provider-tunnel** are configured for **bgp-evpn-mpls** with **root-and-leaf**, the system will send an IMET-P2MP-IR route, that is, an IMET route with a composite P2MP+IR tunnel type.
- When **no ingress-repl-inc-mcast-advertisement** and **assisted-replication replicator** are configured, the system will send IMET-AR routes, but IMET-IR routes will not be sent.

Default ingress-repl-inc-mcast-advertisement

ip-route-advertisement

Syntax **ip-route-advertisement** [**incl-host**]

no ip-route-advertisement

Context config>service>vpls>bgp-evpn

Description This command enables and disables the advertisement of IP prefixes in EVPN.If enabled, any active route in the R-VPLS VPRN route table will be advertised in EVPN using the VPLS BGP configuration. The interface host addresses are not advertised in EVPN unless the **ip-route-advertisement incl-host** command is enabled.

Default no ip-route-advertisement

Parameters **incl-host** — Specifies to advertise the interface host addresses in EVPN.

isid-route-target

Syntax **isid-route-target**
no isid-route-target

Context config>service>vpls>bgp-evpn

Description This command enables the context for the configuration of isid-range to route-target associations.

isid-range

Syntax **isid-range** *from* [**to** *to*] {**auto-rt** | **route-target** *rt*}
no isid-range *from*

Context config>service>vpls>bgp-evpn>isid-route-target

Description This command creates a range of ISIDs associated with a specified route-target that is advertised with BMAC-ISID and IMET-ISID routes for the ISID. The route-target can be explicitly configured or automatically assigned by the system if the **auto-rt** option is configured. Auto routes assignment is based on RFC 7623 as follows:

<2-byte-as-number>:<4-byte-value>, where 4-byte-value = 0x30+ISID

The **no** form of the command deletes the **isid-range** and its association with the **route-target**.

The **no** form is the default action, which advertises the BMAC-ISID and IMET-ISID routes with the B-VPLS configured route-target.

Default no isid-range

Parameters *from* — Specifies the start of the ISID range.

Values 1 to 16777215

to — Specifies the end of the ISID range. If it is not configured, the range is comprised of (only) the ISID specified in the *to* option.

Values 1 to 16777215

auto-rt — Automatically generates an ISID-derived **route-target** in the format: AS_number:0x30+ISID.

route-target — Specifies an explicit route target.

Values rt - target:{<ip-addr:comm-val>|<2byte-as-number:extcomm-val>|<4byte-asnumber:comm-val>}

ip-addr: a.b.c.d

comm-val: [0 to 65535]

2byte-as-number: [0 to 65535]

ext-comm-val: [0 to 4294967295]

4byte-asnumber: [0 to 4294967295]

local-ac-name

Syntax [**no**] **local-ac-name** *ac-name*

Context config>service>epipe>bgp-evpn

Description This command enables and disables the context in which the local Ethernet tag value is configured.

Default no local-ac-name

Parameters *ac-name* — Specifies the name of the local attachment circuit.

remote-ac-name

Syntax [**no**] **remote-ac-name** *ac-name*

Context config>service>epipe>bgp-evpn

Description This command enables and disables the context in which the remote Ethernet tag value is configured.

Default no remote-ac-name

Parameters *ac-name* — Specifies the name of the remote attachment circuit.

eth-tag

Syntax [**no**] **eth-tag** *tag-value*

Context config>service>epipe>bgp-evpn>local-ac-name
config>service>epipe>bgp-evpn>remote-ac-name

Description This command configures the Ethernet tag value. When configured in the **local-ac-name** context, the system will use the value in the advertised AD per-EVI route sent for the attachment circuit. When configured in the **remote-ac-name context**, the system will compare that value with the eth-tag value of the imported AD per-EVI routes for the service. If there is a match, the system will create an EVPN destination for the Epipe.

Parameters *tag-value* — Specifies the Ethernet tag value of the attachment circuit.

Values 1 to 16777215

mac-advertisement

Syntax [**no**] **mac-advertisement**

Context config>service>vpls>bgp-evpn

Description The mac-advertisement command enables the advertisement in BGP of the learned macs on SAPs and SDP bindings. When the mac-advertisement is disabled, the local macs will be withdrawn in BGP.

Default mac-advertisement

mac-duplication

Syntax **mac-duplication**

Context config>service>vpls>bgp-evpn

Description This command enables the context to configure the BGP EVPN MAC duplication parameters.

detect

Syntax **detect num-moves** *num-moves* **window** *minutes*

Context config>service>vpls>bgp-evpn>mac-duplication

Description The **mac-duplication** featured is always enabled by default. This command modifies the default behavior. **mac-duplication** monitors the number of moves of a MAC address for a period of time (window).

Default num-moves 5 window 3

Parameters *num-moves* — Identifies the number of MAC moves in a VPLS service. The counter is incremented when a specified MAC is locally relearned in the FDB or flushed from the FDB due to the reception of a better remote EVPN route for that MAC.

Values 3 to 10

Default 3

minutes — Specifies the length of the window in minutes.

Values 1 to 15

Default 3

retry

Syntax **retry** *minutes*
no retry

Context config>service>vpls>bgp-evpn>mac-duplication

Description Specifies the timer after which the MAC in hold-down state is automatically flushed and the mac-duplication process starts again. This value is expected to be equal to two times or more than that of window.

If **no** retry is configured, this implies that, when mac-duplication is detected, MAC updates for that MAC will be held down till the user intervenes or a network event (that flushes the MAC) occurs.

Default 9

Parameters *minutes* — Specifies the BGP EVPN MAC duplication retry in minutes.

Values 2 to 60

black-hole-dup-mac

Syntax **black-hole-dup-mac**
no black-hole-dup-mac

Context config>service>vpls>bgp-evpn>mac-duplication

Description The **black-hole-dup-mac** command is disabled by default. If enabled, a duplicated MAC detected in the network is programmed as a black-hole MAC in the FDB and displayed in the **show service id fdb detail** command as follows:

- Source-Identifier—black-hole
- Type—EvpnD:P

Because the MAC is now programmed in the FDB as a black-hole, all received frames with MAC DA matching the duplicate MAC are discarded. The duplicate black-hole MACs are installed as Protected, therefore, all received frames with MAC SA matching the duplicate MAC are discarded by default.

A BGP-EVPN (MPLS or VXLAN) shutdown is required to add or remove the **black-hole-dup-mac** command.

The **no** form of the command removes the feature, and duplicate MACs are no longer programmed as black-hole MACs.

Default no black-hole-dup-mac

mpls

Syntax **mpls**

Context config>service>vpls>bgp-evpn
config>service>epipe>bgp-evpn

Description This command enables the context to configure the BGP EVPN MPLS parameters.

auto-bind-tunnel

Syntax **auto-bind-tunnel**

Context config>service>vpls>bgp-evpn>mpls
config>service>epipe>bgp-evpn>mpls

Description This command enables the context to configure automatic binding of a BGP-EVPN service using tunnels to MP-BGP peers.

The **auto-bind-tunnel** node is simply a context to configure the binding of EVPN routes to tunnels. The user must configure the **resolution** option to enable auto-bind resolution to tunnels in TTM. The following configurations are available:

- If the **resolution** option is explicitly set to **disabled**, the auto-binding to the tunnel is removed.
- If **resolution** is set to **any**, then any supported tunnel type in EVPN context will be selected following TTM preference.
- If one or more explicit tunnel types are specified using the **resolution-filter option**, then only these tunnel types will be selected again following the TTM preference.

resolution

Syntax **resolution** {**disabled** | **any** | **filter**}

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel

Description This command configures the resolution mode in the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.

Parameters **any** — Enables the binding to any supported tunnel type in a BGP-EVPN MPLS context following TTM preference.

disabled — Disables the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.

filter — Enables the binding to the subset of tunnel types configured under **resolution-filter**.

resolution-filter

Syntax **resolution-filter**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel

Description This command enables the context that allows the configuration of the subset of tunnel types that can be used in the resolution of BGP-EVPN routes within the automatic binding of BGP-EVPN MPLS service to tunnels to MP-BGP peers.

The following tunnel types are supported in a BGP-EVPN MPLS context in order of preference: RSVP, SR-TE, LDP, SR-ISIS, SR-OSPF, BGP and UDP.

The **ldp** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

The **rsvp** value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

When the **sr-isis** (**sr-ospf**) value is enabled, a SR tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF) instance.

The **sr-te** value instructs the code to search for the best metric SR-TE LSP to the address of the BGP next-hop. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple SR-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

The **bgp** value instructs BGP EVPN to search for a BGP LSP to the address of the BGP next-hop. If the user does not enable the BGP tunnel type, inter-area or inter-as prefixes will not be resolved.

The **udp** value instructs BGP EVPN to search for a UDP LSP to the address of the BGP next-hop.

Note: UDP tunnels are created through import policies with action **create-udp-tunnel**.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

bgp

Syntax [**no**] **bgp**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description Selects the BGP tunnel type.

ldp

Syntax [**no**] **ldp**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description Selects the LDP tunnel type.

rsvp

Syntax [**no**] **rsvp**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description Selects the RSVP-TE tunnel type.

sr-isis

Syntax [**no**] **sr-isis**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description Selects the Segment Routing (SR) tunnel type programed by an ISIS instance in TTM.

sr-ospf

Syntax [**no**] **sr-ospf**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description Selects the Segment Routing (SR) tunnel type programed by an OSPF instance in TTM.

sr-te

Syntax [**no**] **sr-te**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description Selects the Segment Routing (SR) Traffic Engineered (SR-TE) LSP programmed in TTM.

udp

Syntax [**no**] **udp**

Context config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description Selects the MPLS-over-UDP tunnel type programmed in TTM.

bgp-instance

Syntax **bgp-instance** [*bgp-instance*]

Context config>service>vpls>bgp-evpn>vxlan
config>service>vpls>bgp-evpn>mpls

Description This command associates **bgp-evpn vxlan** and/or **bgp-evpn mpls** to a specific BGP instance that has been previously created. While **bgp-evpn mpls** can be associated with instance 1 or 2, **bgp-evpn vxlan** can only be associated with instance 1.

Default 1

Parameters *bgp-instance* — Specifies the value associated with the instance.

Values 1, 2

control-word

Syntax [**no**] **control-word**

Context config>service>vpls>bgp-evpn>mpls
config>service>epipe>bgp-evpn>mpls

Description This command enables the transmission and reception of the **control-word**. As defined in RFC7432, the use of the control-word helps avoid frame disordering.

It is enabled or disabled for all EVPN-MPLS destinations at the same time.

Default no control-word

ecmp

Syntax **ecmp** *value*

Context config>service>epipe>bgp-evpn>mpls
config>service>vpls>bgp-evpn>mpls

Description This command controls the number of paths to reach a specified MAC address when that MAC in the FDB is associated to a remote all-active multi-homed ethernet-segment.

The configuration of 2 or more ecmp paths to a specified MAC enables the “aliasing” function described in RFC7432.

Parameters *value* — Specifies the number of paths allowed to the same multi-homed MAC address, assuming the MAC is located in an all-active multi-homed ethernet-segment.

Values 0 to 32

Default 0

entropy-label

Syntax [**no**] **entropy-label**

Context config>service>vpls>bgp-evpn>mpls
config>service>epipe>bgp-evpn>mpls

Description If **entropy-label** is configured, the Entropy label and Entropy Label Indicator are inserted in packets for which at least one LSP in the stack for the far-end of the tunnel used by the service has advertised entropy label capability. If the tunnel is RSVP type, **entropy-label** can also be controlled under the **config>router>mpls** or **config>router>mpls>lsp** context.

The entropy label is mutually exclusive with the hash label feature. The entropy label cannot be configured on a spoke-sdp or service where the hash label feature has already been configured unless no hash label is set, and vice-versa.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

Context config>service>vpls>bgp-evpn>mpls
config>service>epipe>bgp-evpn>mpls

Description This command allows the system to preserve the vlan-id and 802.1p bits of the service-delimiting qtag in a new tag added in the customer frame before sending it to the EVPN-MPLS destinations.

This command may be used in conjunction with the **sap ingress vlan-translation** command. If so used, the configured translated vlan-id will be the vlan-id sent to the EVPN-MPLS destinations as opposed to the service-delimiting tag vlan-id. If the ingress SAP/SDP binding is 'null'-encapsulated, the output vlan-id and pbits will be zero.

Default no force-vlan-vc-forwarding

ingress-replication-bum-label

Syntax [**no**] **no-ingress-replication-bum-label**

Context config>service>vpls>bgp-evpn>mpls

Description This command allows the user to configure the system so that a separate label is sent for BUM (Broadcast, Unknown unicast and Multicast) traffic in a specified service. By default (**no ingress-replication-bum-label**), the same label is used for unicast and flooded BUM packets when for-warding traffic to remote PEs.

When saving labels, this might cause transient traffic duplication for all-active multi-homing. By enabling **ingress-replication-bum-label**, the system will advertise two labels per EVPN VPLS instance, one for unicast and one for BUM traffic. The ingress PE will use the BUM label for flooded traffic to the advertising egress PE, so that the egress PE can determine if the unicast traffic has been flooded by the ingress PE. Depending on the scale required in the network, the user may choose between saving label space or avoiding transient packet duplication sent to an all-active multi-homed CE for certain macs.

Default no ingress-replication-bum-label

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>bgp-evpn>mpls
config>service>epipe>bgp-evpn>mpls

Description This command controls the administrative state of EVPN-MPLS in the service.

split-horizon-group

Syntax **split-horizon-group** *name*
no split-horizon-group

Context config>service>vpls>bgp-evpn>mpls

Description This command allows the user to configure an explicit split-horizon-group for all BGP-EVPN MPLS destinations that can be shared by other SAPs and/or spoke-SDPs. The use of explicit split-horizon-groups for EVPN-MPLS and spoke-SDPs allows the integration of VPLS and EVPN-MPLS networks.

If the **split-horizon-group** command for **bgp-evpn>mpls>** is not used, the default split-horizon-group (that contains all the EVPN destinations) is still used, but it is not possible to refer to it on SAPs/spoke-SDPs. User-configured split-horizon-groups can be configured within the service context. The same group-name can be associated to saps, spoke-sdps, pw-templates, pw-template-bindings and EVPN-MPLS destinations. The configuration of **bgp-evpn>mpls> split-horizon-group** will only be allowed if **bgp-evpn>mpls** is shutdown; no changes are allowed when bgp-evpn>mpls is **no shutdown**.

When the SAPs and/or spoke-SDPs (manual or BGP-AD-discovered) are configured within the same **split-horizon-group** as the EVPN-MPLS endpoints, MAC addresses will still be learned on them but they will not be advertised in BGP-EVPN. If provider-tunnel is enabled in the bgp-evpn service, the SAPs and SDP-bindings that share the same split-horizon-group of the EVPN-MPLS provider-tunnel will be brought operationally down if the point-to-multipoint tunnel is operationally up.

Default no split-horizon-group

Parameters *name* — Specifies the split-horizon-group name.

send-evpn-encap

Syntax **send-evpn-encap** [**mpls**] [**mplsoudp**]
no send-evpn-encap

Context config>service>vpls>bgp-evpn>mpls
config>service>vpls>bgp-evpn>vxlan

Description This command configures the encapsulation to be advertised with the EVPN routes for the service. The encapsulation is encoded in RFC5512-based tunnel encapsulation extended communities.

When used in the **bgp-evpn**>**mpls** context, the supported options are none (**no send-evpn-encap**), **mpls**, **mplsoudp** or both.

When used in the **bgp-evpn**>**vxlan** context, the supported options are **send-evpn-encap** (the router signals a VXLAN value) or **no send-evpn-encap** (no encapsulation extended community is sent).

Default send-evpn-encap mpls (in the config>service>vpls>bgp-evpn>mpls context)

send-evpn-encap (in the config>service>vpls>bgp-evpn>vxlan context)

Parameters **mpls** — Specifies the MPLS-over-UDP encapsulation value in the RFC5512 encapsulation extended community.

mplsoudp — Specifies the MPLS encapsulation value in the RFC5512 encapsulation extended community.

unknown-mac-route

Syntax [**no**] **unknown-mac-route**

Context config>service>vpls>bgp-evpn

Description This command enables the advertisement of the unknown-mac-route in BGP. This will be coded in an EVPN MAC route where the MAC address is zero and the MAC address length 48. By using this unknown-mac-route advertisement, the user may decide to optionally turn off the advertisement of MAC addresses learned from saps and sdp-bindings, hence reducing the control plane overhead and the size of the FDB tables in the data center. All the receiving NVEs supporting this concept will send any unknown-unicast packet to the owner of the unknown-mac-route, as opposed to flooding the unknown-unicast traffic to all other nodes part of the same VPLS. Although the 7750 SR, 7450 ESS, or 7950 XRS can be configured to generate and advertise the unknown-mac-route, the router will never honor the unknown-mac-route and will flood to the vpls flood list when an unknown-unicast packet arrives to an ingress sap/sdp-binding.

Use of the unknown-mac-route is only supported for BGP-EVPN VXLAN.

Default no unknown-mac-route

vxlan

Syntax **vxlan vni** *vni-id* [**create**] [**instance** *instance-id*]
no vxlan vni *vni-id*

Context config>service>vpls
config>service>epipe

Description This command enables the use of VXLAN in the VPLS or Epipe service.

Parameters *vni-id* — Specifies the VXLAN network identifier configured in the VPLS or Epipe service. In VPLS services, all the EVPN advertisements (MAC routes and inclusive multicast routes) will encode the configured VNI in the Ethernet Tag field of the NLRI.

Values 1 to 16777215

The VPLS service will be operationally up when the **vxlan vni** *vni-id* is successfully created. However, **bgp-evpn** must be enabled so that VXLAN bindings can be established and MAC learning and flooding can happen on them.

instance-id — Specifies the VXLAN instance ID.

Values 1

create — Mandatory keyword that creates a VXLAN instance.

egr-vtep

Syntax **egr-vtep** {*ip-address* | *ipv6-address*}
no egr-vtep

Context config>service>epipe>vxlan

Description This command configures the static destination VTEP IP used when originating VXLAN packets for the service.

Parameters *ip-address* — Specifies the IPv4 address used as the destination VTEP when originating VXLAN packets for the service.

ipv6-address — Specifies the IPv6 address used as the destination VTEP when originating VXLAN packets for the service.

oper-group

Syntax **oper-group** *name*
no oper-group

Context config>service>epipe>vxlan>egr-vtep

Description This command associates an operational group to the VXLAN static egress VTEP. If the egress VTEP IP disappears from the routing table, the oper-group status will become operationally down.

The operational group must be monitored in a different service and not in the service where it is defined.

Parameters *name* — Specifies the name of the **oper-group**, up to a maximum of 32 characters.

assisted-replication

Syntax **assisted-replication** {**replicator** | **leaf**} [**replicator-activation-time** *seconds*]
no assisted-replication

Context config>service>vpls>vxlan

Description This command enables the Assisted Replication (AR) function for VXLAN tunnels in the service. The execution of this command triggers the BGP EVPN to send an update containing the inclusive multicast route for the service and the AR type=AR Replicator (AR-R) or AR Leaf (AR-L).

The Replicators switch the VXLAN traffic back to VXLAN destinations when the IP destination address matches their own AR-IP address. Leaf nodes select a Replicator node and send all the Broadcast or Multicast frames to it so that the Replicator can replicate the traffic on their behalf.

Enabling or disabling the AR function, or changing the role between the replicator and leaf requires the BGP EVPN MPLS to be shutdown.

If the **leaf** parameter is configured, the system creates a Broadcast or Multicast (BM) destination to the selected AR-R and Unknown Unicast (U) destinations to the rest of the VTEPs. If no replicator exists, the leaf creates BUM bindings to all the VTEPs.

If the **replicator** parameter is configured, the system will create BUM destinations to the remote leafs, Regular Network Virtualization Edge routers (RNVE), and other AR-Rs. The system will perform assisted replication for traffic from known VTEPs only (that is, where the routes have been received and programmed toward a VTEP).

The **no** version of this command removes the AR function from the service.

Default no assisted-replication

Parameters **replicator-activation-time** *seconds* — Optional parameter that can be added to the leaf parameter. It specifies the wait time before the leaf can begin sending traffic to a new replicator and is used to allow some time for the replicator to learn about the leaf.

Values 1 to 255

Default 0 seconds (indicates **no replicator-activation-time** and no delay in sending packets to the AR-R)

replicator | leaf — Selects the AR role of the router for the service.

vxlan

Syntax **vxlan**

Context config>service>vpls>bgp-evpn

Description This command enables the context to configure the VXLAN parameters when BGP EVPN is used as the control plane.

send-imet-ir-on-ndf

Syntax **send-imet-ir-on-ndf**
no send-imet-ir-on-ndf

Context config>service>vpls>bgp-evpn>vxlan

Description This command controls the advertisement of Inclusive Multicast Ethernet Tag (IMET) routes for ingress replication in the case where the PE is Non-DF for a specified network interconnect VXLAN virtual ES. When enabled, the router will advertise IMET-IR routes even if the PE is NDF. This attracts BUM traffic but also speeds up convergence in case of DF failure.

The **no** form of this command withdraws the advertisement of the IMET-IR route on the network interconnect VXLAN NDF router.

Default send-imet-ir-on-ndf

vxlan-src-vtep

Syntax **vxlan-src-vtep** {*ip-address* | *ipv6-address*}
no vxlan-src-vtep

Context config>service>vpls
config>service>epipe

Description This command enables the router to use the configured IP address as the tunnel source IP address (source VTEP) when originating VXLAN-encapsulated frames for this service. This IP address is also used to set the BGP NLRI next-hop in EVPN route advertisements for the service.

Default no vxlan-src-vtep

Parameters *ip-address* — Specifies the non-system IPv4 address that terminates VXLAN for a service.

ipv6-address — Specifies the IPv6 address that terminates VXLAN for a service.

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>bgp-evpn>vxlan

Description This command enables and disables the automatic creation of VXLAN auto-bindings by BGP-EVPN.

Default shutdown

pbb

Syntax **pbb**

Context config>service>vpls

Description This command enables the context where the PBB parameters are configured.

leaf-source-bmac

Syntax **leaf-source-bmac** *ieee-address*
no leaf-source-bmac

Context config>service>pbb

Description This command enables the use of PBB-EVPN E-Tree. The **leaf-source-bmac** address must be configured before any I-VPLS E-Tree type can be created. The **leaf-source-bmac** address is used as the BMAC SA in all PBB frames that encapsulate customer frames generated from leaf AC SAPs and spoke-SDPs. When configured, the B-VPLS service accepts PBB traffic destined to the source BMAC as well as the leaf source BMAC address.

The **no** version of this command removes the leaf source BMAC address.

Default no leaf-source-bmac

Parameters *ieee-address* — Specifies the MAC address assigned to the leaf source BMAC. The parameter is entered in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

send-bvpls-evpn-flush

Syntax **send-bvpls-evpn-flush**
no send-bvpls-evpn-flush

Context config>service>vpls>pbb

Description This command triggers ISID-based CMAC-flush signaling in the PBB-EVPN. When the command is enabled in an I-VPLS service, a BMAC/ISID route is sent for the I-VPLS ISID.

Default no send-bvpls-evpn-flush

use-es-bmac

Syntax **use-es-bmac**

Context config>service>vpls>pbb

Description This command is only supported in B-VPLS instances where BGP-EVPN is enabled and controls the source BMAC used by the system for packets coming from the SAP or spoke-SDPs when they belong to an EVPN ethernet-segment.

If enabled, the system will use a source BMAC derived from the source-bmac (high order four bytes) and the least significant two bytes configured in **config>service>system>bgp-evpn>ethernet-segment>source-bmac-lsb** for all the packets coming from the local ethernet-segment.

If **no use-es-bmac** is configured, the system will use the regular source-bmac (provided by the **config>service>vpls>pbb>source-bmac** command, or the chassis bmac if the source-bmac is not configured).

Default no use-es-bmac

provider-tunnel

Syntax **provider-tunnel**

Context config>service>vpls

Description This command enables the context to configure the use of a P2MP LSP to forward Broadcast, Unknown unicast, and Multicast (BUM) packets of a VPLS or B-VPLS instance. The P2MP LSP is referred to as the Provider Multicast Service Interface (PMSI).

inclusive

Syntax **inclusive**

Context config>service>vpls>provider-tunnel

Description This command enables the context to configure the use of a P2MP LSP as the default tree for forwarding Broadcast, Unknown unicast, and Multicast (BUM) packets of a VPLS or B-VPLS instance. The P2MP LSP is referred to, in this case, as the Inclusive Provider Multicast Service Interface (I-PMSI).

When enabled, this feature relies on BGP Auto-Discovery (BGP-AD), BGP-VPLS or BGP-EVPN to discover the PE nodes participating in a specified VPLS/B-VPLS instance. In the case of BGP-AD or BGP-VPLS, the BGP route contains the information required to signal both point-to-point (P2P) PWs used to forward unicast known Ethernet frames, and the RSVP or mLDP P2MP LSP used to forward the BUM frames. In the case of BGP-EVPN, the EVPN IMET route contains the information to set up the mLDP P2MP LSP and may also contain the information that enables the remote leaf-only nodes to setup an EVPN destination to the sending PE.

Note: The provider-tunnel for a specified service must be configured with an owner protocol (BGP-AD, BGP-VPLS or BGP-EVPN); only one owner must be configured. Use the **owner {bgp-ad|bgp-vpls|bgp-evpn-mpls}** command to configure an owner.

With an mLDP I-PMSI, each leaf node will initiate the signaling of the mLDP P2MP LSP upstream using the P2MP FEC information in the I-PMSI tunnel information discovered through the BGP.

If IGMP or PIM snooping are configured on the VPLS/B-VPLS instance, multicast packets matching a L2 multicast Forwarding Information Base (FIB) record will also be forwarded over the P2MP LSP.

Use the **mldp** command to enable the use of an LDP P2MP LSP as the I-PMSI for forwarding Ethernet BUM and IP multicast packets in a VPLS instance:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>mldp

When a **no shutdown** is performed under the context of the inclusive node and the expiration of a delay timer, BUM packets will be forwarded over an automatically signaled mLDP P2MP LSP.

Use the **root-and-leaf** command to configure the node to operate as both root and leaf in the VPLS instance:

config>service>vpls [b-vpls]>provider-tunnel>inclusive>root-and-leaf

The node behaves as a leaf-only node by default. For the I-PMSI of type mLDP, the leaf-only node will join I-PMSI rooted at other nodes it discovered but will not include a PMSI Tunnel Attribute in BGP route update messages. This way a leaf-only node will forward packets to other nodes in the VPLS/B-VPLS using the point-to-point spoke-sdps in the case of BGP-AD or BGP-VPLS, or using EVPN destinations in the case of BGP-EVPN.

Note: Either BGP-AD/VPLS or BGP-EVPN must be enabled in the VPLS/B-VPLS instance otherwise the execution of the **no shutdown** command under the context of the inclusive node will fail and the I-PMSI will not come up.

If the P2MP LSP instance goes down, the VPLS/B-VPLS immediately reverts the forwarding of BUM packets to the P2P PWs or EVPN destinations (in the case of BGP-EVPN). Performing a shutdown under the context of the inclusive node will allow the user to restore BUM packet forwarding over the P2P PWs or EVPN destinations.

This feature is supported with VPLS and B-VPLS; it is not supported with I-VPLS. Although Routed VPLS is supported, routed traffic cannot be sent over the I-PMSI tree.

data-delay-interval

Syntax **data-delay-interval** *seconds*
no data-delay-interval

Context config>service>vpls>provider-tunnel>inclusive

Description This command enables the context to configure the I-PMSI data delay timer.

For an mLDP P2MP LSP, the delay timer is started as soon as the P2MP FEC corresponding to the I- PMSI is resolved and installed at the root node. When configuring a value at the root node, the user must factor the configured **IGP-LDP sync timer (config>router>if>ldp-sync-timer)** on the network interfaces. This is required because the mLDP P2MP LSP may move to a different interface at the expiry of the sync timer as the routing upstream of the LDP Label Mapping message may change when the sync timer expires and the interface metric is restored.

When the data delay timer expires, the VPLS/B-VPLS begins forwarding BUM packets over the P2MP LSP instance even if all the paths are not up.

The **no** version of this command reinstates the default value for the delay timer.

Parameters *seconds* — Specifies the delay-time in seconds.

Values 3 to 180

Default 15

mldp

Syntax [**no**] **mldp**

Context config>service>vpls>provider-tunnel>inclusive

Description This command enables the context to configure the parameters of an LDP P2MP LSP used for forwarding Broadcast, Unicast unknown and Multicast (BUM) packets of a VPLS or B-VPLS instance.

root-and-leaf

Syntax [**no**] **root-and-leaf**

Context config>service>vpls>provider-tunnel>inclusive

Description This command enables the node to operate as both root and leaf of the I-PMSI in a specified VPLS/B-VPLS instance.

By default, a node will behave as a leaf-only node. For the I-PMSI of type mLDP, the leaf-only node will join I-PMSI rooted at other nodes it discovered but will not include a PMSI tunnel attribute in BGP route update messages. This way a leaf-only node will forward packets to other nodes in the VPLS/B-VPLS using the point-to-point spoke-sdp's or the EVPN destinations.

The **no** version of the command reinstates the default value.

owner

Syntax [**no**] **owner** {**bgp-ad** | **bgp-vpls** | **bgp-evpn-mpls**}

Context config>service>vpls>provider-tunnel>inclusive

Description This command selects the owner protocol of the inclusive PMSI tunnel in the service. Only one of the protocols will support the provider tunnel.

The *bgp-vpls* and *bgp-evpn-mpls* parameters cannot be configured together in the same service. Although *bgp-ad* and *bgp-evpn* can coexist in the same service, *bgp-ad* cannot be configured as the owner of the provider-tunnel. In addition, the following applies to this configuration:

- The owner must be explicitly set before the provider-tunnel can be **no shutdown**.
- If the owner is *bgp-ad*, then bgp-evpn mpls and bgp-evpn vxlan will fail to **no shutdown**.
- The provider-tunnel must be shutdown to change the owner; on the fly change is not allowed.

Default no owner

Parameters *bgp-ad* — Specifies that bgp-ad is the owner of the provider-tunnel.

bgp-vpls — Specifies that bgp-vpls is the owner of the provider-tunnel.

bgp-evpn-mpls — Specifies that bgp-evpn-mpls is the owner of the provider-tunnel.

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>provider-tunnel>inclusive

Description This command administratively enables and disables the service.

proxy-arp

Syntax [**no**] **proxy-arp**

Context config>service>vpls

Description This command enables the context to configure the proxy-ARP parameters in a VPLS service.

Default no proxy-arp

proxy-nd

Syntax [**no**] **proxy-nd**

Context config>service>vpls

Description This command enables the context to configure the proxy-ND parameters in a VPLS service.

Default no proxy-arp

age-time

Syntax [**no**] **age-time** *seconds*

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description This command specifies the aging timer per proxy-ARP/proxy-ND entry for dynamic entries. When the aging expires, the entry is flushed. The age is reset when a new ARP/GARP/NA for the same MAC-IP is received. If the corresponding FDB MAC entry is flushed, the proxy-ARP/proxy-ND entry goes inactive and subsequent ARP/NS lookups are treated as "missed". EVPN will withdraw the IP→MAC if the entry goes inactive. The **age-time** should be set at *send-refresh * 3* to ensure that no active entries are unnecessarily removed.

Default no age-time

Parameters *seconds* — Specifies the age-time in seconds.

Values 60 to 86400

dup-detect

Syntax **dup-detect** [**anti-spoof-mac** *mac-address*] **window** *minutes* **num-moves** *count* **hold-down** [*minutes* | **max**]
dup-detect [**anti-spoof-mac** *mac-address*] **window** *minutes* **num-moves** *count* **hold-down** [*minutes* | **max**] [**static-black-hole**]

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description This command enables a mechanism that detects duplicate IPs and ARP/ND spoofing attacks. Attempts (relevant to dynamic and EVPN entry types) to add the same IP (different MAC) are monitored for **window** <*minutes*>. When <*count*> is reached within that **window**, the proxy-ARP/ND entry for the suspected IP is marked as duplicate. An alarm is also triggered. This condition is cleared when **hold-down** time expires (max does not expire) or a **clear** command is issued.

If the **anti-spoof-mac** is configured, the proxy-ARP/ND offending entry's MAC is replaced with this <*mac-address*> and advertised in an unsolicited GARP/NA for local SAP/SDP-bindings, and in EVPN to remote PEs. This mechanism assumes that the same **anti-spoof-mac** is configured in all the PEs for the same service and that traffic with destination **anti-spoof-mac** received on SAPs/SDP-bindings will be dropped. An ingress **mac-filter** may be configured to drop traffic to the **anti-spoof-mac**.

The **anti-spoof-mac** can also be combined with the **static-black-hole** option. To use a black-hole MAC entry for the **anti-spoof-mac** function in a proxy-ARP/ND service, the following must be configured:

- **static-black-hole** option for the **anti-spoof-mac**
- a static black-hole MAC using the same MAC address used for the **anti-spoof-mac: static-mac mac <mac-address> create black-hole** command.

When both **anti-spoof-mac** and **static-black-hole** commands are configured, the MAC is advertised in EVPN as Static. Locally, the MAC will be shown in the FDB as CStatic and associated with a black-hole.

The combination of the **anti-spoof-mac** and the **static-black-hole** options ensures that any frame arriving in the system with MAC DA=**anti-spoof-mac** will be discarded, regardless of the ingress endpoint type (SAP/SDP-binding or EVPN) and without the need for a filter.

If the user wants to redirect the traffic with MAC DA=**anti-spoof-mac** instead of discarding it, redirect filters should be configured on saps/sdp-bindings instead of the **static-black-hole** option.

If the **static-black-hole** option is not configured for the **anti-spoof-mac**, the behavior is as follows:

- The **anti-spoof-mac** is not programmed in the FDB.
- Any attempt to add a Static MAC (or any other MAC) with the **anti-spoof-mac** value will be rejected by the system.
- A mac-filter is needed to discard traffic with MAC DA=**anti-spoof-mac**.

Any changes to the configuration of **anti-spoof-mac** require proxy-arp or proxy-nd to first be shut down. See ARP/ND Snooping and Proxy Support for more information.

Default dup-detect window 3 num-moves 5 hold-down 9

Parameters **window** *minutes* — Specifies the window size in minutes.

Values 1 to 15

Default 3

count — Specifies the number of moves required so that an entry is declared duplicate.

Values 3 to 10

Default 5

hold-down *minutes* — Specifies the hold-down time for a duplicate entry.

Values 2 to 60

Default 9

hold-down max — Specifies permanent hold-down time for a duplicate entry.

mac-address — Specifies the optional anti-spoof-mac to use.

dynamic

Syntax **dynamic** *ip-address* [**create**]
no dynamic *ip-address*

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description This command creates a dynamic IP that can be associated to a MAC list. The configured dynamic IP is only converted to a dynamic entry when the resolve process for the IP has passed successfully.

A summary of the IP resolution process is as follows:

- A resolve message is sent for the configured IP as soon as the dynamic IP is configured. The message is sent with a configurable frequency of 1 to 60 minutes (using the **resolve** command); the default value is 5 minutes. The actual resolve interval is a “jittered” value of the configured interval.
- The resolve message is an ARP-request or NS message flooded to all the non-EVPN endpoints in the service, irrespective of the status of the **unknown-arp-request-flood-evpn** or **unknown-ns-flood-evpn** commands.
- The router sends resolve messages at the configured frequency until a dynamic entry for the IP is created in the proxy-ARP or proxy-ND table. The IP entry is created only if all of the following conditions are true.
  - An ARP, GARP, or NA message is received for the configured IP.
  - The associated MAC exists in the configured MAC list for the IP.

    If the MAC list is empty or not configured, the router does not create an entry for the IP.
- After a dynamic entry is created in the proxy-ARP or proxy-ND table, the IP->MAC entry is advertised in the EVPN.

The **no** form of the command deletes the dynamic IP and the associated proxy-ARP or proxy-ND entry, if it exists.

Parameters *ip-address* — Specifies the IPv4 or IPv6 address.

Values ip-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
where:
x: [0 to FFFF]H
d: [0 to 255]D

evpn-route-tag

Syntax **evpn-route-tag** *tag*
no evpn-route-tag

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description This command configures a local route tag that can be used on export policies to match MAC/IP routes generated by the proxy-ARP or proxy-ND module. For example, if a new active dynamic proxy-ARP entry is added to the proxy-ARP table and **evpn-route-tag** is 10, an export policy that matches on tag 10 and adds a site-of-origin community SOO-1, allows the router to advertise the MAC/IP route for the proxy-ARP entry with community SOO-1.

The **no** form of this command removes the route tag for the generated EVPN MAC/IP routes.

Parameters *tag* — Specifies the route tag, in either decimal or hexadecimal form.

Values 1 to 255

mac-list

Syntax **mac-list** *name*
no mac-list

Context config>service>vpls>proxy-arp>dynamic
config>service>vpls>proxy-nd>dynamic

Description This command associates a previously created MAC list to a dynamic IP. The MAC list is created using the **config**>**service**>**proxy-arp-nd**>**mac-list** command.

The **no** form of the command deletes the association of the MAC list and the dynamic IP.

Parameters *name* — The name of the MAC list previously created using the **config**>**service**>**proxy-arp-nd**>**mac-list** command.

resolve

Syntax **resolve** *minutes*

Context config>service>vpls>proxy-arp>dynamic
config>service>vpls>proxy-nd>dynamic

Description This command configures the frequency at which a resolve message is sent. The resolve message is an ARP-request or NS message flooded to all the non-EVPN endpoints in the service irrespective of the current status of the **unknown-arp-request-flood-evpn** or **unknown-ns-flood-evpn** commands.

Parameters *minutes* — Specifies the frequency in minutes at which the **resolve** message is issued.

Values 1 to 60

Default 5

dynamic-arp-populate

Syntax [**no**] **dynamic-arp-populate**

Context config>service>vpls>proxy-arp

Description This command enables the addition of dynamic entries to the proxy-ARP table (disabled by default). When executed, the system will populate proxy-ARP entries from snooped GARP/ ARP messages on SAPs/SDP-bindings. These entries will be shown as dynamic.

When disabled, dynamic-arp entries will be flushed from the proxy-ARP table. Enabling dynamic-arp-populate is only recommended in networks with a consistent configuration of this command in all the PEs.

Default no dynamic-arp-populate

garp-flood-evpn

Syntax [**no**] **garp-flood-evpn**

Context config>service>vpls>proxy-arp

Description This command controls whether the system floods GARP-requests / GARP-replies to the EVPN. The GARPs impacted by this command are identified by the sender's IP being equal to the target's IP and the MAC DA being broadcast.

The **no** form of the command only floods to local saps/binds but not to EVPN destinations.

Disabling this command is only recommended in networks where CEs are routers that are directly connected to the PEs. Networks using aggregation switches between the host/ routers and the PEs should flood GARP messages in the EVPN to ensure that the remote caches are updated and the BGP does not miss the advertisement of these entries.

Default garp-flood-evpn

send-refresh

Syntax [**no**] **send-refresh** *seconds*

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description If enabled, this command will make the system send a refresh at the configured time. A refresh message is an ARP-request message that uses 0s as sender's IP for the case of a proxy-ARP entry. For proxy-ND entries, a refresh is a regular NS message using the chassis-mac as MAC source-address.

Default no send-refresh

Parameters *seconds* — Specifies the send-refresh in seconds.

Values 120 to 86400

static

Syntax [**no**] **static** *ip-address ieee-address*

Context config>service>vpls>proxy-arp

Description This command configures static entries to be added to the table. A static MAC-IP entry requires the addition of the MAC address to the FDB as either learned or CStatic (conditional static MAC) in order to become active.

Parameters *ip-address* — Specifies the IPv4 address for the static entry.

ieee-address — Specifies a 48-bit MAC address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

table-size

Syntax **table-size** *table-size*

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description This command adds a table-size limit per service. By default, the table-size limit is 250; it can be set up to 16k entries per service. A non-configurable implicit high watermark of 95% and low watermark of 90% exists, per service and per system. When those watermarks are reached, a syslog/trap is triggered. When the system/service limit is reached, entries for a specified IP can be replaced (a different MAC can be learned and added) but no new IP entries will be added, regardless of the type (Static, evpn, dynamic). If the user attempts to change the **table-size** value to a value that cannot accommodate the number of existing entries, the attempt will fail.

Default 250

Parameters *table-size* — Specifies the table-size as number of entries for the service.

Values 1 to 16384

unknown-arp-request-flood-evpn

Syntax [**no**] **unknown-arp-request-flood-evpn**

Context config>service>vpls>proxy-arp

Description This command controls whether unknown ARP-requests are flooded into the EVPN network. By default, the system floods ARP-requests, including EVPN (with source squelching), if there is no active proxy-arp entry for the requested IP.

The **no** form of the command will only flood to local SAPs/SDP-bindings and not to EVPN destinations.

Default unknown-arp-request-flood-evpn

dynamic-nd-populate

Syntax [**no**] **dynamic-nd-populate**

Context config>service>vpls>proxy-nd

Description This command enables the addition of dynamic entries to the proxy-ND table. The command is disabled by default. When executed, the system will populate proxy-ND entries from snooped Neighbor Advertisement (NA) messages on SAPs/SDP-bindings, in addition to the entries coming from EVPN (if the EVPN is enabled). These entries will be shown as dynamic, as opposed to EVPN entries or static entries.

When disabled, dynamic-ND entries will be flushed from the proxy-ND table. Enabling **dynamic-nd-populate** is only recommended in networks with a consistent configuration of this command in all the PEs.

Default no dynamic-nd-populate

evpn-nd-advertise

Syntax **evpn-nd-advertise** {**host** | **router**}

Context config>service>vpls>proxy-nd

Description This command enables the advertisement of static or dynamic entries that are learned as host or routers (only one option is possible in a specified service), and determines the R flag (host or router) when sending Neighbor Advertisement (NA) messages for existing EVPN entries in the proxy-ND table.

This command cannot be modified without **proxy-nd** *shutdown*.

Default evpn-nd-advertise router

Parameters **host** — Enables the advertisement of static or dynamic entries that are learned as host.

router — Enables the advertisement of static or dynamic entries that are learned as routers.

host-unsolicited-na-flood-evpn

Syntax [**no**] **host-unsolicited-na-flood-evpn**

Context config>service>vpls>proxy-nd

Description This command controls whether the system floods host unsolicited Neighbor Advertisements to the EVPN. The NA messages impacted by this command are NA messages with the following flags: S=0 and R=0.

The **no** form of the command will only flood to local saps/binds but not to the EVPN destinations. This is only recommended in networks where CEs are routers that are directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood unsolicited NA messages in the EVPN to ensure that the remote caches are updated and the BGP does not miss the advertisement of these entries.

Default host-unsolicited-na-flood-evpn

router-unsolicited-na-flood-evpn

Syntax [**no**] **router-unsolicited-na-flood-evpn**

Context config>service>vpls>proxy-nd

Description This command controls whether the system floods router unsolicited Neighbor Advertisements to EVPN. The NA messages impacted by this command are NA messages with the following flags: S=0 and R=1.

The **no** form of the command will only flood to local saps/binds but not to EVPN destinations. This is only recommended in networks where CEs are routers directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood unsolicited NA messages in EVPN to ensure that the remote caches are updated and BGP does not miss the advertisement of these entries.

Default router-unsolicited-na-flood-evpn

static

Syntax **static** *ipv6-address ieee-address* {**host** | **router**}
no static *ipv6-address*

Context config>service>vpls>proxy-nd

Description This command configures static entries to be added to the table. A static MAC-IP entry requires the addition of the MAC address to the FDB as either dynamic or CStatic (Conditional Static MAC) in order to become active. Along with the IPv6 and MAC, the entry must also be configured as either host or router. This will determine if the received NS for the entry will be replied with the R flag set to 1 (router) or 0 (host).

Default router-unsolicited-na-flood-evpn

Parameters *ipv6-address* — Specifies the IPv6 address for the static entry.

ieee-address — Specifies a 48-bit MAC address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

host — Specifies that the entry is type “host”.

router — Specifies that the entry is type “router”.

unknown-ns-flood-evpn

Syntax [**no**] **unknown-ns-flood-evpn**

Context config>service>vpls>proxy-nd

Description This command controls whether unknown Neighbor Solicitation messages are flooded into the EVPN network. By default, the system floods NS (with source squelching) to SAPs/SDP-bindings including EVPN, if there is no active proxy-nd entry for the requested IPv6.

The **no** form of the command will only flood to local SAPs/SDP-bindings but not to EVPN destinations.

Default unknown-ns-flood-evpn

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description This command enables and disables the proxy-ARP and proxy-nd functionality. ARP/GARP/ND messages will be snooped and redirected to the CPM for lookup in the proxy-ARP/proxy-ND table. The proxy-ARP/proxy-ND table is populated with IP->MAC pairs received from different sources (EVPN, static, dynamic). When the **shutdown** command is issued, it flushes the dynamic/EVPN dup proxy-ARP/proxy-ND table entries and instructs the system to stop snooping ARP/ND frames. All the static entries are kept in the table as *inactive*, regardless of their previous *Status*.

Default shutdown

disable-send-bvpls-evpn-flush

Syntax **disable-send-bvpls-evpn-flush**
no disable-send-bvpls-evpn-flush

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command disables the ISID-based CMAC-flush indication when the corresponding SAP or spoke-SDP enters the operationally down state.

If the **send-bvpls-evpn-flush** is properly enabled, the **no** version of the command enables BMAC/ISID route updates to be sent when the SAP or spoke-SDP is operationally down.

Default no disable-send-bvpls-evpn-flush

static-mac

Syntax **static-mac**

Context config>service>vpls

Description A set of conditional static MAC addresses can be created within a VPLS supporting bgp-evpn. Conditional Static Macs are also supported in B-VPLS with SPBs. Unless they are configured as **black-hole**, conditional Static Macs are dependent on the SAP/SDP state.

This command allows the assignment of a set of conditional Static MAC addresses to a SAP/spoke-SDP or **black-hole**. In the FDB, the static MAC is then associated with the active SAP or spoke-SDP.

When configured in conjunction with SPBM services, Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

Static MACs configured in a bgp-evpn service are advertised as protected (EVPN will signal the MAC as protected).

mac

Syntax **mac** *ieee-address* [**create**] **sap** *sap-id* **monitor** *fwd-status*
mac *ieee-address* [**create**] **spoke-sdp** *sdp-id:vc-id*] **monitor** *fwd-status*
mac *ieee-address* [**create**] **black-hole**
no mac *ieee-address*

Context config>service>vpls>static-mac

Description This command assigns a conditional static MAC address entry to an SPBM B-VPLS SAP/ spoke-SDP or **black-hole**, allowing external MACs for single and multi-homed operation.

This command also assigns a conditional static MAC address entry to an EVPN VPLS SAP/ spoke-SDP or a black-hole on the 7450 ESS or 7750 SR.

When configured in conjunction with SPBM services, Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

Parameters *ieee-address* — Specifies the static MAC address to SAP/SDP-binding or **black-hole**.

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx). It cannot be all zeros.

sap-id — Specifies the SAP ID.

sdp-id — Specifies the SDP ID

vc-id — Specifies the virtual circuit ID.

create — This keyword is mandatory while creating a static MAC.

black-hole — This keyword creates a static FDB entry for the MAC address to black-hole traffic.

fwd-status — Specifies that this static MAC will be installed in the FDB based on the forwarding status of the SAP or spoke-SDP.

vxlan

Syntax **vxlan**

Context config>service>vprn

Description This command enables the context to configure VXLAN parameters in the VPRN.

tunnel-termination

Syntax **tunnel-termination** {*ip-address* | *ipv6-address*} **fpe** *fpe-id* [**create**]
no tunnel-termination {*ip-address* | *ipv6-address*}

Context config>service>vprn>vxlan

Description This command instructs the system to redirect traffic to the corresponding PXC interface associated with the configured FPE when the destination IP address matches the configured tunnel termination IP address. Because the IP address is registered, the system can respond to ICMP packets directed to it.

The **no** version of this command removes the non-system IP address as the tunnel termination IP address.

Parameters *ip-address* | *ipv6-address* — Specifies the non-system IPv4 or IPv6 address that terminates the VXLAN.

Values ip-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
where:
x: [0 to FFFF]H
d: [0 to 255]D

fpe-id — Specifies the FPE identifier associated with the PXC port or LAG that processes and terminates the VXLAN.

Values 1 to 64

create — Mandatory keyword to create the FPE.

evpn-tunnel

Syntax [**no**] **evpn-tunnel**

Context config>service>vprn>if>vpls

Description This command enables and disables the evpn-tunnel mode for the attached R-VPLS. When enabled, no IP address will be required under the same interface.

Default no evpn-tunnel

vsd-domain

Syntax **vsd-domain** *name*
no vsd-domain

Context config>service>vpls
config>service>vprn

Description This command associates a previously configured vsd-domain to an existing VPRN or VPLS service. The vsd-domain is a tag used between the VSD and the 7750 SR, 7450 ESS, or 7950 XRS to correlate configuration parameters to a service.

Parameters *name* — Specifies the vsd-domain name.

vsd

Syntax **vsd**

Context config>service
config>service

Description This command provides the context for the vsd configuration.

domain

Syntax **domain** *name* [**type** {**l2-domain** | **vrf-gre** | **vrf-vxlan** | **l2-domain-irb**}] [**create**]
no domain *name*

Context config>service>vsd

Description This command configures a vsd-domain that can be associated to a VPLS or VPRN service.

Parameters *name* — Specifies the name of the vsd-domain. 32 characters maximum.

l2-domain — Assigns the l2-domain type to the domain. l2-domain-type domains must be associated to a VPLS service.

vrf-gre — Assigns the vrf-gre type to the domain. vrf-gre-type domains must be associated to a VPRN service.

vrf-vxlan — Assigns the vrf-vxlan type to the domain. vrf-vxlan-type domains must be associated to a VPLS service.

l2-domain-irb — Assigns the l2-domain-irb type to the domain. l2-domain-irb-type domains must be associated to a VPLS service.

create — This keyword is mandatory when creating the vsd-domain.

description

Syntax **description** *description-string*
no description

Context config>service>vsd>domain

Description This command provides a description for a vsd-domain. This description must be added before the domain is activated using the **no shutdown** command.

Parameters *description-string* — Specifies the text for the description.

service-range

Syntax **service-range** *svc-id* **to** *svc-id*
no service-range

Context config>service>vsd

Description This command configures the range of service identifiers that the system allows for dynamic services configured by python, when the Nuage VSD sends the service configuration parameters for the VSD fully-dynamic integration model

Parameters *svc-id* — Specifies the start and end service identifier values.

Values 1 to 2147483647

shutdown

Syntax [**no**] **shutdown**

Context config>service>vsd>domain

Description This command enables or disables a domain. A description must be provided before no shutdown is executed.

system-id

Syntax **system-id** *name*
no system-id

Context config>system>vsd

Description This command configures the DC GW system-id that is used for the configuration from VSD. VSD will identify the DC GW based on this identifier, hence it must be unique per VSD.

Parameters *name* — Specifies the name of the DC GW.

xmpp

Syntax **xmpp**

Context config>system

Description This command provides the context for the xmpp configuration.

server

Syntax **server** *xmpp-server-name* [**domain-name** *fqdn*] [**username** *user-name*] [**password** *password*] [**create**] [**service-name** *service-name*]
server *xmpp-server-name* [**domain-name** *fqdn*] [**username** *user-name*] [**password** *password*] [**create**] [**router** *router-instance*]
no server *xmpp-server-name*

Context config>system>xmpp

Description This command configures the XMPP server as well as the Jabber ID that the 7750 SR, 7450 ESS, or 7950 XRS will use for the XMPP communication with the server. The system uses DNS to resolve the configured domain-name.

no server *name* will remove all the dynamic configurations in all the services.

Parameters *xmpp-server-name* — Specifies the name of the server in lower-case letters.

fqdn — Specifies the Fully Qualified Domain Name of the server.

user-name — Specifies the user-name part of the Jabber ID.

password — Specifies the password part of the Jabber ID's user.

create — Keyword used to create the server instance.

router-instance — Specifies the router name or service ID used to identify the router instance.

Values

router-instance: *router-name* or *vprn-svc-id*

router-name	"Base", "management"
vprn-svc-id	1 to 2147483647

Default Base

service-name — Specifies the service name, up to 64 characters.

shutdown

Syntax [**no**] **shutdown**

Context config>system>xmpp>server

Description This command enables or disables the communication with a specified XMPP server. When the xmpp server is properly configured, **no shutdown** instructs the system to establish a TCP session with the XMPP server through the management interface first. If it fails to establish communication, the 7750 SR, 7450 ESS, or 7950 XRS uses an in-band communication and its system IP as source IP address. Shutdown does not remove the dynamic configurations.

security

Syntax **security**

Context config>system

Description This command enables the context for the configuration of the security parameters in the system.

cli-script

Syntax **cli-script**

Context config>system>security

Description This command enables the context for the configuration of the security parameters in the system.

authorization

Syntax **authorization**

Context config>system>security>cli-script

Description This command enables the context for the configuration of the authorization parameters for the cli-scripts in the system.

vsd

Syntax **vsd**

Context config>system>security>cli-script>authorization

Description This command enables the context for the configuration of the authorization parameters related to VSD in the system.

cli-user

Syntax **cli-user** *user-name*
no cli-user

Context config>system>security>cli-script>authorization>vsd

Description This command configures the cli-user for the configuration coming from VSD (fully dynamic VSD integration model). The user-profile determines what CLI set of commands can be executed by the VSD. This set of commands is a sub-set of the white-list of commands allowed by the system for the or VSD. You can use the **tools dump service vsd-services command-list** to check the white-list of commands.

Parameters *user-name* — Specifies the user-name that the VSD will use when adding a configuration to the system.

password

Syntax **password**

Context config>system>security>password

Description This command enables the context for the configuration of the passwords in the system.

vsd-password

Syntax **vsd-password** *password* [{**hash** | **hash2**}]
no vsd-password

Context config>system>security>password

Description This command configures the password required to access the **enable-vsd-config** mode. The **enable-vsd-config** mode allows editing of services provisioned by the VSD in fully dynamic mode (or by the **tools perform service vsd evaluate-script** command

Parameters *password* — Specifies the password required to login as authorized user in the **enable-vsd-config** mode.

hash — Specifies that the primary hashing sequence should be used.

hash2 — Specifies that the secondary hashing sequence should be used.

router

Syntax **router**

Context config

Description This command enables the context for the configuration of the base router in the system.

bgp

Syntax **bgp**

Context config>router

Description This command enables the context for the configuration of the base router BGP parameters in the system.

group

Syntax **group** *name*

Context config>router>bgp

Description This command enables the context for the configuration of a BGP group in the base router.

Parameters *name* — Specifies the name of the BGP group.

neighbor

Syntax **neighbor** *ip-address*

Context config>router>bgp>group

Description This command enables the context for the configuration of a BGP group neighbor in the base router.

Parameters *ip-address* — Specifies the IP address of the BGP group neighbor.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x: 0 to FFFF (hexadecimal)
	d: 0 to 255 (decimal)
	interface: 32 characters max. Mandatory for link local addresses

def-recv-evpn-encap

Syntax **def-recv-evpn-encap** {**mpls** | **vxlan**}

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command defines how the BGP will treat a received EVPN route without RC5512 BGP encapsulation extended community. If no encapsulation is received, BGP will validate the route as MPLS or VXLAN depending on how this command is configured.

Default no def-recv-evpn-encap

Parameters **mpls** — Specifies that **mpls** is the default encapsulation value in the case where no RFC5512 extended community is received in the incoming BGP-EVPN route.

vxlan — Specifies that **vxlan** is the default encapsulation value.

python

Syntax **python**

Context config

Description This command enables the context for the configuration of the Python parameters in the system.

python-policy

Syntax **python-policy** *name*

Context config>python

Description This command enables the context for the configuration of the Python policy parameters in the system.

Parameters *name* — Specifies the name of the Python policy.

vsd

Syntax **vsd script** *script*
no vsd

Context config>python

Description This command defines the python script for the Python policy sent by the VSD.

Parameters *script* — Specifies the VSD script that points at the python-script command.

enable-vsd-config

Syntax [**no**] **enable-vsd-config**

Context <root>

Description This command allows editing of VSD services just like normal services. As this is an action that should only be executed by authorized personnel, the activation of this command is protected by the use of a password, defined under **configure system security password vsd-password**.

5.6.2.2 Show Configuration Commands

provider-tunnel-using

Syntax **provider-tunnel-using leaf-only** [**bgp-ad** | **bgp-vpls** | **bgp-evpn-mpls**]
provider-tunnel-using root-and-leaf [**bgp-ad** | **bgp-vpls** | **bgp-evpn-mpls**]

Context show>service

Description This command displays the list of provider tunnels existing in the router for all services. The output can be filtered based on the provider tunnel owner.

Parameters **leaf-only** — Displays the leaf-only provider tunnels for all services.

root-and-leaf — Displays the root and leaf provider tunnels for all services.

bgp-ad — Filters the provider tunnels owned by BGP AD services.

bgp-vpls — Filters the provider tunnels owned by BGP VPLS services.

bgp-evpn-mpls — Filters the provider tunnels owned by BGP EVPN-MPLS services.

Output

Sample Output

```
A:PE-76# show service provider-tunnel-using root-and-leaf
=======================================================
Provider-Tunnel Using (Root-and-Leaf)
=======================================================
SvcId      SdpId             Owner       Admin Oper
                                         State State
-------------------------------------------------------
300        32767:4294967294  bgpEvpnMpls Up    Up
-------------------------------------------------------
Number of Root-and-Leaf : 1
```

```
=====================================================
A:PE-76# show service provider-tunnel-using root-and-leaf bgp-evpn-mpls
=====================================================
Provider-Tunnel Using (Root-and-Leaf)
=====================================================
SvcId      SdpId             Owner       Admin Oper
                                         State State
-----------------------------------------------------
300        32767:4294967294  bgpEvpnMpls Up    Up
-----------------------------------------------------
Number of Root-and-Leaf : 1
=====================================================
```

proxy-arp-nd

Syntax **proxy-arp-nd**

Context show>service

Description This command enables the context to configure the service-level **proxy-arp-nd** commands.

mac-list

Syntax **mac-list**
mac-list *name*
mac-list *name* **associations**

Context config>service>proxy-arp-nd

Description This command displays MAC address list information including MAC lists, MAC list details, and associations used in the **proxy-arp-nd** context.

Parameters *name* — Name of the MAC address list for which the detailed information is shown; the name can be up to 32 characters.

associations — Mandatory keyword to display the service ID and dynamic IP to which the MAC list is associated.

Output

Sample Output

```
*A:PE-3# show service proxy-arp-nd mac-list
===============================================================================
MAC List Information
===============================================================================
MAC List Name                    Last Change         Num Macs    Num Assocs
-------------------------------------------------------------------------------
list-1                           12/20/2016 09:21:13  3           1
-------------------------------------------------------------------------------
Number of Entries: 1
```

```
===============================================================================
*A:PE-3# show service proxy-arp-nd mac-list "list-1"
===============================================================================
MAC List MAC Addr Information
===============================================================================
MAC Addr                                Last Change
-------------------------------------------------------------------------------
00:ca:fe:ca:fe:01                       12/20/2016 09:21:13
00:ca:fe:ca:fe:02                       12/20/2016 09:21:13
00:ca:fe:ca:fe:03                       12/20/2016 09:21:13
-------------------------------------------------------------------------------
Number of Entries: 3

===============================================================================
*A:PE-3# show service proxy-arp-nd mac-list "list-1" associations
===============================================================================
MAC List Associations
===============================================================================
Service Id                    IP Addr
-------------------------------------------------------------------------------
5                             10.0.0.1
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

service-using

Syntax **service-using** [**vsd**]
service-using [**origin vsd**]

Context show>service

Description This command displays the services matching certain usage properties. If no optional parameters are specified, all services defined on the system are displayed.

Parameters **vsd** — Displays the VSD domain tags used and the associated service identifier.

origin vsd — Displays the services created by the VSD fully-dynamic integration model. Python will create the service after receiving the relevant parameters from VSD.

Output

Sample Output

```
*A:PE1# show service service-using vsd
==========================================
Services-using VSD Domain
==========================================
Svc Id     Domain
------------------------------------------
64000      L2-DOMAIN-5
------------------------------------------
Number of services using VSD Domain: 1
==========================================
```

```
*A:PE1# show service service-using origin vsd
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
64000        VPLS      Up   Up   1          evi64000
-------------------------------------------------------------------------------
Matching Services : 1
-------------------------------------------------------------------------------
===============================================================================
```

system

Syntax **system**

Context show>service

Description This command enables the context to display the system parameters.

bgp-evpn

Syntax **bgp-evpn** [**ethernet-segment**]
bgp-evpn ethernet-segment name *name* [**all**] [**evi** *evi*] [**isid** *isid*]

Context show>service>system

Description This command shows all the information related to the base EVPN instance, including the base RD used for ES routes, the ethernet-segments or individual ethernet-segment information.

Parameters **ethernet-segment** — Displays information for Ethernet segments.

name — Specifies the name of an Ethernet segment for which to show information. 28 characters maximum.

all — Displays all available information for the specified Ethernet segment.

evi — Displays information for the specified EVI.

Values 1 to 65535

isid — Displays information for the specified ISID.

Values 1 to 16777215

Output

Sample Output

```
*A:PE1# show service system bgp-evpn
```

```
===============================================================================
Service BGP EVPN Information
===============================================================================
Evpn Route Dist.   : 192.0.2.69:0
===============================================================================

*A:PE1# show service system bgp-evpn ethernet-segment
===============================================================================
Service Ethernet Segment
===============================================================================
Name                            ESI                             Admin     Oper
-------------------------------------------------------------------------------
ESI-71                          01:00:00:00:00:71:00:00:00:01 Enabled   Up
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================

*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-71" all
===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-71
Admin State             : Enabled           Oper State         : Up
ESI                     : 01:00:00:00:00:71:00:00:00:01
Multi-homing            : allActive         Oper Multi-homing  : allActive
Source BMAC LSB         : 71-71
ES BMac Tbl Size        : 8                 ES BMac Entries    : 1
Lag Id                  : 1
ES Activation Timer     : 0 secs
Exp/Imp Route-Target    : target:00:00:00:00:71:00

Svc Carving             : auto
ES SHG Label            : 262142
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   no
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.69
1                                       192.0.2.72
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------

===============================================================================
ISID Information
===============================================================================
```

```
ISID                SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
20001               20001               0                   no
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
ISID                                    DF Address
-------------------------------------------------------------------------------
20001                                   192.0.2.69
20001                                   192.0.2.72
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
BMAC Information
===============================================================================
SvcId                                   BMacAddress
-------------------------------------------------------------------------------
20000                                   00:00:00:00:71:71
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

ethernet-segment

Syntax **ethernet-segment**
ethernet-segment name *name* [**all**]
ethernet-segment name *name* **evi** [*evi*]
ethernet-segment name *name* **isid** [*isid*]
ethernet-segment name *name* **virtual-ranges**

Context show>service>system>bgp-evpn

Description This command enables the context to display the ethernet-segment parameters.

Parameters *name* — Specifies the name of an Ethernet segment for which to show information; maximum28 characters are allowed.

all — Displays all available information for the specified Ethernet segment.

evi — Displays information for the specified EVI.

Values 1 to 65535

isid — Displays information for the specified ISID.

Values 1 to 16777215

virtual-ranges — Displays the vc-id, qtag, s-tag, or c-tag per s-tag virtual ranges configured on the virtual Ethernet segment.

Output

Sample Output

```
*A:PE-2# show service system bgp-evpn ethernet-segment name "vES-23"
===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : vES-23
Eth Seg Type            : Virtual
Admin State             : Enabled           Oper State         : Up
ESI                     : 01:23:23:23:23:23:23:23:23:23
Multi-homing            : allActive         Oper Multi-homing  : allActive
ES SHG Label            : 262141
Source BMAC LSB         : 00-23
ES BMac Tbl Size        : 8                 ES BMac Entries    : 0
Lag Id                  : 1
ES Activation Timer     : 3 secs (default)
Svc Carving             : manual            Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref
-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  100            12/20/2016 09:21:08      100            Enabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
===============================================================================
*A:PE-2# show service system bgp-evpn ethernet-segment name "vES-23" evi
===============================================================================
EVI Information
===============================================================================
EVI                SvcId              Actv Timer Rem      DF
-------------------------------------------------------------------------------
5                  5                  0                   yes
30                 30                 0                   yes
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
*A:PE-2# show service system bgp-evpn ethernet-segment name "vES-23" evi 5
===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
5             5             0                   yes 12/20/2016 09:21:24
===============================================================================
===============================================================================
DF Candidates                           Time Added
-------------------------------------------------------------------------------
192.0.2.2                               12/20/2016 09:21:21
192.0.2.3                               12/20/2016 09:21:52
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

```
*A:PE-2# show service system bgp-evpn ethernet-segment name "vES-23" virtual-ranges
===============================================================================
Q-Tag Ranges
===============================================================================
Q-Tag Start         Q-Tag End           Last Changed
-------------------------------------------------------------------------------
5                   11                  12/20/2016 09:21:08
30                  30                  12/20/2016 09:21:08
-------------------------------------------------------------------------------
Number of Entries: 2
===============================================================================
===============================================================================
VC-Id Ranges
===============================================================================
VC-Id Start         VC-Id End           Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================
===============================================================================
S-Tag Ranges
===============================================================================
S-Tag Start         S-Tag End           Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================
===============================================================================
S-Tag C-Tag Ranges
===============================================================================
S-Tag Start         C-Tag Start         C-Tag End       Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================
===============================================================================
Vxlan Instance Service Ranges
===============================================================================
Svc Range Start          Svc Range End            Last Changed
-------------------------------------------------------------------------------
500                      500                      06/07/2017 15:10:59
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

vsd

Syntax **vsd**

Context show>service

Description This command enables the context for the vsd parameters.

domain

Syntax **domain** *domain-name* **association**

Context show>service>vsd

Description This command shows all the parameters related to a VSD domain created by the user or by VSD.

Parameters *domain-name* — Specifies the name of the VSD domain. 64 characters maximum.

association — Displays associations for the specified VSD domain.

Output

Sample Output

```
*A:PE71(1)# show service vsd domain
===============================================================================
VSD Domain Table
===============================================================================
Name                                Type           Origin    Admin
-------------------------------------------------------------------------------
L2-DOMAIN-5                         l2Domain       vsd       inService
-------------------------------------------------------------------------------
Number of domain entries: 1
===============================================================================


*A:PE71(1)# show service vsd domain "L2-DOMAIN-5"
===============================================================================
VSD Information
===============================================================================
Name               : L2-DOMAIN-5
Description        : L2-DOMAIN-5
Type               : l2Domain                  Admin State   : inService
Last Error To Vsd  : (Not Specified)
Last Error From Vsd: (Not Specified)

Statistics
-------------------------------------------------------------------------------
Last Cfg Chg Evt   : 07/15/2015 21:20:23       Cfg Chg Evts  : 1
Last Cfg Update    : 07/15/2015 21:20:23       Cfg Upd Rcvd  : 1
Last Cfg Done      : 07/15/2015 21:20:23
Cfg Success        : 1                         Cfg Failed    : 0
Last Recd Params   : script = {'domain' : '', 'vn
                   : i' : '64000', 'rt' : 'target
                   : :64000:64000', 'rte' : 'targ
                   : et:64000:64000', 'servicetyp
                   : e' : 'L2DOMAIN', 'metadata'
                   : : 'rd=1:1, sap=1/1/10:3000 '
                   : }
Last Exec Params   : script = {'domain' : '', 'vn
                   : i' : '64000', 'rt' : 'target
                   : :64000:64000', 'rte' : 'targ
                   : et:64000:64000', 'servicetyp
                   : e' : 'L2DOMAIN', 'metadata'
```

```
                    : : 'rd=1:1, sap=1/1/10:3000 '
                    : }
===============================================================================


*A:PE71(1)# show service vsd domain "L2-DOMAIN-5" association
============================================================
Service VSD Domain
============================================================
Svc Id      Svc Type  Domain Type    Domain Admin   Origin
------------------------------------------------------------
64000       vpls      l2Domain       inService      vsd
------------------------------------------------------------
Number of entries: 1
============================================================
```

root-objects

Syntax **root-objects**

Context show>service>vsd

Description This command displays the root objects created by vsd.

Output

Sample Output

```
*A:PE1# show service vsd root-objects
===============================================================================
VSD Dynamic Service Root Objects
===============================================================================
OID Prefix                  : svcRowStatus
OID index                   : .64000
Snippet name                : script
Snippet instance            : L2-DOMAIN-5
Orphan time                 : N/A
-------------------------------------------------------------------------------
No. of Root Objects: 1
===============================================================================
```

script

Syntax **script**

Context show>service>vsd

Description This command enables the context to show dynamic services script information.

snippets

Syntax **snippets** [**detail**]

Context show>service>vsd>script

Description This command displays the dynamic services snippets information. The CLI output generated by a single VSD service Python function call is a snippet instance.

Parameters **detail** — Displays detailed information.

Output

Sample Output

```
*A:PE1# show service vsd script snippets name "script"
===============================================================================
VSD Dynamic Services Snippets
===============================================================================
Name                        Instance                      Ref-count  Dict-len
-------------------------------------------------------------------------------
script                      L2-DOMAIN-5                   0          126
-------------------------------------------------------------------------------
No. of Snippets: 1
===============================================================================


*A:PE1# show service vsd script snippets name "script" detail
===============================================================================
VSD Dynamic Service Snippets
===============================================================================
Snippet            : script:L2-DOMAIN-5
-------------------------------------------------------------------------------
reference-count    : 0
dictionary-length  : 126

Root-object
-----------
oid                : 0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0

Reserved-id
-----------
id                 : service-id:64000
-------------------------------------------------------------------------------
No. of Snippets: 1
===============================================================================
```

statistics

Syntax **statistics**

Context show>service>vsd>script

Description This command displays vsd service script statistics. Only non-zero values are shown. The script statistics can be cleared with the "**clear service statistics vsd**" command.

Output

Sample Output

```
*A:PE1# show service vsd script statistics
===============================================================================
VSD Dynamic Services Script Statistics
===============================================================================
Description                                                          Counter
-------------------------------------------------------------------------------
python scripts with 0 retries due to timeout                         1
setup    - jobs launched                                             1
setup    - jobs handled                                              1
setup    - success                                                   1
-------------------------------------------------------------------------------
No. of VSD Script Statistics: 4
-------------------------------------------------------------------------------
Last Cleared Time: N/A
===============================================================================
```

summary

Syntax **summary**

Context show>service>vsd

Description This command displays the global configuration summary for vsd services.

Output

Sample Output

```
*A:PE1# show service vsd summary
===============================================================================
VSD Information
===============================================================================
Service Range
Start              : 64000                        End            : 65000
===============================================================================
===============================================================================
VSD Domain Table
===============================================================================
Name                                  Type           Origin    Admin
-------------------------------------------------------------------------------
L2-DOMAIN-5                           l2Domain       vsd       inService
-------------------------------------------------------------------------------
Number of domain entries: 1
===============================================================================
```

bgp

Syntax **bgp** *bgp-instance*

Context show>service>id

Description This command displays all the information for a specified BGP instance in a service.

Parameters *bgp-instance* — Specifies the BGP instance.

Output

Sample Output

```
*A:PE-1# show service id 7000 bgp 1
===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 1:1
Oper Route Dist      : 1:1
Oper RD Type         : configured
Rte-Target Import    : None                 Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi           Oper RT Import   : 64500:7000
Oper RT Exp Origin   : derivedEvi           Oper RT Export   : 64500:7000
PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
*A:PE-1# show service id 7000 bgp 2
===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 2:2
Oper Route Dist      : 2:2
Oper RD Type         : configured
Rte-Target Import    : None                 Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi           Oper RT Import   : 64500:7000
Oper RT Exp Origin   : derivedEvi           Oper RT Export   : 64500:7000
-------------------------------------------------------------------------------
===============================================================================
```

bgp-evpn

Syntax **bgp-evpn**

Context show>service>id

Description This command displays the **bgp-evpn** configured parameters for a specified service, including the admin status of VXLAN, the configuration for mac-advertisement and unknown-mac-route, as well as the mac-duplication parameters. The command shows the duplicate MAC addresses that mac-duplication has detected.

This command also shows whether the **ip-route-advertisement** command (and the **incl-host** parameter) is enabled. If the service is BGP-EVPN MPLS, the command also shows the parameters corresponding to EVPN-MPLS.

Output

Sample Output

```
# bgp-evpn vxlan service

*A:DutA# show service id 1 bgp-evpn
===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
VXLAN Admin Status : Enabled            Creation Origin    : manual
MAC Dup Detn Moves : 5                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 9                  Number of Dup MACs : 1
IP Route Advertise*: Enabled            Include hosts      : Disabled
-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses             Time Detected
-------------------------------------------------------------------------------
00:12:12:12:12:00                            01/17/2014 16:01:02
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
BGP EVPN MPLS Information
===============================================================================
Admin Status       : Disabled
Force Vlan Fwding  : Disabled           Control Word       : Disabled
Split Horizon Group: (Not Specified)
Ingress Rep BUM Lbl: Disabled           Max Ecmp Routes    : 0
Ingress Ucast Lbl  : N/A                Ingress Mcast Lbl  : N/A
Entropy Label      : Disabled
===============================================================================
===============================================================================
BGP EVPN MPLS Auto Bind Tunnel Information
===============================================================================
Resolution         : disabled
Filter Tunnel Types: (Not Specified)
===============================================================================

# bgp-evpn mpls service

*A:DutA# show service id 1 bgp-evpn
===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
CFM MAC Advertise  : Enabled
VXLAN Admin Status : Disabled           Creation Origin    : manual
MAC Dup Detn Moves : 3                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 9                  Number of Dup MACs : 0
IP Route Advertise*: Disabled

EVI                : 1
```

```
-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses            Time Detected
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.

===============================================================================
BGP EVPN MPLS Information
===============================================================================
Admin Status       : Enabled
Force Vlan Fwding  : Disabled           Control Word       : Disabled
Split Horizon Group: (Not Specified)
Ingress Rep BUM Lbl: Disabled           Max Ecmp Routes    : 4
Ingress Ucast Lbl  : 262142             Ingress Mcast Lbl  : 262142
Entropy Label      : Disabled
===============================================================================

===============================================================================
BGP EVPN MPLS Auto Bind Tunnel Information
===============================================================================
Resolution         : any
Filter Tunnel Types: (Not Specified)
===============================================================================
```

isid-route-target

Syntax **isid-route-target**

Context show>service>id>bgp-evpn

Description This command displays a list of the auto-derived or configured ISID-based route-targets per B-VPLS service. The entries show the ISID ranges and association to either an auto-rt or an actual configured route-target.

The auto-rt display format is: <2-byte-as-number>:<4-byte-value>, where: 4-byte-value = 0x30+ISID.

Output

Sample Output

```
*A:PE-2# show service id 10 bgp-evpn isid-route-target
===============================================================================
EVPN ISID RT Information
===============================================================================
Start     End       RT type    Route Target               Last Chgd
Range     Range
-------------------------------------------------------------------------------
11        11        auto       N/A                        10/03/2016 16:19:46
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

evpn-mpls

Syntax **evpn-mpls** [**esi** *esi*] [**es-bmac** *ieee-address*]

Context show>service>id

Description This command displays the existing EVPN-MPLS destinations for a specified service and all related information. The command allows filtering based on **esi** (for EVPN multi-homing) and **es-bmac** (for PBB-EVPN multi-homing) to display the EVPN-MPLS destinations associated to an ESI.

Parameters *esi* — Specifies an ESI by which to filter the displayed information.

ieee-address — Specifies a 48-bit MAC address by which to filter information. The parameter is entered in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

Output

Sample Output

```
*A:PE1# show service id 1 evpn-mpls
===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                 Transport
-------------------------------------------------------------------------------
192.0.2.69      262140        0           Yes             07/15/2015 19:44:07
                ldp
192.0.2.69      262141        2           No              07/15/2015 19:44:07
                ldp
192.0.2.70      262139        0           Yes             07/15/2015 19:44:07
                ldp
192.0.2.70      262140        1           No              07/15/2015 19:44:07
                ldp
192.0.2.72      262140        0           Yes             07/15/2015 19:44:07
                ldp
192.0.2.72      262141        1           No              07/15/2015 19:44:07
                ldp
192.0.2.73      262139        0           Yes             07/15/2015 19:44:09
                ldp
192.0.2.254     262142        1           Yes             07/15/2015 19:44:31
                bgp
-------------------------------------------------------------------------------
Number of entries : 8
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:71:00:00:00:01   2                       07/15/2015 20:41:09
01:74:13:00:74:13:00:00:74:13   1                       07/15/2015 20:41:07
-------------------------------------------------------------------------------
```

```
Number of entries: 2
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS ES BMAC Dest
===============================================================================
vBmacAddr                       Num. Macs               Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================


*A:PE1# show service id 1 evpn-mpls esi 01:00:00:00:00:71:00:00:00:01
===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:71:00:00:00:01   2                       07/15/2015 20:41:09
===============================================================================
===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport
-------------------------------------------------------------------------------
192.0.2.69              262141                  07/15/2015 20:41:09
                        ldp
192.0.2.72              262141                  07/15/2015 20:41:09
                        ldp
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================


A:PE3# show service id 20000 evpn-mpls es-bmac 00:00:00:00:71:71
===============================================================================
BGP EVPN-MPLS ES BMAC Dest
===============================================================================
vBmacAddr                       Num. Macs               Last Change
-------------------------------------------------------------------------------
00:00:00:00:71:71               1                       07/15/2015 19:44:10
===============================================================================
===============================================================================
BGP EVPN-MPLS ES BMAC Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport
-------------------------------------------------------------------------------
192.0.2.69              262138                  07/15/2015 19:44:10
                        ldp
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

esi

Syntax **esi** *esi*

Context show>service>id>evpn-mpls

Description This command shows the remote Ethernet segment identifiers (ESIs) as well as the BGP-EVPN MPLS destinations associated to them.

Parameters *esi* — Specifies a 10-byte ESI.

Output

Sample Output

```
*A:PE71(1)# show service id 1 evpn-mpls esi 01:00:00:00:00:71:00:00:00:01
===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:71:00:00:00:01   1                       07/17/2015 18:31:27
===============================================================================
===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport
-------------------------------------------------------------------------------
192.0.2.69              262141                  07/17/2015 18:31:26
                        ldp
192.0.2.72              262141                  07/17/2015 18:31:26
                        ldp
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

es-bmac

Syntax **es-bmac** *ieee-address*

Context show>service>id>evpn-mpls

Description This command shows the remote Ethernet segment BMACs as well as the BGP-EVPN MPLS destinations associated to them.

Parameters *ieee-address* — Specifies a 48-bit MAC address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

Output

Sample Output

```
*A:PE70(4)# show service id 20000 evpn-mpls es-bmac 00:00:00:00:71:71
===============================================================================
BGP EVPN-MPLS ES BMAC Dest
===============================================================================
vBmacAddr                       Num. Macs               Last Change
-------------------------------------------------------------------------------
00:00:00:00:71:71               1                       07/15/2015 19:50:22
===============================================================================
===============================================================================
BGP EVPN-MPLS ES BMAC Dest TEP Info
===============================================================================
TEP Address               Egr Label               Last Change
                          Transport
-------------------------------------------------------------------------------
192.0.2.69                262138                  07/15/2015 19:50:22
                          ldp
192.0.2.72                262136                  07/15/2015 19:50:22
                          ldp
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

es-pbr

Syntax **es-pbr**

Context show>service>id

Description When a filter with an **action forward esi** is successfully added to a VPLS service and the PE receives an EVPN Auto-Discovery route for the configured ESI, this command displays the PBR VXLAN bindings auto-created, including the ESI, the VXLAN VTEP:VNI and the status of the binding.

Output

Sample Output

```
A:PE1# show service id 301 es-pbr
===============================================================================
L2 ES PBR
===============================================================================
ESI                           Users      Status
                                         VTEP:VNI
-------------------------------------------------------------------------------
ff:00:00:00:00:00:00:00:00:01 1          Active
                                         192.0.2.72:7272
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

proxy-arp

Syntax **proxy-arp** [*ip-address*] [**detail**]
proxy-arp [*ip-address*] **dynamic**

Context show>service>id

Description This command displays, in a table, the existing proxy-ARP entries for a particular service. The table is populated by EVPN MAC routes that contain a MAC and an IP address, as well as static entries or dynamic entries from snooped ARP messages on access SAP or SDP-bindings.

A 7750 SR, 7450 ESS, or 7950 XRS that receives an ARP request from a SAP or SDP-binding performs a lookup in the proxy-ARP table for the service. If a match is found, the router replies to the ARP and does not allow ARP flooding in the VPLS service. If a match is not found, the ARP is flooded within the service if the configuration allows it.

The command allows for specific IP addresses to be displayed. Dynamic IP entries associated to a MAC list are displayed with the corresponding MAC list and resolve timers information.

Parameters *ip-address* — Specifies an IP address.

Values a.b.c.d

detail — Displays detailed information.

dynamic — Displays detailed information about dynamic entries.

Output

Sample Output

```
*A:PE-3# show service id 5 proxy-arp
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : 120 secs
Table Size        : 250                 Total             : 1
Static Count      : 0                   EVPN Count        : 0
Dynamic Count     : 1                   Duplicate Count   : 0
Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None
EVPN
-------------------------------------------------------------------------------
Garp Flood        : enabled             Req Flood         : enabled
Static Black Hole : disabled
EVPN Route Tag    : 10
-------------------------------------------------------------------------------
*A:PE-3# show service id 5 proxy-arp detail
```

```
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : 120 secs
Table Size        : 250                 Total             : 1
Static Count      : 0                   EVPN Count        : 0
Dynamic Count     : 1                   Duplicate Count   : 0
Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None
EVPN
-------------------------------------------------------------------------------
Garp Flood        : enabled             Req Flood         : enabled
Static Black Hole : disabled
EVPN Route Tag    : 10
-------------------------------------------------------------------------------
===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
10.0.0.1            00:ca:fe:ca:fe:01   dyn       active    12/20/2016 12:38:28
-------------------------------------------------------------------------------
Number of entries : 1


===============================================================================
*A:PE-3# show service id 5 proxy-arp dynamic
===============================================================================
Proxy ARP Dyn Cfg Summary
===============================================================================
IP Addr                                 Mac List
-------------------------------------------------------------------------------
10.0.0.1                                list-1
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================


*A:PE-3# show service id 5 proxy-arp dynamic 10.0.0.1
===============================================================================
Proxy ARP Dyn Cfg Detail
===============================================================================
IP Addr        Mac List                           Resolve Time   Remaining
                                                  (mins)          Resolve Time
                                                                  (secs)
-------------------------------------------------------------------------------
10.0.0.1       list-1                             1              0
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

proxy-nd

Syntax **proxy-nd** [*ipv6-address*] [**detail**]
proxy-nd [*ipv6-address*] **dynamic**

Context show>service>id

Description This command displays, in a table, the existing proxy-ND entries for a particular service. The table is populated by the EVPN MAC routes containing a MAC and an IPv6 address, as well as static entries or dynamic entries from snooped NA messages on access SAP or SDP-bindings.

A 7750 SR, 7450 ESS, or 7950 XRS that receives a Neighbor Solicitation (NS) from a SAP or SDP-binding performs a lookup in the proxy-ND table for the service. If a match is found, the router replies to the NS and does not allow NS flooding in the VPLS service. If a match is not found, the NS is flooded in the service if the configuration allows it.

The command allows for specific IPv6 addresses to be shown. Dynamic IPv6 entries associated to a MAC list are shown with the corresponding MAC list and resolve timers information.

Parameters *ipv6-address* — Specifies an IPv6 address.

Values ipv6-address:
x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
where:
x - [0 to FFFF]H
d - [0 to 255]D

detail — Displays detailed information.

dynamic — Displays detailed information about dynamic entries.

Output

Sample Output

```
*A:PE-2#  show service id 5 proxy-nd
-------------------------------------------------------------------------------
Proxy ND
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : 120 secs
Table Size        : 250                 Total             : 1
Static Count      : 0                   EVPN Count        : 0
Dynamic Count     : 1                   Duplicate Count   : 0
Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None
```

```
EVPN
------------------------------------------------------------------------------
Unknown NS Flood  : enabled             ND Advertise      : Router
Rtr Unsol NA Flood: enabled             Host Unsol NA Fld : enabled
EVPN Route Tag    : 10
------------------------------------------------------------------------------


*A:PE-2#  show service id 5 proxy-nd detail
------------------------------------------------------------------------------
Proxy ND
------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : 120 secs
Table Size        : 250                 Total             : 1
Static Count      : 0                   EVPN Count        : 0
Dynamic Count     : 1                   Duplicate Count   : 0
Dup Detect
------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None
EVPN
------------------------------------------------------------------------------
Unknown NS Flood  : enabled             ND Advertise      : Router
Rtr Unsol NA Flood: enabled             Host Unsol NA Fld : enabled
EVPN Route Tag    : 10
------------------------------------------------------------------------------
==============================================================================
VPLS Proxy ND Entries
==============================================================================
IP Address               Mac Address       Type Status Rtr/ Last Update
                                                      Host
------------------------------------------------------------------------------
200::1                   00:ca:fe:ca:fe:01 dyn  active Rtr  12/20/2016 14:04:10
------------------------------------------------------------------------------
Number of entries : 1
==============================================================================


*A:PE-2#  show service id 5 proxy-nd dynamic
==============================================================================
Proxy ND Dyn Cfg Summary
==============================================================================
IP Addr                                 Mac List
------------------------------------------------------------------------------
200::1                                  list-1
------------------------------------------------------------------------------
Number of Entries: 1
==============================================================================


*A:PE-2#  show service id 5 proxy-nd dynamic 200::1
==============================================================================
Proxy ND Dyn Cfg Detail
==============================================================================
IP Addr                                 Mac List
 Resolve Time(mins)                      Remaining Resolve Time(secs)
------------------------------------------------------------------------------
200::1                                  list-1
```

```
1                                       0
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

vxlan

Syntax **vxlan**
vxlan assisted-replication replicator
vxlan instance *instance* **oper-flags**

Context show>service>id

Description This command displays the VXLAN bindings auto-created in a specified service. A VXLAN binding is composed of the remote VTEP (VXLAN Termination Endpoint) and the corresponding egress VNI (VXLAN Network Identifier) to identify the service at the egress node. The command shows the number of MACs associated to each binding as well as the operational status and if the binding is part of the multicast list. The binding will be operationally down when the VTEP address is not found in the base routing table (the VTEP address cannot be reached). A binding will be part of the multicast list if a valid BGP EVPN inclusive multicast route exists for it.

A VXLAN binding can be associated with the following types of Mcast values.

- BM — Refers to the capability of the binding to send Broadcast or Multicast to the remote VTEP. This binding type is setup to AR Replicator nodes from a Leaf node.
- BUM — Refers to the capability of the binding to send Broadcast or Multicast to the remote VTEP. This binding type is setup to AR Replicator nodes from a Leaf node.
- U — Refers to the capability of the binding to send Unknown Unicast to the VTEP. This binding type is setup from Leaf nodes to other Leaf and RNVE nodes.
- "-" — Specifies that the binding can only be used for known unicast traffic.

Parameters **assisted-replication replicator** — Displays all the discovered candidate AR Replicators for the service and the replicator that has been selected by the leaf to send the BM traffic. The list of replicators is ordered by VTEP address and VNI. This command is only supported on the nodes configured as leaf.

The "In Use" column indicates whether the replicator has been selected for the service. When selecting a replicator for the service, the candidate list is ordered by VTEP IP (lowest IP is ordinal 0) and VNI. A modulo function of the service ID and the number of candidate PEs will give the selected replicator for a specified service.

The "Pending Time" column shows the remaining seconds till the node starts sending the BM traffic to the replicator. This time is configurable by the replicator-activation-time parameter.

instance — Specifies the VXLAN instance.

Values 1

Output

Sample Output

```
A:PE6# show service id 8001 vxlan
===============================================================================
Vxlan Src Vtep IP: N/A
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 801
Creation Origin: manual
Assisted-Replication: none
RestProtSrcMacAct: none
===============================================================================
VPLS VXLAN service Network Specifics
===============================================================================
Ing Net QoS Policy : none                          Vxlan VNI Id     : 801
Ingress FP QGrp    : (none)                        Ing FP QGrp Inst : (none)
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI  Num. MACs   Mcast Oper  L2
                                                                    State PBR
-------------------------------------------------------------------------------
192.0.2.4                             801         2           BUM   Up    No
192.0.2.5                             801         1           BUM   Up    No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
===============================================================================
A:PE6# show service id 7000 vxlan
===============================================================================
Vxlan Src Vtep IP: N/A
===============================================================================
VPLS VXLAN, Ingress VXLAN Network Id: 7000
Creation Origin: manual
Assisted-Replication: leaf      Replicator-Activation-Time: None
RestProtSrcMacAct: none
===============================================================================
VPLS VXLAN service Network Specifics
===============================================================================
Ing Net QoS Policy : none                          Vxlan VNI Id     : 7000
Ingress FP QGrp    : (none)                        Ing FP QGrp Inst : (none)
===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI  Num. MACs   Mcast Oper  L2
                                                                    State PBR
-------------------------------------------------------------------------------
4.4.4.4                               7000        0           BM    Up    No
5.5.5.5                               7000        0           -     Up    No
192.0.2.4                             7000        0           U     Up    No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 3
-------------------------------------------------------------------------------

A:PE6# show service id 7000 vxlan assisted-replication replicator
===============================================================================
Vxlan AR Replicator Candidates
===============================================================================
VTEP Address           Egress VNI     In Use  In Candidate List Pending Time
-------------------------------------------------------------------------------
```

```
4.4.4.4                7000           yes     yes              0
5.5.5.5                7000           no      yes              0
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================

A:PE-2# show service id 500 vxlan instance 1 oper-flags

===============================================================================
VPLS VXLAN oper flags
===============================================================================
MhStandby                              : false
===============================================================================
```

evpn-mpls

Syntax **evpn-mpls** [*tep-ip-address*]

Context show>service

Description This command shows the remote EVPN-MPLS tunnel endpoints in the system.

Parameters *tep-ip-address* — Specifies the IP address of a tunnel endpoint.

Values a.b.c.d

Output

Sample Output

```
*A:PE70(4)# show service evpn-mpls
===============================================================================
EVPN MPLS Tunnel Endpoints
===============================================================================
EvpnMplsTEP Address EVPN-MPLS Dest      ES Dest             ES BMac Dest
-------------------------------------------------------------------------------
192.0.2.69          3                   1                   1
192.0.2.71          2                   0                   0
192.0.2.72          3                   1                   1
192.0.2.73          2                   1                   0
192.0.2.254         1                   0                   0
-------------------------------------------------------------------------------
Number of EvpnMpls Tunnel Endpoints: 5
-------------------------------------------------------------------------------
===============================================================================
*A:PE70(4)# show service evpn-mpls
<TEP ip-address>
 192.0.2.69   192.0.2.71   192.0.2.72   192.0.2.73   192.0.2.254


*A:PE70(4)# show service evpn-mpls 192.0.2.69
===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
```

```
Service Id                              Egr Label
------------------------------------------------------------------------------
1                                       262140
1                                       262141
20000                                   262138
------------------------------------------------------------------------------
==============================================================================

==============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
==============================================================================
Service Id          Eth Seg Id                              Egr Label
------------------------------------------------------------------------------
1                   01:00:00:00:00:71:00:00:00:01           262141
------------------------------------------------------------------------------
==============================================================================
==============================================================================
BGP EVPN-MPLS ES BMac Dest
==============================================================================
Service Id              ES BMac                 Egr Label
------------------------------------------------------------------------------
20000                   00:00:00:00:71:71       262138
------------------------------------------------------------------------------
==============================================================================
```

vxlan

Syntax **vxlan** [*ip-address*]

Context show>service

Description This command displays the VXLAN bindings auto-created in a specified service. A VXLAN binding is composed of the remote VTEP (VXLAN Termination Endpoint) and the corresponding egress VNI (VXLAN Network Identifier) to identify the service at the egress node. The command shows the number of MACs associated to each binding as well as the operational status and if the binding is part of the multicast list. The binding will be operationally down when the VTEP address is not found in the base routing table (the VTEP address cannot be reached). A binding will be part of the multicast list if a valid BGP EVPN inclusive multicast route exists for it.

A VXLAN binding can be associated with the following types of Mcast values.

- BM — Refers to the capability of the binding to send Broadcast or Multicast to the remote VTEP. This binding type is setup to AR Replicator nodes from a Leaf node.
- BUM — Refers to the capability of the binding to send Broadcast or Multicast to the remote VTEP. This binding type is setup to AR Replicator nodes from a Leaf node.
- U — Refers to the capability of the binding to send Unknown Unicast to the VTEP. This binding type is setup from Leaf nodes to other Leaf and RNVE nodes.
- "-" — Specifies that the binding can only be used for known unicast traffic.

Parameters *ip-address* — Specifies the remote VTEP address for the VXLAN binding.

Values ipv4-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
where:
x: [0 to FFFF]H
d: [0 to 255]D

Output

Sample Output

```
===============================================================================
A:PE6# show service vxlan
===============================================================================
VXLAN Tunnel Endpoints (VTEPs)
===============================================================================
VTEP Address                                Number of Egress VNIs    Oper
                                                                     State
-------------------------------------------------------------------------------
2.2.2.2                                     1                        Up
4.4.4.4                                     2                        Up
5.5.5.5                                     1                        Up
192.0.2.2                                   1                        Up
192.0.2.3                                   1                        Up
192.0.2.4                                   2                        Up
192.0.2.5                                   2                        Up
-------------------------------------------------------------------------------
Number of VTEPs: 7
-------------------------------------------------------------------------------
===============================================================================
A:PE6# show service vxlan 2.2.2.2
===============================================================================
VXLAN Tunnel Endpoint: 2.2.2.2
===============================================================================
Egress VNI                  Service Id              Oper State
-------------------------------------------------------------------------------
4000                        4000                    Up
-------------------------------------------------------------------------------
===============================================================================
```

vxlan-instance-using

Syntax **vxlan-instance-using ethernet-segment** [*name*]

Context show>service

Description This command displays the services and VXLAN instances associated with a specified virtual ES, as well as its operational status.

Parameters *name* — Specifies the virtual ES name, up to 28 characters.

Output

Sample Output

```
A:PE-2# show service vxlan-instance-using ethernet-segment "vES23"
===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
SvcId                           VXLAN Instance                  Status
-------------------------------------------------------------------------------
500                             1                               DF
===============================================================================
A:PE-2# show service vxlan-instance-using ethernet-segment
===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
SvcId       VXLAN Instance      ES Name                         Status
-------------------------------------------------------------------------------
500         1                   vES23                           DF
===============================================================================
```

server

Syntax **server** [*name*]

Context show>system>xmpp

Description This command shows the connectivity to the XMPP server, including the configured parameters and statistics. When the user provides the name of the server, a detailed view is shown.

Parameters *name* — Specifies the name of the XMPP server. 32 characters maximum.

Output

Sample Output

```
:sr12U-46-PE2# show system xmpp server
===============================================================================
XMPP Server Table
===============================================================================
Name                            User Name          State
 XMPP FQDN                       Last State chgd    Admin State
-------------------------------------------------------------------------------
vsd1-hy                         cspTest             Functional
 vsd1-hy.alu-srpm.us             0d 22:42:15        inService
-------------------------------------------------------------------------------
No. of XMPP server's: 1
===============================================================================
B:Dut# show system xmpp server "vsd1-hy"
===============================================================================
XMPP Server Table
===============================================================================
XMPP FQDN          : vsd1-hy.alu-srpm.us
```

```
XMPP Admin User    : cspTest
XMPP Oper User     : cspTest
State Lst Chg Since: 0d 22:40:16        State              : Functional
Admin State        : Up                 Connection Mode    : outOfBand
Auth Type          : md5
IQ Tx.             : 306                IQ Rx.             : 306
IQ Error           : 72                 IQ Timed Out       : 0
IQ Min. Rtt        : 100 ms             IQ Max. Rtt        : 450 ms
IQ Ack Rcvd.       : 234
Push Updates Rcvd  : 41                 VSD list Upd Rcvd  : 91
Msg Tx.            : 279                Msg Rx.            : 207
Msg Ack. Rx.       : 135                Msg Error          : 72
Msg Min. Rtt       : 0 ms               Msg Max. Rtt       : 450 ms
Sub Tx.            : 1                  UnSub Tx.          : 0
Msg Timed Out      : 0

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

vsd

Syntax **vsd** [*entry*]

Context show>system
show>system>xmpp

Description This command shows the connectivity to the VSD server, including the configured parameters and statistics. When the user provides the entry number of the VSD server, a detailed view for that specific server is shown, including statistics.

Parameters *entry* — Specifies the entry number of the VSD server.

Values 0 to 4294967295

Output

Sample Output

```
:Dut# show system vsd
===============================================================================
VSD Information
===============================================================================
System Id               : SR12U-46-PE
GW Last Audit Tx Time   : 03/07/2000 04:07:06

Gateway Publish-Subscribe Information
-------------------------------------------------------------------------------
Subscribed              : True
Subscriber Name         : nuage_gateway_id_SR12U-46-PE
Last Subscription Time  : 03/06/2000 05:27:06
===============================================================================


*B:Dut# show system xmpp vsd
===============================================================================
Virtual Services Directory Table
```

```
===============================================================================
Id User Name                    Uptime             Status
-------------------------------------------------------------------------------
1  cna@vsd1-hy.alu-srpm.us/nua* 0d 22:45:39        Available
-------------------------------------------------------------------------------
No. of VSD's: 1
===============================================================================


*B:Dut# show system xmpp vsd 1
===============================================================================
VSD Server Table
===============================================================================
VSD User Name      : cna@vsd1-hy.alu-srpm.us/nuage
Uptime             : 0d 22:45:41        Status             : Available
Msg Tx.            : 282                Msg Rx.            : 209
Msg Ack. Rx.       : 136                Msg Error          : 73
Msg TimedOut       : 0                  Msg MinRtt         : 70 ms
Msg MaxRtt         : 450 ms

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

domain

Syntax **domain** [*domain-name*] [**association**]

Context show>system>vsd

Description This command shows the different VSD domains configured in the system. If association is added, the VSD domain to service association is shown. If a specific domain-name is used, configuration event statistics are shown.

Parameters *domain-name* — Specifies a VSD domain for which to display information.

association — Displays all VSD domain-to-service associations.

Output

Sample Output

```
B:Dut# show service vsd domain
===============================================================================
VSD Domain Table
===============================================================================
Name                                Type           Origin    Admin
-------------------------------------------------------------------------------
nuage_401                           l2DomainIrb    manual    inService
nuage_402                           l2Domain       manual    inService
nuage_501                           l2Domain       manual    inService
nuage_502                           l2Domain       manual    inService
-------------------------------------------------------------------------------
Number of entries: 4
===============================================================================
*B:Dut# show service vsd domain "nuage_501"
===============================================================================
```

```
VSD Information
===============================================================================
Name               : nuage_501
Description        : nuage_501_l2_domain
Type               : l2Domain                    Admin State   : inService
Last Error To Vsd  : (Not Specified)
Last Error From Vsd: (Not Specified)

Statistics
-------------------------------------------------------------------------------
Last Cfg Chg Evt   : 01/01/2000 00:00:11         Cfg Chg Evts  : 0
Last Cfg Update    : 01/01/2000 00:00:11         Cfg Upd Rcvd  : 0
Last Cfg Done      : 01/01/2000 00:00:11
Cfg Success        : 0                           Cfg Failed    : 0
===============================================================================
*B:Dut# show service vsd domain "nuage_501" association
============================================================
Service VSD Domain
============================================================
Svc Id      Svc Type  Domain Type    Domain Admin   Origin
------------------------------------------------------------
501         vpls      l2Domain       inService      manual
------------------------------------------------------------
Number of entries: 1
============================================================
*B:sr12U-46-PE2# show service vsd domain association
==========================================
Services-using VSD Domain
==========================================
Svc Id     Domain
------------------------------------------
501        nuage_501
502        nuage_502
------------------------------------------
Number of services using VSD Domain: 2
==========================================
```

vxlan

Syntax **vxlan**

Context show>service>system

Description This command shows the global VXLAN configuration in the system. In particular, the command displays the configured assisted-replication IP address and the VXLAN tunnel-termination addresses, if the system terminates VXLAN tunnels in addresses that are not the same as the system IP address.

Output

Sample Output

```
A:PE1# show service system vxlan
===============================================================================
System VXLAN Information
```

```
===============================================================================
Asstd Repl Ip Address.                :
===============================================================================
===============================================================================
Vxlan Tunnel Termination
===============================================================================
Tunnel Term IP                                FPE ID    Last Change
-------------------------------------------------------------------------------
11.11.11.1                                    1         06/22/2016 14:18:55
-------------------------------------------------------------------------------
Number of Entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

redundancy

Syntax **redundancy**

Context show

Description This command enables the context for the display of global redundancy parameters.

bgp-evpn-multi-homing

Syntax **bgp-evpn-multi-homing**

Context show>redundancy

Description This command shows the information related to the EVPN global timers.

Output

Sample Output

```
*A:PE2# show redundancy bgp-evpn-multi-homing
===============================================================================
Redundancy BGP EVPN Multi-homing Information
===============================================================================
Boot-Timer              : 10 secs
Boot-Timer Remaining    : 0 secs
ES Activation Timer     : 3 secs
===============================================================================
```

5.6.2.3 Clear Commands

mac-dup-detect

Syntax **mac-dup-detect** {*ieee-address* | **all**}

Context clear>service>id>evpn

Description This command clears a duplicate MAC and restarts the MAC duplication process. It also clears black-hole MACs.

Parameters *ieee-address* — Specifies the MAC address.

all — Specifies that the **clear** command applies to all duplicate MACs.

domain

Syntax **domain** [*name*]

Context clear>service>statistics>vsd

Description This command clears the statistics shown in the **show service vsd domain** *name* command.

Parameters *name* — Specifies the VSD domain name.

scripts

Syntax **scripts**

Context clear>service>statistics>vsd

Description This command clears the statistics shown in the **show service vsd script statistics** command.

server

Syntax **server** [*xmpp-server-name*]

Context clear>system>statistics>xmpp

Description This command clears the statistics shown in the **show system xmpp server** *name* command.

Parameters *xmpp-server-name* — Specifies the XMPP server name.

ver

Syntax **server** [*xmpp-server-name*]

Context clear>system>statistics>xmpp

Description This command clears the statistics shown in the **show system xmpp server** *name* command.

Parameters *xmpp-server-name* — Specifies the XMPP server name.

5.6.2.4 Debug Commands

xmpp

Syntax **xmpp** [**connection**] [**gateway**] [**message**] [**vsd**] [**iq**] [**all**]
no xmpp

Context debug>system

Description This command enables the debug for XMPP messages sent or received by the 7750 SR, 7450 ESS, or 7950 XRS.

Parameters **connection** — Filters only the messages related to the XMPP connection.

gateway — Filters the messages related to the gateway.

message — Filters only the messages.

vsd — Filters the VSD messages.

iq — Filters the IQ messages between the gateway and the VSD.

all — Includes all the above.

vsd

Syntax **vsd**

Context debug

Description This command enables the context for the debug vsd commands.

scripts

Syntax **scripts**

scripts event [**cli**] [**errors**] [**executed-cmd**] [**state-change**] [**warnings**]
scripts instance *instance* **event** [**cli**] [**errors**] [**executed-cmd**] [**state-change**] [**warnings**]

Context debug>vsd

Description This command enables the debug of the VSD fully dynamic integration scripts.

event

Syntax [**no**] **event**

Context debug>vsd>scripts

Description This command enables/disables the generation of all script debugging event output: cli, errors, execute-cmd, warnings, state-change.

instance

Syntax [**no**] **instance** *instance*

Context debug>vsd>scripts

Description This command enables/disables the generation of script debugging for a specific instance

Parameters *instance* — Specifies the instance name.

cli

Syntax [**no**] **cli**

Context debug>vsd>scripts>event
debug>vsd>scripts>instance

Description This command enables/disables the generation of a specific script debugging event output: **cli**.

errors

Syntax [**no**] **errors**

Context debug>vsd>scripts>event
debug>vsd>scripts>instance

Description This command enables/disables the generation of a specific script debugging event output: **errors**.

executed-cmd

Syntax [**no**] **executed-cmd**

Context debug>vsd>scripts>event
debug>vsd>scripts>instance

Description This command enables/disables the generation of a specific script debugging event output: **execute-cmd**.

state-change

Syntax [**no**] **state-change**

Context debug>vsd>scripts>event
debug>vsd>scripts>instance

Description This command enables/disables the generation of a specific script debugging event output: **state-change**.

warnings

Syntax [**no**] **warnings**

Context debug>vsd>scripts>event
debug>vsd>scripts>instance

Description This command enables/disables the generation of a specific script debugging event output: **warnings**.

5.6.2.5 Tools Commands

service

Syntax **service**

Context tools>dump

Description Use this command to configure tools to display service dump information.

id

Syntax **id** *service-id*

Context tools>dump

Description Use this command to configure parameters to display service ID information.

Parameters *service-id* — Specifies the service ID.

vxlan

Syntax **vxlan** [**clear**]

Context tools>dump>service>id

Description This command displays the number of times a service could not add a VXLAN binding or <VTEP, Egress VNI> due to the following limits:

- The per-system VTEP limit has been reached
- The per-system <VTEP, Egress VNI> limit has been reached
- The per-service <VTEP, Egress VNI> limit has been reached
- The per-system Bind limit: Total bind limit or vxlan bind limit has been reached.

Parameters **clear** — Clears the per-system VTEP, per-system VTEP Egress VNI, per-service VTEP Egress VNI, and per-system Bind statistics.

Output

Sample Output

```
*A:PE63# tools dump service id 3 vxlan
VTEP, Egress VNI Failure statistics at 000 00:03:55.710:
statistics last cleared at 000 00:00:00.000:
      Statistic            |    Count
---------------------------+-------------
                      VTEP |           0
             Service Limit |           0
              System Limit |           0
   Egress Mcast List Limit |           0
Duplicate VTEP, Egress VNI |           1
```

dup-vtep-egrvni

Syntax **dup-vtep-egrvni** [**clear**]

Context tools>dump>service>vxlan

Description This command dumps the <VTEP, VNI> bindings that have been detected as duplicate attempts, i.e. an attempt to add the same binding to more than one service. The commands provides a **clear** option.

Parameters **clear** — Clears the VTEP VNI bindings that have been detected as duplicate attempts.

Output

Sample Output

```
*A:PE71# tools dump service vxlan dup-vtep-egrvni
Duplicate VTEP, Egress VNI usage attempts at 000 00:03:41.570:
1. 10.1.1.1:100
```

usage

Syntax **usage**

Context tools>dump>service>id>evpn

Description This command shows the maximum number of EVPN-tunnel interface IP next-hops per R-VPLS as well as the current usage for a specified R-VPLS service.

Output

Sample Output

```
*A:PE71# tools dump service id 504 evpn usage
Evpn Tunnel Interface IP Next Hop: 1/8189
```

domain-to-vsd-mapping

Syntax **domain-to-vsd-mapping**

Context tools>dump>service

Description This command enables the context for the domain-to-vsd mappings.

domain

Syntax **domain name** *name*

Context tools>dump>service>domain-to-vsd-mapping

Description This command shows mapping of a specified VSD to a vsd-domain.

Parameters *name* — Specifies a VSD domain name.

Output

Sample Output

```
Dut# tools dump service domain-to-vsd-mapping domain name "nuage_501"
```

```
===============================================================================
Domain to VSD Mapping
===============================================================================
Domain name                     VSD
-------------------------------------------------------------------------------
nuage_501                       cna@vsd1-hy.alu-srpm.us/nuage
===============================================================
```

xmpp

Syntax **xmpp**

Context tools>perform>system

Description This command enables the xmpp context.

vsd-refresh

Syntax **vsd-refresh**

Context tools>perform>system>xmpp

Description This command instructs the system to refresh immediately the list of VSDs and not to wait for the next VSD list audit that the system does periodically.

proxy-arp

Syntax **proxy-arp**

Context tools>perform>service>id

Description This command enables the proxy-arp context.

dynamic-resolve

Syntax **dynamic-resolve all** [**force**]
dynamic-resolve *ip-address* [**force**]

Context tools>perform>service>id>proxy-arp

Description This command triggers the resolve procedure for dynamic IP entries. When executed, a resolve message (ARP-request) is issued for the requested IP or, if the **all** option used, for all the configured dynamic IPs.

The **force** option triggers the resolve process even for IPs with an existing entry in the proxy-ARP table.

Parameters *ip-address* — Specifies the IP address.

Values a.b.c.d

all — Runs the command for all configured dynamic IPs.

force — Issues a resolve message even when configured dynamic IP entries are present.

proxy-nd

Syntax **proxy-nd**

Context tools>perform>service>id

Description This command enables the proxy-nd context.

dynamic-resolve

Syntax **dynamic-resolve all** [**force**]
dynamic-resolve *ipv6-address* [**force**]

Context tools>perform>service>id>proxy-nd

Description This command triggers the resolve procedure for dynamic IPv6 entries. When executed, a resolve message (Neighbor Solicitation) is issued for the requested IPv6 or, if the **all** option used, for all the configured dynamic IPv6s.

The **force** option triggers the resolve process even for IPv6 addresses with an existing entry in the proxy-ARP table.

Parameters *ipv6-address* — Specifies the IPv4 or IPv6 address.

Values ip-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
where:
x: [0 to FFFF]H
d: [0 to 255]D

all — Runs the command for all configured dynamic IPv6 addresses.

force — Issues a resolve message even when configured dynamic IP entries are present.

domain

Syntax **domain name** [*name*] **refresh-config**

Context tools>perform>service>vsd

Description This command instructs the system to refresh the configuration of a specified domain immediately instead of waiting for the next audit interval.

Parameters *name* — Specifies the name of the VSD domain.

evaluate-script

Syntax **evaluate-script domain-name** [*domain-name*] **type** [*type*] **action** *script-action* [**vni** *vni-id*] [**rt-i** *ext-community*] [**rt-e** *ext-community*] [**metadata** *metadata*] **policy** *python-policy*

Context tools>perform>service>vsd

Description The command enables the user to test their setup, and modify and tear down Python scripts in a lab environment without the need to be connected to a VSD. The successful execution of the command for action setup will create a VSD domain and the corresponding configuration, just as the system would do when the parameters are received from VSD.

Parameters *domain-name* — Specifies the VSD domain name. 64 characters maximum.

type — Specifies the VSD domain type.

Values l2-domain, vrf-gre, vrf-vxlan, l2-domain-irb

script-action — Specifies the action taken for Python scripts.

Values setup, modify, teardown

vni-id — Specifies the VNI ID.

Values 1 to 16777215

rt-i *ext-community* — Specifies the internal route-target (RT-i).

Values target: {*ip-addr*:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4byte-asnumber*:*comm-val*}
ip-addr: a.b.c.d
comm-val: 0 to 65535
2byte-asnumber: 0 to 65535
ext-comm-val: 0 to 4294967295
4byte-asnumber: 0 to 4294967295

rt-e *ext-community* — Specifies the external route-target (RT-e).

Values target: {*ip-addr*:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4byte-asnumber*:*comm-val*}
ip-addr: a.b.c.d

comm-val: 0 to 65535
2byte-asnumber: 0 to 65535
ext-comm-val: 0 to 4294967295
4byte-asnumber: 0 to 4294967295

metadata — Specifies the opaque *key=value* pairs. 500 characters maximum.

python-policy — Specifies the name of the Python script used to translate the VSD parameters into a configuration. 32 characters maximum.

Output

Sample Output

```
*A:PE1# tools perform service vsd evaluate-script domain-name "L2-DOMAIN-5" type l2-
domain action setup policy "py-l2" vni 64000 rt-i target:64000:64000 rt-
e target:64000:64000 metadata "rd=1:1, sap=1/1/10:3000"

Success
```

fd-domain-sync

Syntax **fd-domain-sync** {**full** | **diff**}

Context tools>perform>service>vsd

Description This command instructs the system to audit the VSD and retrieve either the "DIFF" list or the "FULL" list of domains in the VSD.

Parameters **full** — Retrieves the full VSD domain list.

diff — Retrieves the diff VSD domain list.

bgp-evpn

Syntax **bgp-evpn**

Context tools>dump>service>system

Description This command enables the context for the bgp-evpn base instance.

ethernet-segment

Syntax **ethernet-segment** *name* **evi** *evi* **df**
ethernet-segment *name* **isid** *isid* **df**

Context tools>dump>service>system>bgp-evpn

Description This command shows the computed DF PE for a specified EVI or ISID.

Parameters *name* — Specifies the name of the Ethernet segment. 32 characters maximum.

evi — Specifies the EVI.

Values 1 to 65535

isid — Specifies the ISID.

Values 1 to 16777215

Output

Sample Output

```
*A:PE2# tools dump service system bgp-evpn ethernet-segment "ESI-71" evi 1 df
[07/15/2015 21:52:08] Computed DF: 192.0.2.72 (Remote) (Boot Timer Expired: Yes)
*A:PE2# tools dump service system bgp-evpn ethernet-segment "ESI-71" isid 20001 df
[07/15/2015 21:52:21] Computed DF: 192.0.2.72 (Remote) (Boot Timer Expired: Yes)
```

evpn

Syntax **evpn**

Context tools>dump>service

Description This command enables the context for the global evpn parameters.

usage

Syntax **usage**

Context tools>dump>service>evpn

Description This command displays the consumed VXLAN EVPN resources in the system.

Output

Sample Output

```
*A:PE71# tools dump service evpn usage

EVPN usage statistics at 000 02:01:03.810:

MPLS-TEP                                        :             5
VXLAN-TEP                                       :             2
Total-TEP                                       :      7/  8191

Mpls Dests (TEP, Egress Label + ES + ES-BMAC)   :            16
Mpls Etree Leaf Dests                           :             1
Vxlan Dests (TEP, Egress VNI)                   :             2
```

```
Total-Dest                                  :     18/131071

Sdp Bind +  Evpn Dests                      :     20/196607
ES L2/L3 PBR                                :      0/ 32767
Evpn Etree Remote BUM Leaf Labels           :             3
```

vsd-services

Syntax **vsd-services**

Context tools>dump>service

Description This command enables the context for vsd-services commands.

command-list

Syntax **command-list**

Context tools>dump>service>vsd-services

Description This command displays the list of CLI nodes allowed in the VSD fully dynamic provisioning model. Python will have access to the shown nodes.

When access is granted to a node, all commands in that node are allowed; however, CLI nodes are only allowed if explicitly listed. Nodes in CLI are shown with a "+" in the CLI.

While you can navigate special "Pass through nodes" via these nodes, the commands in that node are not implicitly allowed. When configured in a service through VSD, these commands will not be shown in the 'info' output of the **config** command.

Note: A 'node' implies leaf-nodes and leaf-table nodes in reality. A 'Leaf-table' is a sub-table that looks like a leaf (i.e. it is entered/displayed as a one-liner). An example of leaf-table node is /**configure router policy-options prefix-list x prefix 0.0.0.0/0** - since you can have multiple instances of prefixes.

6 Pseudowire Ports

This chapter provides information about pseudowire ports (PW ports), process overview, and implementation notes.

6.1 Overview

A PW port is primarily used to provide PW termination with the following characteristics:

- Provide access (SAP) based capabilities to a PW which has traditionally been a network port based concept within SR OS. For example, PW payload can be extracted onto a PW-port-based SAPs with granular queuing capabilities (queuing per SAP). This is in contrast with traditional PW termination on network ports where queuing is instantiated per physical port on egress or per MDA on ingress.
- Lookup dot1q and qinq VLAN tags underneath the PW labels and map the traffic to different services.
- Terminate subscriber traffic carried within the PW on a BNG. In this case PW-port-based SAPs are instantiated under a group interface with Enhanced Subscriber Management (ESM). In this case, a PW-port-based SAP is treated as any other regular SAP created directly on a physical port with full ESM capabilities.

Mapping between PWs and PW ports is performed on one-to-one basis.

There are two modes in which PW port can operate:

- A PW port bound to a specific physical port (I/O port) — A successful mapping between the PW and PW port requires that the PW terminates on the same physical port (I/O port) to which the PW port is bound. In this mode of operation, PW ports do not support re-routing of PWs between the I/O ports. For example, if a PW is rerouted to an alternate physical port due to a network failure, the PW port will become non-operational.
- A PW port independent of the physical port (I/O port) on which the PW is terminated. This capability relies on FPE functionality and hence the name FPE based PW port. The benefit of such PW port is that it can provide services in cases where traffic within PW is rerouted between I/O ports due to a network failure.

When the PW port is created, the mapping between the PW port and PW will depend on the mode of operation and application.

PW port creation:

```
configure
    pw-port <id>
        encap-type {dot1q|qinq}
```

Similar to any other Ethernet-based port, the PW port supports two encapsulation types, dot1q and qinq.

Ether-type on a PW port is not configurable and it is set to a fixed value of 0x8100 for dot1q and qinq encapsulation.

6.2 PW Port Bound to a Physical Port

In this mode of operation, the PW port is bound to a specific physical port through an SDP binding context:

```
configure
    service
        sdp 1 mpls create
            far-end 10.10.10.10
            ldp
            binding
                port 1/1/1
                pw-port 1 vc-id 11 create
                    egress
                    shaping inter-dest-id vport-1
```

In this example, pw-port 1 is bound to a physical port 1/1/1.This PW port is mapped to the PW with vc-id 11 under the sdp 1 which must be terminated on port 1/1/1.A PW port is shaped by a virtual port scheduler (Vport) construct named vport-1 configured under port 1/1/1. SAPs created under such PW ports can be terminated in ESM, Layer 3 IES/VPRN interface or in an Epipe.

6.3 FPE-Based PW Port

The FPE based PW-port is primarily used to extract a PW payload onto an access based PW-port SAP, independent of the network I/O ports. FPE uses Port Cross-Connect (PXC) ports and provides an anchoring point for PW-port, independent of I/O ports, the term anchored PW-port can be interchangeably used with the term FPE based PW-port.

The following are examples of applications which rely on FPE based PW-port:

- ESM over PW where MPLS/GRE based PW can be rerouted between I/O ports on an SR OS node without affecting ESM service
- Granular QoS per PW since the PW payload is terminated on an access based PW-port SAP 1738 → ingress/egress queues are created per SAP (as opposed to per network port on egress and per MDA on network ingress)
- PW-SAP with MPLS resiliency, where the LSP used by the PW terminated on a PW Ports is protected using MPLS mechanisms such as FRR and could therefore use any port on the system
- PW-port using LDP-over-RSVP tunnels
- A PW Port using a BGP VPWS

Although the primary role of FPE based PW-port is to terminate an external PW, in certain cases PW-port can be used to terminate traffic from regular SAP on I/O ports. This can be used to:

- Separate service termination point from the SAPs which are tied to I/O ports.
- Distribute load from a single I/O port to multiple line cards based on S-Tag (traffic from each S-tag can be mapped to a separate PW associated with different PXCs residing on different line cards).

6.3.1 Cross-Connect Between the External PW and the FPE-Based PW-Port

PW payload delivery from the I/O ports to the FPE based PW-port (and SAP) is facilitated via an internal cross-connect which is built on top of PXC sub-ports. Such cross-connect allows for mapping between PWs and PW-ports even in cases where PW payloads have overlapping VLANS.

This concept is shown in Figure 199.

Figure 199 **Multiplexing PWs over PXC-Based Internal Cross-Connect**

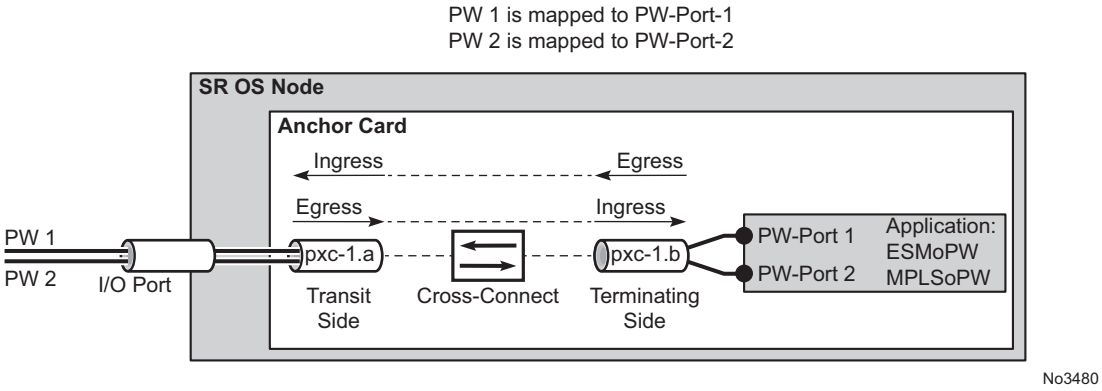

Parameters associated with the PXC sub-ports or PXC based LAGs (QoS, lag-profiles, etc.) are accessible/configurable via CLI. For example, the operator may apply an egress port-scheduler on sub-port pxc-1.b in Figure 199 in order to manage the sum of the bandwidth from associated PW-ports (PW-ports 1 and 2).

To avoid confusion during configuration of PXC sub-ports /LAGs, a clear definition of reference points on the cross-connect created via FPE is required:

- Terminating side of the cross-connect is closer to PW-ports (.b side)
- Transit side of the cross-connect is closer to I/O ports (.a side)

Since the creation of the cross-connect on FPE based PW-ports is highly automated through and FPE configurations, the SR OS system will:

- Assign PXC sub-ports .a to the transit side, and PXC sub-ports .b to the terminating side in case that a single PXC is used; see Figure 200.

Figure 200 **Assign PXC Sub Ports**

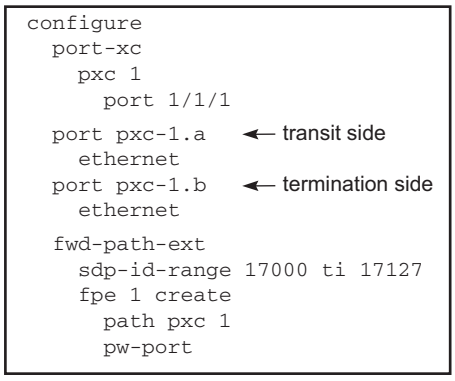

- Assign the xc-a LAG to the transit side, and the xc-b LAG to the terminating side if that a PXC based LAG is used; see Figure 201.

Figure 201 **Assign the LAG**

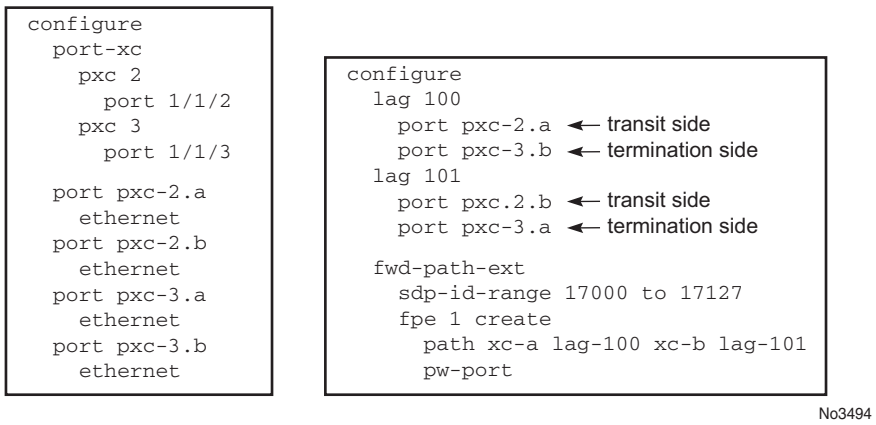

xc-a and xc-b can be associated with any PXC based LAG ID. For example, the following path configuration is allowed: xc-a with lag-id 100 (which includes pxc sub-ports pxc-2.a and pxc-3.b) and xc-b with lag-id 101 (which includes pxc sub-ports pxc-3.a and pxc-2.b). Regardless of the pxc sub-ports that are assigned to respective LAGs, the xc-a side of the path is used as the transit side of the cross-connect, while the xc-b side of the path is used as the termination side of the cross-connect.

6.3.2 PXC-Based PW-Port — Building the Cross-Connect

From a logical perspective, the internal cross-connect that maps the external PW to a PW-port is implemented as a switched Epipe service (**vc-switching**). This switched Epipe service switches an external PW to the internal PW that is terminated on a FPE based PW-port. In this fashion, the PW-port becomes independent of the I/O ports.

Assuming that PXC and PW-port are already configured in the system, the following are the three main configuration steps required to terminate the payload carried over external PW on the PW-port SAP:

1. Auto-setup of the internal transport tunnel over which the cross-connect is built
2. Auto-setup of the internal PW, switching the external PW to the internal PW and terminating the PW on the FPE based PW-port
3. Terminating the service on the PW-SAP

The status of the internally built constructs can be examines via various show commands (for example **show service id** <*epipe-id*> **1 sdp**). The internal SDP id is allocated from the user space. To avoid conflict between the user provisioned SDP ids and the system provisioned SDP ids, a range of SDP ids that will be used for internal consumption must be reserved in advance. This is accomplished via the **sdp-id-range** commands under the **config>fwd-path-ext** hierarchy.

Configuration steps necessary to build PW-port based cross-connect over PXC are shown in the following diagrams (a single PXC is used in this example).

6.3.2.1 Building the Internal Transport Tunnel

The **fpe** command instructs the SR OS system to build an LSP tunnel over the PXC. This tunnel is used to multiplex PW traffic to respective PW-ports. Each external PW is switched to an internal PW (on top of this tunnel) and its payload is off-loaded to a respective PW-port.

After the **fpe** is configured (refer to “Forwarding Path Extensions” in the *Interface Configuration Guide*), the SR OS system will automatically configure steps 1, 2 and 3 in Figure 202. The objects created in steps 1, 2 and 3 can be seen via show commands. However, they are not visible to the operator in the configure branch of CLI.

The significance of the pw-port command under the fpe is to inform the system about the kind of cross-connect that needs to be built over PXC – in this case this cross-connect is PW-port specific. Applications other than PW-port may require different functionality over PXCs and this will be reflected by a different command under the fpe CLI hierarchy (for example vxlan-termination command instead of pw-port).

Note: the IP addresses setup on internal interfaces on PXC sub-ports are Martian IP addresses and they are shown in CLI as fpe_<id>.a and fpe_<id>.b.

Figure 202 **Building the Internal LSP over PXCs**

```
configure
  port-xc
    pxc 2
      port 1/1/2
    pxc 3
      port 1/1/3

  port pxc-2.a
    ethernet
  port pxc-2.b
    ethernet
  port pxc-3.a
    ethernet
  port pxc-3.b
    ethernet
```

```
configure
  lag 100
    port pxc-2.a  <-- transit side
    port pxc-3.b  <-- termination side
  lag 101
    port pxc.2.b  <-- transit side
    port pxc-3.a  <-- termination side

  fwd-path-ext
    sdp-id-range 17000 to 17127
    fpe 1 create
      path xc-a lag-100 xc-b lag-101
      pw-port
```

No3494

6.3.2.2 Mapping the External PW to the PW-Port

Mapping between the external PW and the FPE based PW-port is performed via an Epipe of type vc-switching.

The user configurable Epipe (id 5 in this example) will aid in setting up steps 4, 5 and 6 in Figure 203:

1. An internal PW is automatically added to the user configured Epipe 5
2. A bind is created between the internal PW and the PW-port attached to PXC.
3. External PW is switched to the internal PW.

At this stage, the external PW is mapped to the **pw-port 1**, as shown in Figure 203.

The **spoke-sdp 17000:1** and the binding under SDP 17001 (**spoke-sdp 17001:100001**) created in steps 4 and 5 (Figure 203) can be seen via show commands. However, they are not visible to the operator in the configure branch of CLI.

Figure 203 **Mapping Between the External PW and the PXC Based PW-Port**

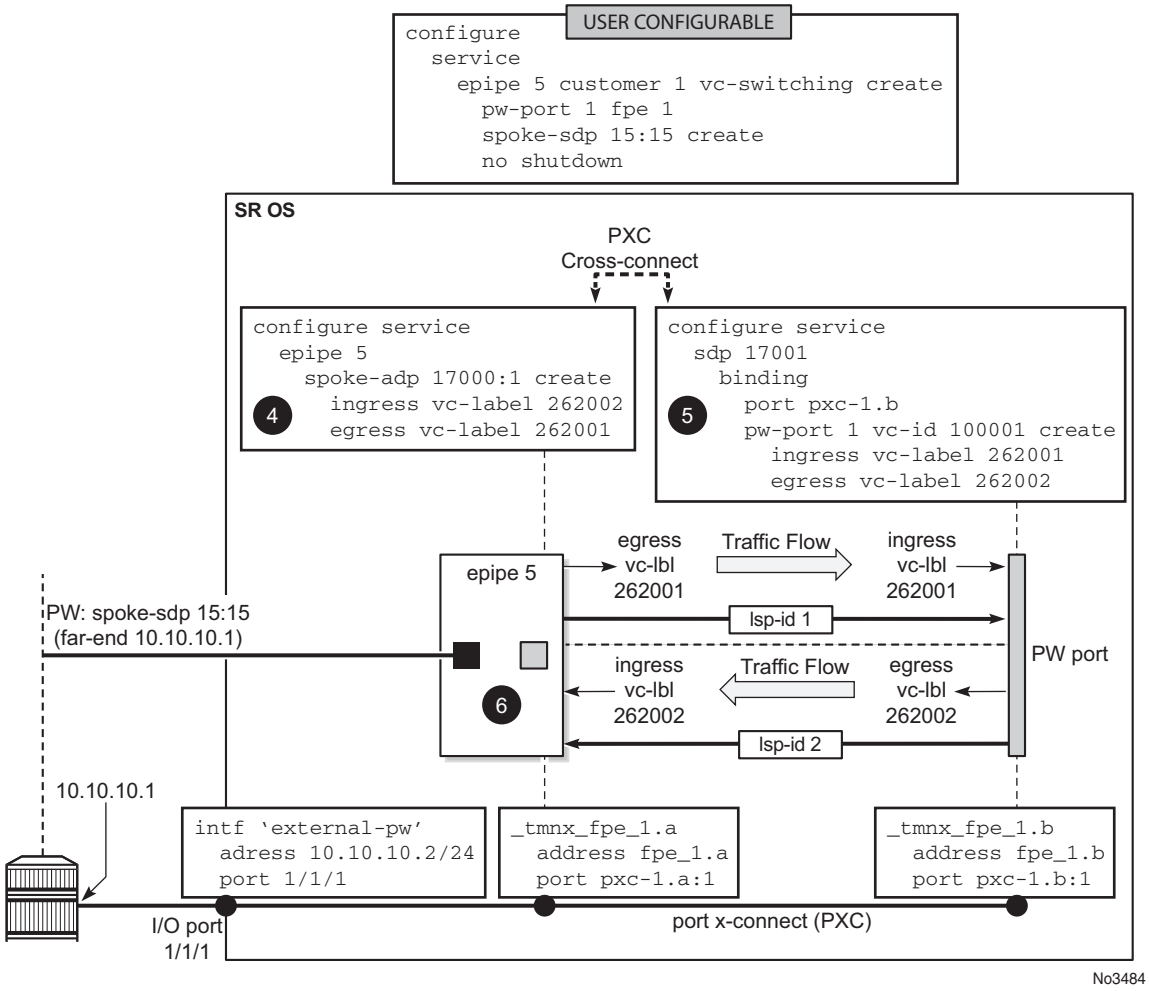

6.3.2.3 Terminating the Service on PW-SAP

In the final step, PW-port SAP is applied to a service (Figure 204).

Figure 204 **Service Termination on PW-SAP**

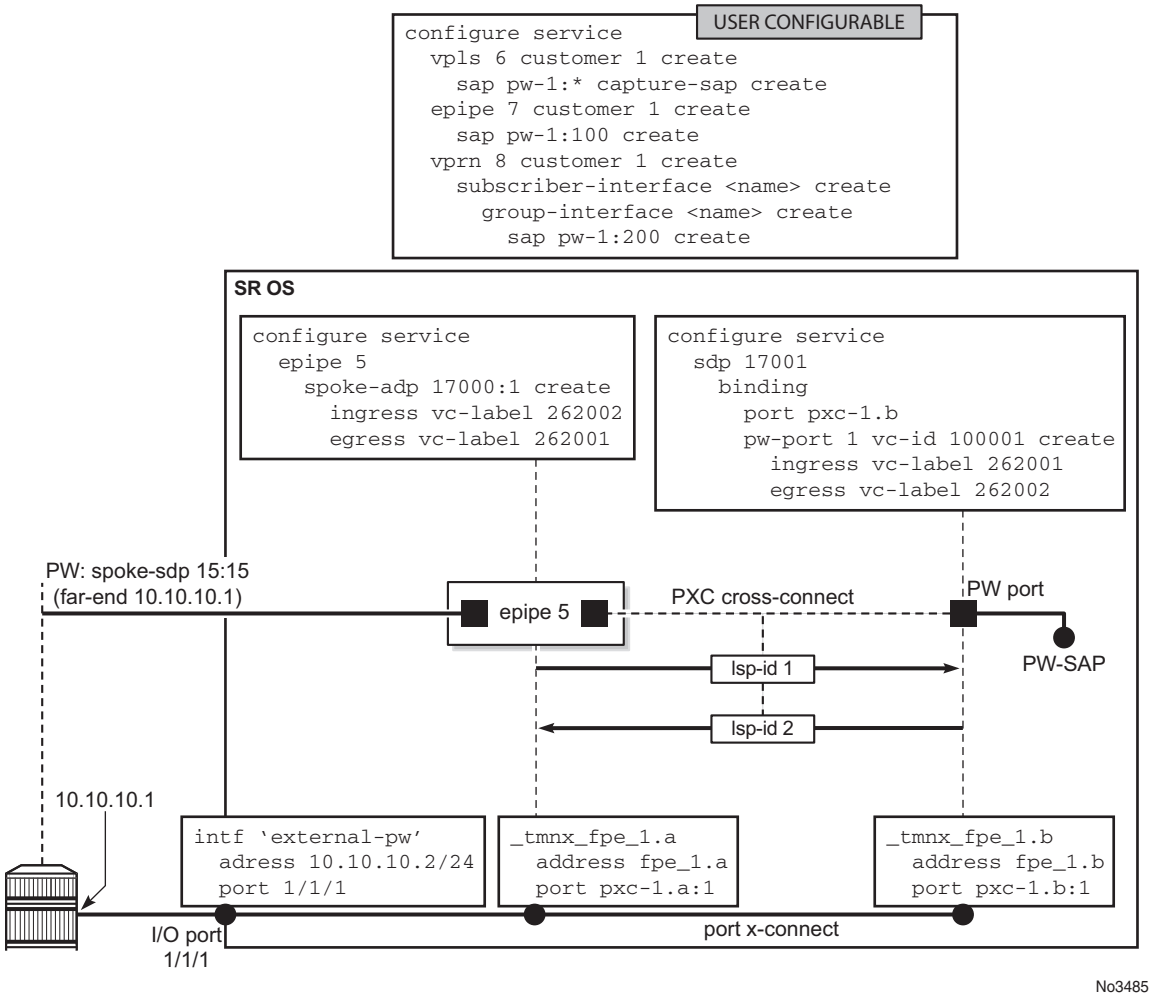

6.3.3 FPE-Based PW-port Operational State

The FPE-based PW-port operational state is driven by the ability of the stitching service (see the following notes) to forward traffic. This includes the stitching service's operational status and, if the external PW is TLDP signaled, the PW status bits. The operational flag for a non-operational PW-port is set to *stitchingSvcTxDown*.

Transitioning of the PW-port into down state due to a PXC failure (for example physical port fails), will bring the stitching service down with the following result:

- In case of TLDP-signaled PW, the *psnIngressFault* and *psnEgressFault* PW status bits is propagated to the remote end, indicating that the local stitching service is down.
- In case of EVPN, the EVPN route will be withdrawn, indicating that the local stitching service is down.
- In case of BGP-VPWS, the BGP-VPWS the ‘D’ bit of the Layer 2 Information Extended Community flag field is set, indicating that the local stitching service is down.

Note: The stitching service in this context is an Epipe service in vc-switching mode for BGP-VPWS or TLDP signaled PW, as follows:

```
configure
      service epipe <epipe-id> customer <cust-id> vc-switching [create]
            pw-port <pw-port-id> fpe <fpe-id>
            spoke-sdp <sdp-id:vc-id> [create]
                  or
            bgp-vpws
                  :
```

Note: The stitching service in this context is an Epipe service in a non-vc-switching mode for EVPN, as follows:

```
configure
      service epipe <epipe-id> customer <cust-id> create
            pw-port <pw-port-id> fpe <fpe-id>
            bgp-evpn
                  :
```

6.3.4 QoS

QoS fundamentals for the case where multiple PWs are multiplexed over a single cross-connect are shown in Figure 205.

Figure 205 **QoS on FPE-Based PW-Port**

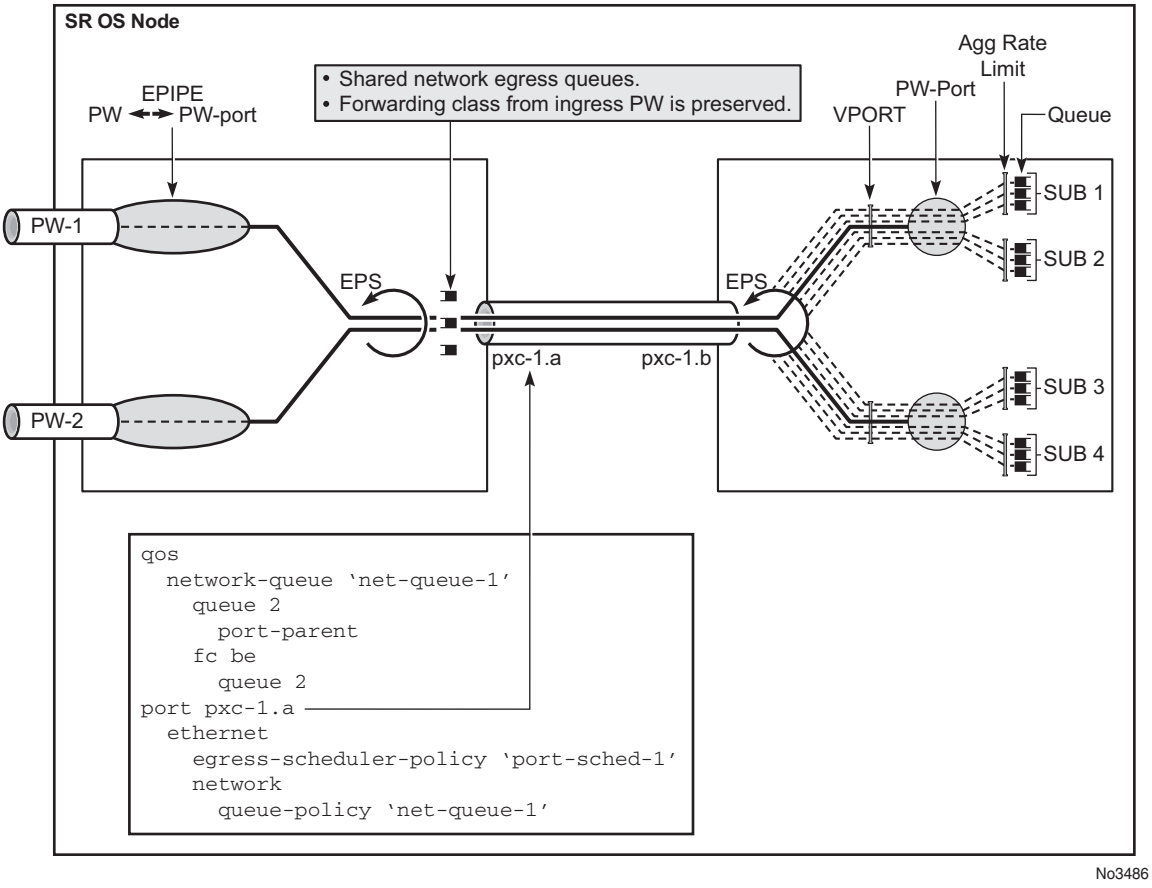

Egress QoS may be applied on both sides of the cross-connect (PXC sub-ports .a and .b) to control congestion on the cross-connect itself. This can be accomplished via an Egress Port Scheduler (EPS) applied to each sub-port.

EPS applied to pxc-1.a (transit side) will manage congestion on the cross-connect for traffic coming from the external PWs. A single set of queues will be shared by all PWs utilizing this cross-connect in this direction.

EPS applied to the pxc-1.b (terminating side) will be used to manage congestion on the cross-connect for traffic going toward the PWs (leaving the SR OS node). A set of queues will be dedicated to each PW-port SAP.

QoS on PXC sub-ports is described in the "PXC" section of the *7450 ESS, 7750 SR, and 7950 XRS Interface Configuration Guide*.

6.3.4.1 Preservation of Forwarding Class Across PXC

The internal cross-connect utilized by FPE based PW-port is relying on an MPLS tunnel built over internal network interfaces configured on PXCs. Those internal network interfaces are using a default network policy 1 for egress traffic classification, remarking and marking purposes. Since the PXC cross-connect is MPLS based, the EXP bits in newly added MPLS header will be marked according to the default network policy (for brevity reasons, only the relevant parts of the network policy are shown here).

```
*A:node-1>config>qos>network# info detail
----------------------------------------------
            description "Default network QoS policy."
            scope template
            egress
                fc af
                    lsp-exp-in-profile 3
                    lsp-exp-out-profile 2
                exit
                fc be
                    lsp-exp-in-profile 0
                    lsp-exp-out-profile 0
                exit
                fc ef
                    lsp-exp-in-profile 5
                    lsp-exp-out-profile 5
                exit
                fc h1
                    lsp-exp-in-profile 6
                    lsp-exp-out-profile 6
                exit
                fc h2
                    lsp-exp-in-profile 4
                    lsp-exp-out-profile 4
                exit
                fc l1
                    lsp-exp-in-profile 3
                    lsp-exp-out-profile 2
                exit
                fc l2
                    lsp-exp-in-profile 1
                    lsp-exp-out-profile 1
                exit
                fc nc
                    lsp-exp-in-profile 7
                    lsp-exp-out-profile 7
                exit
            exit
----------------------------------------------
```

As seen in this excerpt from the default network egress policy, the forwarding classes AF and L1 marks the EXP bits with the same values. This renders the forwarding classes AF and L1 set on one side of PXC, indistinguishable from each other on the other side of the PXC.

This effectively reduces the number of forwarding classes from 8 to 7 in deployment scenarios where the QoS treatment of traffic depends on preservation of forwarding classes across PXC. In other words, in such scenarios, one of the forwarding classes AF or L1 should not be used.

6.3.5 Statistics on the FPE based PW-Port

An FPE-based PW-port is associated with an internal spoke-SDP as described in PXC-Based PW-Port — Building the Cross-Connect and FPE-Based PW-port Operational Stateh. Statistics for the number of forwarded/dropped packets/octets per direction on a PW-port are therefore maintained per this internal spoke-SDP. Octets field counts octets in customer frame (including customer's Ethernet header with VLAN tags).

The following command is used to display PW-port statistics along with the status of the internal spoke-SDP associated with the PW-Port:

```
*A:Dut-B# show pw-port 3 statistics
===============================================================================
Service Destination Point (Sdp Id 17000 Pw-Port 3)
===============================================================================
SDP Binding port     : pxc-1.b
VC-Id                : 100003               Admin Status       : up
Encap                : dot1q                Oper Status        : up
VC Type              : ether
Admin Ingress label  : 262135               Admin Egress label : 262136
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)
Statistics            :
I. Fwd. Pkts.        : 12000                I. Dro. Pkts.      : 0
I. Fwd. Octs.        : 720000               I. Dro. Octs.      : 0
E. Fwd. Pkts.        : 12000                E. Fwd. Octets     : 720000
===============================================================================
```

6.3.6 Intra-Chassis Redundancy Models for PXC-Based PW Port

Intra-chassis redundancy models rely on PXC-based LAG. PXC-based LAG can contain multiple PXCs on the same line card (port redundancy) or PXCs across different line cards (port- and card-level redundancy).

FPE-based PW ports also provide network level-redundancy where MPLS/IP can be rerouted to different I/O ports (due to network failure) without interruption of service.

6.4 L2oGRE Termination on FPE-Based PW Port

L2oGRE termination on an FPE-based PW port allows L2 customer traffic to be transported over an IP network to an SR OS node deeper in the network. In the SR OS node, the customer's payload delivered within L2oGRE tunnel is extracted onto a PW SAP (configured under an interface, group interface, or Epipe) and handed off to an L2 or L3 service. This allows the operator to quickly and simply expand their service offering to their customers over an existing IP network. New service offerings become independent of the existing IP network to which CEs are attached.

For secure operation, the transit IPv4 network should be trusted or alternatively, GRE traffic can be secured by IPSec (L2oGREoIPSec).

Figure 206 shows a typical example where a customer payload is tunneled in GRE through an IPv4 network for an L2 or L3 handoff in an SR OS node that is placed deeper in the network.

Figure 206 **L2oGRE Network Examples**

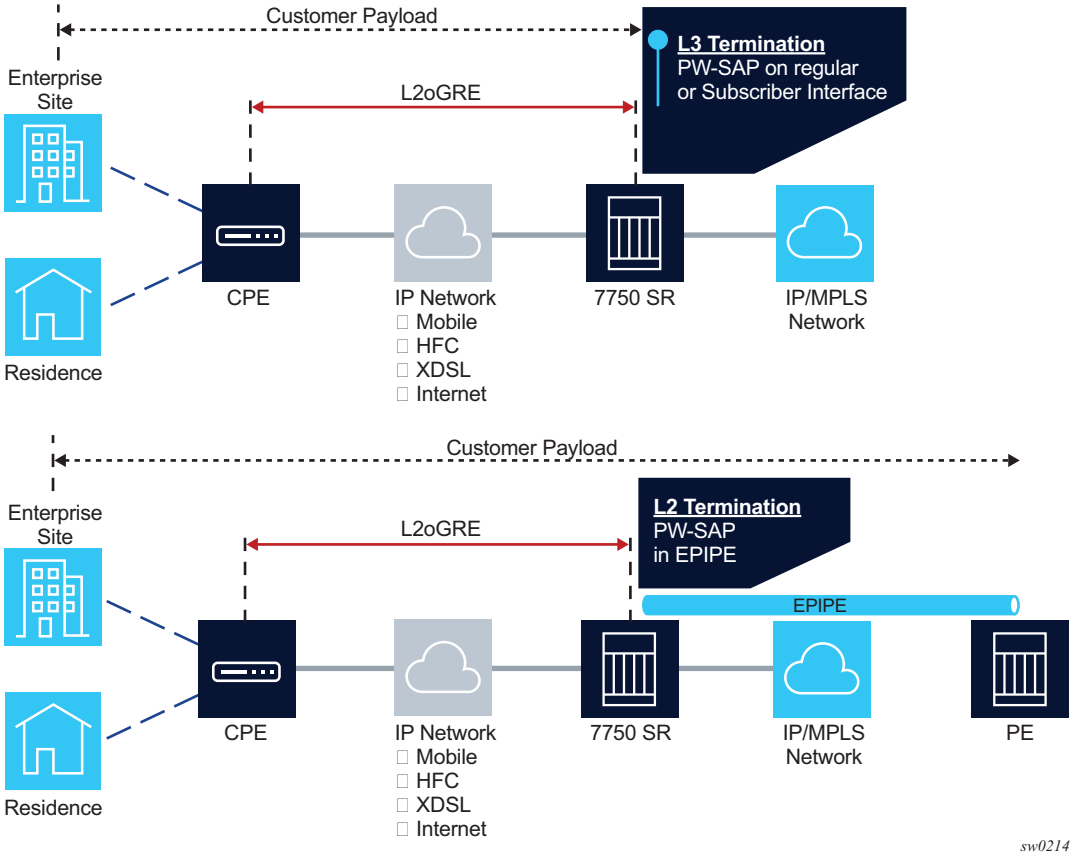

6.4.1 L2oGRE Packet Format

The L2oGRE packet format is as follows:

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |C|       Reserved0       | Ver |         Protocol Type         |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |      Checksum (optional)      |       Reserved1 (Optional)    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The supported GRE header in this context is defined in RFC 2784. The protocol type is set to 0x6558 (bridged Ethernet), and the Checksum and Reserved1 fields are normally omitted. The SR OS can accept headers with those two fields present, but the system omits them when encapsulating packets on transmission. Therefore, the transmitted GRE header length in the SR OS is 4 bytes.

6.4.2 Tracking Payloads and Service Termination Points

A customer payload within L2oGRE can be extracted onto a PW SAP inside an SR OS node. This PW SAP can be configured under an interface, subscriber interface, or an Epipe. Once on a PW SAP, customer traffic can be passed further into the network to its destination using L2 or L3 services. End-to-end L2 and L3 scenarios are described in the following sections.

6.4.2.1 Plain L3 termination

Figure 207 shows an example of plain L3 termination with MTUs.

Figure 207 **L2oGRE MTUs**

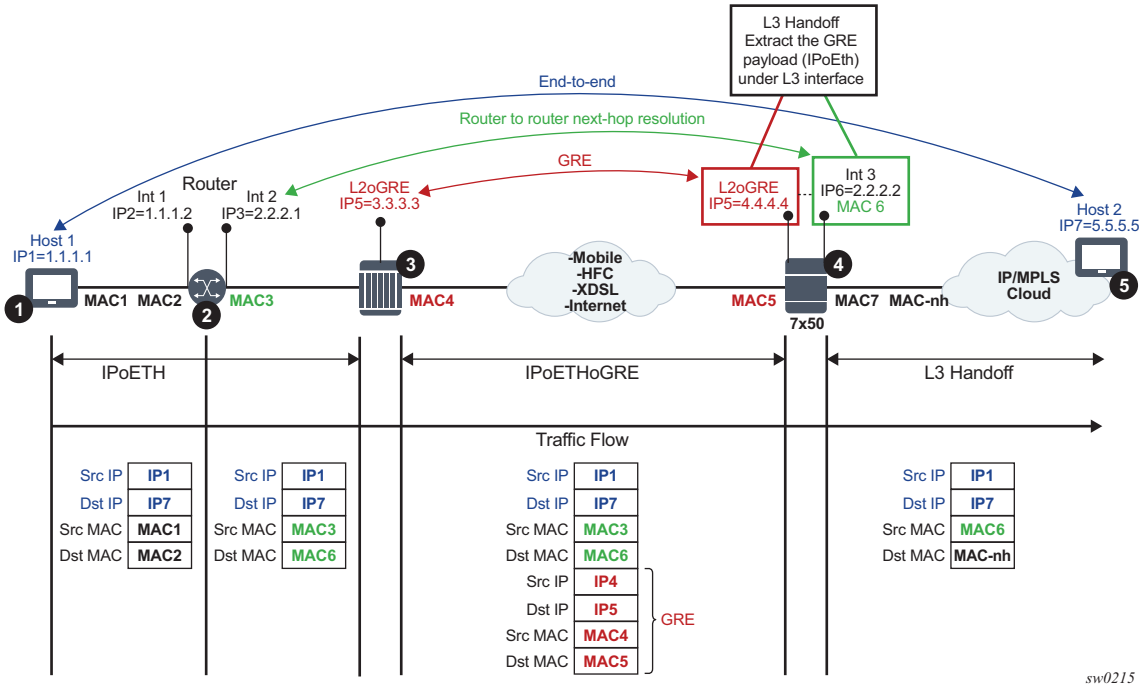

In this example:

- Communication occurs between points 1 and 5.
- There may be a router present at point 2. A router at point 2 would see the SR OS node as L3 next-hop.
- The device in point 3 encapsulates L2 Ethernet frames into GRE and sends them to the SR OS node.
- The SR OS node at point 4 de-encapsulates the packet and performs an L3 lookup on the inner packet in order to deliver it to the destination.

The following is an example where PW SAP is configured under a L3 interface with a PW carrying IP over Ethernet:

```
configure
service vprn 1 customer 1 create
interface example-if
address 192.168.1.1/24
sap pw-1:5.5 create
ingress
filter ip 1000
egress
filter ip 2000
```

6.4.2.2 L2 Termination

Figure 208 shows an example of L2 termination and hand-off.

Figure 208 **L2 Termination and Hand-off**

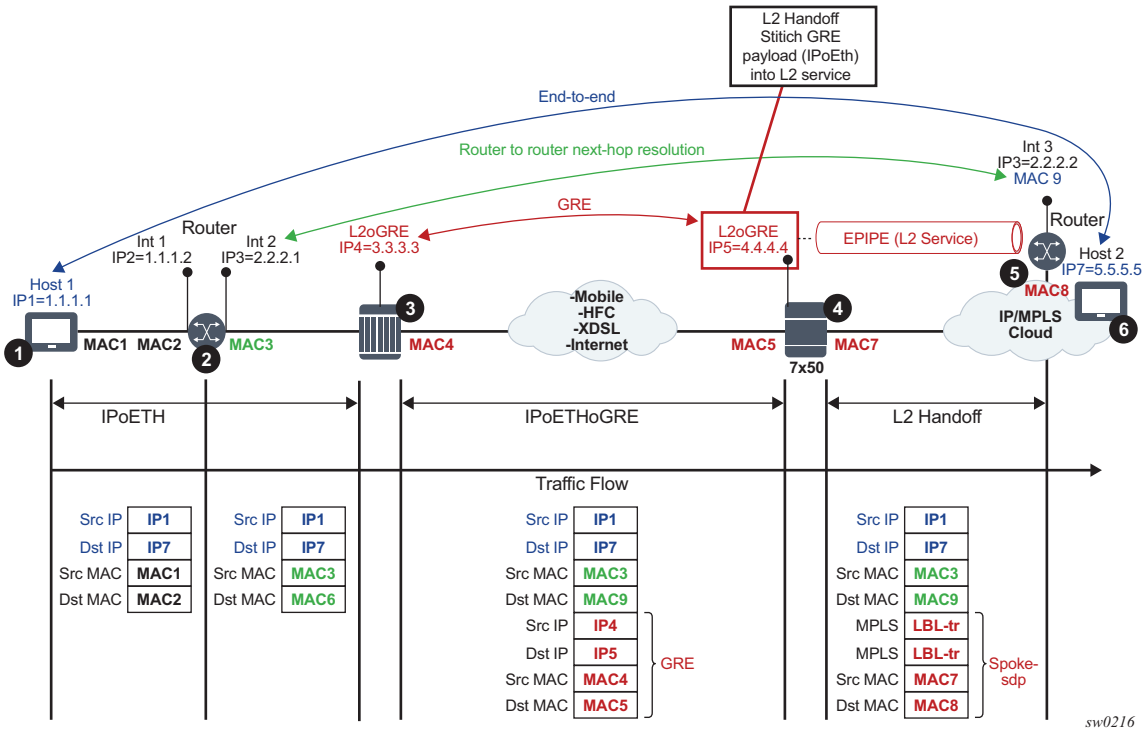

In this example:

- Communication occurs between points 1 and 6.
- There may be a router present at point 2. A router at point 2 would see an L3 device at point 5 as an L3 next-hop. Everything in between is L2.
- The device at point 3 encapsulates L2 Ethernet frames into GRE and sends them to the SR OS node (7x50).
- The SR OS node at point 4 de-encapsulates the packet and sends it into the L2 service that leads to the node at point 5.

The following shows an example of PW SAP configured under an Epipe:

```
configure
service epipe 4 customer 1 create
sap pw-1:2 create
spoke-sdp 1:1
```

6.4.2.3 ESM Termination

The primary case for ESM termination is business services. Figure 209 shows an example of ESM termination.

Figure 209 **ESM Termination**

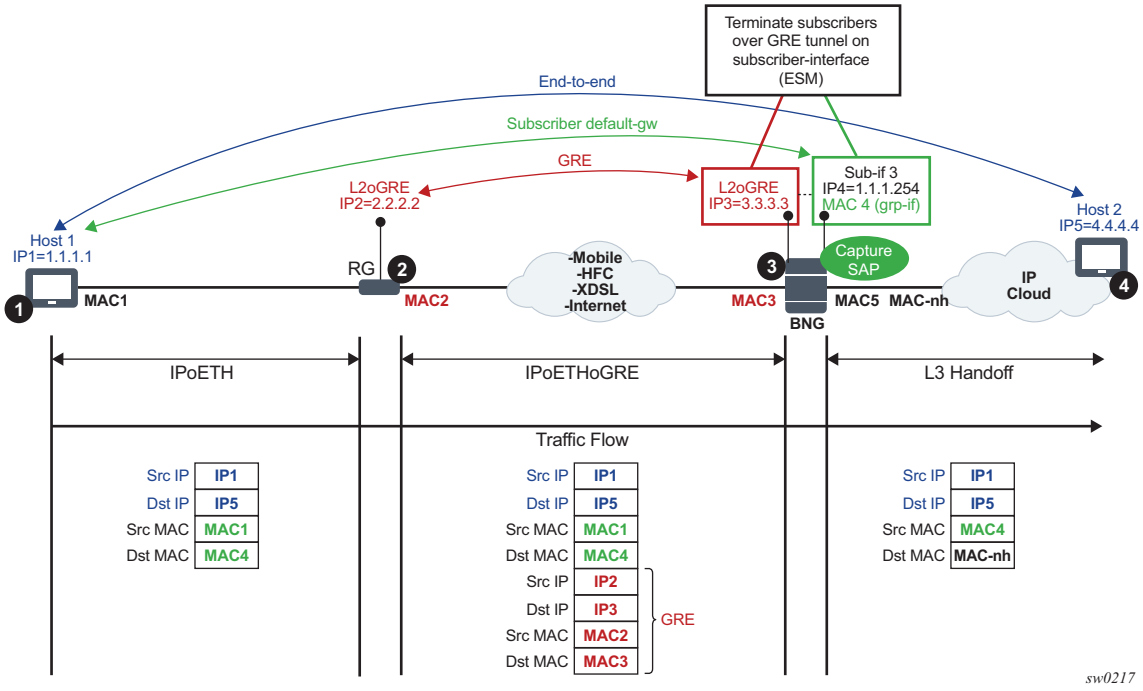

In this example:

- Communication occurs between points 1 and 4.
- RG (Residential GW) at point 2 encapsulates L2 customer frames into GRE and sends them the SR OS node (7x50 BNG).
- The BNG node at point 3 terminates the subscriber traffic, performs an L3 lookup, and sends it to the destination.

The following shows an example where a PW SAP is configured under a subscriber interface:

```
configure
service vprn 3 customer 1 create
interface subscriber-interface <sub-if-name>
address 192.168.1.1/24
group-interface  <grp-if-name>
sap pw-1:10.10 create
```

The following shows an example with the capture of a PW SAP configured:

```
configure
service vpls 2 customer 1 create
trigger-packet dhcp pppoe
sap pw-1:*.* capture-sap create
```

6.4.3 Configuration Steps

L2oGRE tunnels are emulated as an SDP of type **gre-eth-bridged** (shown as GRE-B in the output of relevant show commands). This SDP defines two end-points on the tunnel:

- Far-end IP address

  This defines the IP address of the remote device that terminates the tunnel.
- Local IP address where the tunnel is terminated within an SR OS

  This is a special IP address within an SR OS node that is not associated with any interface. It is only used for L2oGRE tunnel termination.

Binding an L2oGRE tunnel to an FPE-based PW port within the SR OS is performed through an Epipe service. Once the connection is established, the tunnel payload can be extracted to a PW SAP that can be used similarly to a regular SAP under L3 interfaces, subscriber interfaces, or an Epipe.

Table 85 describes the L2oGRE sample configuration steps.

Table 85 **L2oGRE Tunnel Sample Configuration**

Step	Sample CLI	Comments
PXC-based PW Port related configuration		

PW Port creation	`pw-port 1`
	`  encap-type dot1q`
	`  dot1q-etype 0x8100`
		L2oGRE tunnel will be terminated on this PW port.

Table 85 **L2oGRE Tunnel Sample Configuration (Continued)**

Step	Sample CLI	Comments

Port-XC creation	`port-xc`
	`  pxc 1 create`
	`  port 1/1/1`
		This command triggers automatic creation of PXC sub-ports:
		
		`configure`
		`  port pxc-1.a`
		`  port pxc-1.b`
		
		This is where the L2oGRE terminating PW port will be anchored.

Creation of FPE that will be used for PW port anchoring	`fwd-path-ext`
	`  sdp-id-range from 17400 to 17500`
	`  fpe 1 create`
	`  path pxc 1`
	`  pw-port`
		The application under this FPE will be the PW port termination. The use of PW port in this case is versatile and can be used to terminate an L2oGRE or MPLS/GRE-based PW. In this example, it will be used to terminate an L2oGRE tunnel.

L2oGRE tunnel definition		
Configuration of GRE-bridged tunnel termination IPv4 addresses.	`service>system>gre-eth-bridged` `  tunnel-termination 10.1.1.2 fpe 1`	This is a special IPv4 address that is not configured under any L3 interface and it must not overlap with any IPv4 address configured under an L3 interface in Base router. Multiple termination IPv4 addresses are supported.
Configuration of L2oGRE SDP	`service> sdp 2 gre-eth-bridged create` `  far-end 11.1.1.2 local-end 10.1.1.2`	This represents the L2oGRE tunnel within SR OS as defined by the tunnel end-point IPv4 addresses.
Stitching L2oGRE tunnel to an anchored PW port		
Association between the PW port and a PXC port via FPE.	`service>epipe 1` `  pw-port 1 fpe 1`	This commands anchors the PW port 1 to a PXC port referenced in FPE 1.

Table 85 **L2oGRE Tunnel Sample Configuration (Continued)**

Step	Sample CLI	Comments

Binding between L2oGRE tunnel and the PW port	`service>epipe 1`
	`  pw-port 1 fpe 1`
	`  spoke-sdp 2:1`
		L2oGRE will be terminated on a PW port and the L2 payload within the tunnel will be extracted on the PW SAP

PW SAP service association		

Creation of services that use PW SAP	`service>epipe 100`
	`  sap pw-1:100 create`
	
	`service>vprn 101>if`
	`  sap pw-1:101 create`	

6.4.4 Fragmentation and MTU Configuration

L2oGRE traffic is subjected to several MTU checks in the downstream direction (toward the remote end of the L2oGRE tunnel) within the SR OS node, as shown in Figure 209.

Figure 210 **L2oGRE MTUs**

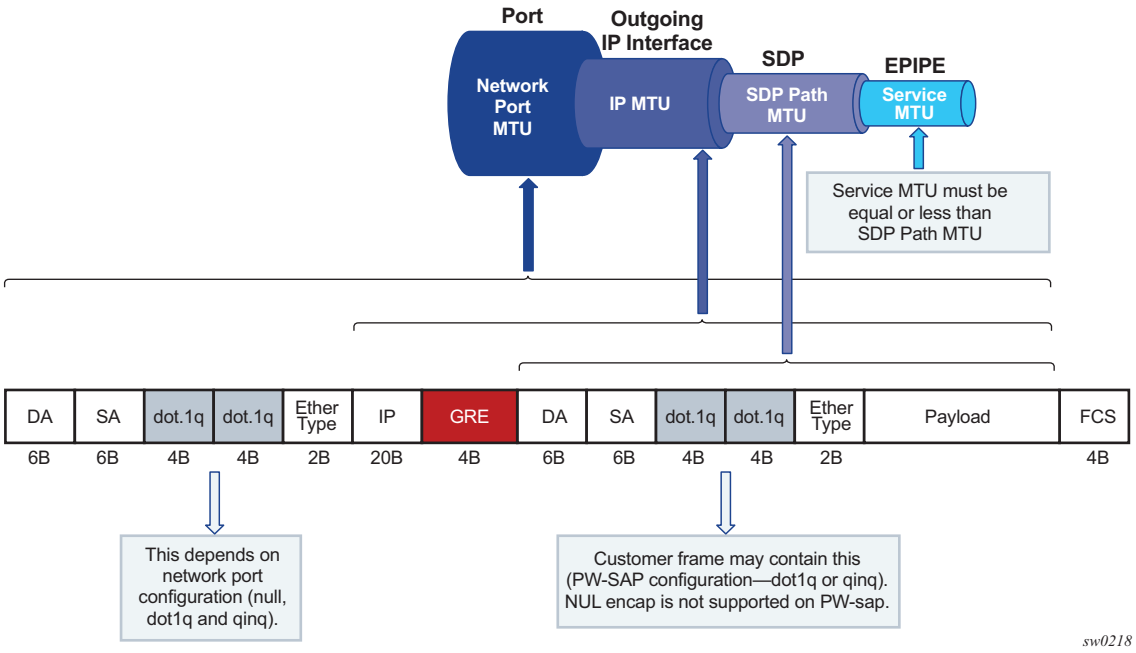

In this example:

- Port MTU represents the maximum frame size on the outgoing physical port.
- IP MTU is the maximum IP packet size on the outgoing IP interface.
- SDP Path MTU represents the maximum size of a frame that is encapsulated with the GRE tunnel. Its value is determined by the smallest MTU size on the path between the two GRE tunnel terminating end-points. The SDP Path MTU is calculated automatically by subtracting transport IP and GRE header bytes (24 bytes) from the configured IP MTU of the outgoing interface.
- Service MTU indicates the maximum frame size that the customer can accept over the service (PW SAP). Its value is determined by the MTU size within the customer's network. The service MTU is configured within the vc-switching Epipe that stitches the L2oGRE spoke-SDP to a PW port. The default value is set to 1514 bytes.

MTU values:

- Port MTU = 1600 bytes (this is operator's configured value)
- IP MTU (of the outgoing interface) = 1600 bytes- 22 bytes= 1578 bytes (this is operator's configured value)
- SDP Path MTU = automatically calculated and set to 1578 bytes - 24 bytes = 1554 bytes.

- Service MTU = This value must be configured to a value no higher than 1554 bytes (SDP Path MTU).

Frames within an SR OS cannot be fragmented on a service or SDP level. However, L2oGRE traffic can be fragmented at the port level and at any downstream point, if the DF bit in the IP header is cleared. The DF bit setting is controlled by the **config>service>sdp>allow-fragmentation** command.

6.4.5 Reassembly

L2oGRE reassembly is supported through a generic reassembly function that requires an MS-ISA. As fragmented traffic enters an SR OS node, it is redirected to an MS-ISA via filters. Once the traffic is reassembled in the MS-ISA, it is re-inserted into the forwarding complex where normal processing continues (as if the non-fragmented traffic originally entered the node).

Table 86 describes the configuration steps to support reassembly for GRE.

Table 86 **Configuring Reassembly For GRE**

Step	Sample CLI	Comments

Creation of a NAT-group that contains MS-ISAs	`configure isa nat-group 1`
	` mda 1/1`
	`mda 2/1`
		The reassembly function is performed in a NAT group that contains one or more MS-ISAs.

Referencing a reassembly group that will be used for traffic in the Base routing context	`configure router` ` reassembly-group 1`	Identification of the reassembly group that will be used for traffic in the Base routing context. Upon reassembly, traffic will be re-inserted in the same (Base) routing context. Reassembly group ID corresponds to the NAT group ID (in this case 1). There can be multiple NAT groups (reassembly groups) configured in the system and this command identifies the reassembly group that will be used in the Base routing context.

Table 86 **Configuring Reassembly For GRE (Continued)**

Step	Sample CLI	Comments

Identifying and directing fragmented traffic to the reassembly function.	`configure filter ip-filter <id>`
	` default-action forward`
	`  entry <id> create`
	`   match  protocol gre`
	`    fragment true`
	`   exit`
	`   action reassemble`
	`  exit`
		Fragmented GRE traffic is identified via a filter and redirected to the reassembly function. This filter must be applied to all ingress interfaces on which GRE traffic is expected to arrive.

6.5 Pseudowire Ports Command Reference

This chapter describes the pseudowire ports (PW-ports) command reference.

6.5.1 Command Hierarchies

6.5.1.1 PW-port Configuration Commands

```
config
    — pw-port id [create]
    — no pw-port
        — description description-string
        — no description
        — dot1q-etype dot1q-etype
        — no dot1q-etype
        — encap-type {dot1q | qinq}
        — no encap-type
        — qinq-etype qinq-etype
        — no qinq-etype
```

6.5.1.2 Redundant Interface Commands

```
config
    — service
        — sdp sdp-id [gre | mpls]
        — no sdp sdp-id
            — binding
                — port [port-id | lag-id]
                — no port
                — pw-port pw-port-id [vc-id vc-id] [create]
                — no pw-port pw-port-id
                    — description description-string
                    — no description
                    — egress
                        — shaper
                            — int-dest-id int-dest-id
                            — no int-dest-id int-dest-id
                            — vport vport-name
                            — no vport vport-name
                    — encap-type {dot1q | qinq}
                    — no encap-type
                    — [no] shutdown
                    — vc-type {ether | vlan}
```

— **no vc-type**
— **vlan-vc-tag** *vlan-id*
— **no vlan-vc-tag**

6.5.1.3 Show Commands

show
— **pw-port**
— **pw-port** [*pw-port-id*] [**detail**]
— **pw-port sdp** [*sdp-id*]
— **pw-port sdp none**

6.5.2 Command Descriptions

6.5.2.1 PW-port Configuration Commands

pw-port

Syntax **pw-port** *id* [**create**]
no pw-port

Context config

Description This command creates a PW-port that can be bound to a physical port or associated with an FPE (anchored PW-port). A PW-port's purpose is to provide, through a PW-SAP, access level (or SAP level) capability to customer traffic that is tunneled to the SR OS node through an IP/MPLS network.

The **no** form of this command removes the **pw-port**.

Default no pw-port

Parameters *id* — ID of the PW-port.

Values 1 to 32767

create — Keyword required to create the configuration context.

description

Syntax **description** *description-string*

no description

Context config>pw-port

Description This command adds a text description to the **pw-port**.

The **no** form of this command removes the text description.

Parameters *description-string* — Specifies the description character string of the configuration context.

Values Any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

dot1q-etype

Syntax **dot1q-etype** *dot1q-etype*
no dot1q-etype

Context config>pw-port

Description This command configures the Dot1q Ethertype on the PW-port. The PW-port is used to extract a customer's Ethernet traffic that is transported in a tunnel over an IP/MPLS network. The **Dot1q-etype** represents the first two bytes (TPID) in the 802.1Q header of a single-tagged Ethernet frame or the inner 802.1Q header of the double-tagged Ethernet frame inside the tunnel.

The **no** form of this command removes the configuration.

Parameters *dot1q-etype* — The value for the **dot1q-etype** field, in hexadecimal format.

Values 0x0600..0xFFFF

Default 0x8100

encap-type

Syntax **encap-type** {**dot1q** | **qinq**}
no encap-type

Context config>pw-port

Description This command configures the encapsulation type on a PW-port. Customer Ethernet frames can be single-tagged or double-tagged, and this command determines the number of tags that the SR OS will check (and strip) on PW-SAP ingress and insert on PW-SAP egress.

The **no** form of this command removes the configuration.

Parameters **dot1q** — Specifies that the encapsulation type is dot1q; used when the customer's Ethernet frame is single-tagged.

qinq — Specifies that the encapsulation type is qinq; used when the customer's Ethernet frame is double-tagged.

Default dot1q

qinq-etype

Syntax **qinq-etype** *qinq-etype*
no qinq-etype

Context config>pw-port

Description This command configures the QinQ Ethertype on the PW-port. The PW-port is used to extract a customer's Ethernet traffic that is transported in a tunnel over an IP/MPLS network. The **qinq-etype** represents the first two bytes (TPID) in the outer 801.1Q header of the double-tagged Ethernet frame inside the tunnel.

The **no** form of this command removes the configuration.

Parameters *qinq-etype* — The value for the **qinq-etype** field, in hexadecimal format.

Values 0x0600..0xFFFF

Default 0x8100

6.5.2.2 SDP Binding Commands

binding

Syntax **binding**

Context config>service>sdp

Description The command enters the context to configure SDP bindings.

port

Syntax **port** [*port-id* | *lag-id*]
no ort

Context config>service>sdp>binding

Description This command specifies the port or lag identifier, to which the PW ports associated with the underlying SDP are bound. If the underlying SDP is re-routed to a port or lag other than the specified one, the PW ports on the SDP are operationally brought down.

The **no** form of the command removes the value from the configuration.

Parameters *port-id* — Specifies the identifier of the port in the slot/mda/port format.

lag-id — Specifies the LAG identifier.

pw-port

Syntax **pw-port** *pw-port-id* [vc-id *vc-id*] [**create**]
no pw-port

Context config>service>sdp>binding

Description This command creates a pseudowire port.

The **no** form of the command removes the pseudowire port ID from the configuration.

Parameters *pw-port-id* — Specifies a unique identifier of the pseudowire port.

Values 1 to 10239

vc-id *vc-id* — Specifies a virtual circuit identifier signaled to the peer.

Values 1 to 4294967295

create — Keyword required to create the configuration context.

description

Syntax **description** *description-string*
no description

Context config>service>sdp>binding>pw-port

Description This command creates a text description stored in the configuration file for a configuration context.

The description command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of the command removes the string from the configuration.

Default no description

Parameters *description-string* — Specifies the description character string of the configuration context.

Values Any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

egress

Syntax **egress**

Context config>service>sdp>binding>pw-port

Description This command enters the context to configure PW-port egress side parameters.

encap-type

Syntax **encap-type** {**dot1q** | **qinq**}
no encap-type

Context config>service>sdp>binding>pw-port

Description This command sets the encapsulation type for the PW-port as dot1q or qinq.

The **no** form of the command reverts to the default value.

Default dot1q

Parameters **dot1q** — Specifies **dot1q** encapsulation type.

qinq — Specifies **qinq** encapsulation type.

shutdown

Syntax **no shutdown**

Context config>service>sdp>binding>pw-port

Description This command changes the administrative status of the PW-port.

Default shutdown

shaper

Syntax [**no**] **shaper**

Context config>service>sdp>binding>pw-port>egress

Description This command configures an egress shaping option for use by a PW port.

The **no** form of the command reverts to the default value.

Default no shaper

int-dest-id

Syntax [**no**] **int-dest-id** *int-dest-id*

Context config>service>sdp>binding>pw-port>egress>shaper

Description This command specifies the intermediate destination string configured for dynamic Vport selection.

This command is only valid for PW ports used for enhanced subscriber management (ESM on PW).

The **no** form of the command removes the configured intermediate destination string.

Default no.int-dest-id

Parameters *int-dest-id* — Specifies a text string that describes the intermediate destination ID.

vport

Syntax **vport** *vport-name*
no vport

Context config>service>sdp>binding>pw-port>egress>shaper

Description This command configures the name of the Vport to be used for the PW port.

This command is valid for PW ports used for enhanced subscriber management (ESM on pseudowire) and pseudowire SAPs on Ethernet ports. It is not valid for pseudowire ports on the HSMDA.

The **no** form of the command removes the configured Vport name.

Default no vport

Parameters *vport-name* — Specifies a text string up to 32 characters in length representing the name of the Vport.

vc-type

Syntax **vc-type** {**ether** | **vlan**}
no vc-type

Context config>service>sdp>binding>pw-port

Description This command sets the forwarding mode for PW-port. The vc-type is signaled to the peer, and must be configured consistently on both ends of the PW. vc-type VLAN is only configurable with dot1q encapsulation on the PW-port. The tag with vc-type vlan only has significance for transport, and is not used for service delineation or ESM. The top (provider tag) is stripped while forwarding out of the PW, and a configured vlan-tag (for vc-type vlan) is inserted when forwarding into the PW. With vc-type ether, the tags if present (max 2), are transparently preserved when forwarding in our out of the PW.

The **no** form of the command reverts to the default value.

Default ether

Parameters **ether** — Specifies **ether** as the virtual circuit (VC) associated with the SDP binding.

vlan — Specifies **vlan** as the virtual circuit (VC) associated with the SDP binding.

vlan-vc-tag

Syntax **vlan-vc-tag** *vlan-id*
no vc-type

Context config>service>sdp>binding>pw-port

Description This command sets tag relevant for vc-type vlan mode. This tag is inserted in traffic forwarded into the PW.

The **no** form of the command reverts to the default value.

Default 0

Parameters *vlan-id* — Specifies the VLAN ID value

Values 0 to 4094

6.5.2.3 Show Commands

pw-port

Syntax **pw-port** [*pw-port-id*] [**detail**]

pw-port sdp *sdp-id*
pw-port sdp none

Context show>pw-port

Description Displays pseudo-wire port information.

If no optional parameters are specified, the command displays a summary of all defined PW ports. The optional parameters restrict output to only ports matching the specified properties.

Parameters *pw-port-id* — Specifies the pseudo-wire port identifier.

Values 1 to 10239

detail — Displays detailed port information that includes all the **pw-port** output fields.

sdp-id — The SDP ID for which to display matching PW port information.

Values 1 to 17407

Output The following is an example of PW port information.

Sample Output

```
*A:ALA-48>config>service# show pw-port
===============================================================================
PW Port Information
===============================================================================
PW Port   Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
1         dot1q          1          1526726657          1
2         qinq           1          1526726658          2
3         dot1q          1          1526726659          3
4         qinq           1          1526726660          4
===============================================================================

*A:ALA-48>config>service# show pw-port 3
===============================================================================
PW Port Information
===============================================================================
PW Port   Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
3         dot1q          1          1526726659          3
===============================================================================
*A:ALA-48>config>service# show pw-port 3 detail
===============================================================================
PW Port Information
===============================================================================
PW Port            : 3
Encap              : dot1q
SDP                : 1
IfIndex            : 1526726659
VC-Id              : 3
Description        : 1-Gig Ethernet dual fiber
===============================================================================
*A:ALA-48>config>pw-port$ show pw-port sdp none

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

```
PW Port Information
===============================================================================
PW Port   Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
5         dot1q                     1526726661
===============================================================================


*A:ALA-48>config>pw-port$ show pw-port sdp 1
===============================================================================
PW Port Information
===============================================================================
PW Port   Encap          SDP        IfIndex             VC-Id
-------------------------------------------------------------------------------
1         dot1q          1          1526726657          1
2         qinq           1          1526726658          2
3         dot1q          1          1526726659          3
4         qinq           1          1526726660          4
===============================================================================
```

The following table describes **show pw-port** output fields:

Table 87 **Subscriber Show PW-Port Field Descriptions**

Label	Description
PW Port	The PW port identifier.
Encap	The encapsulation type of the PW port.
SDP	The SDP identifier.
IfIndex	The interface index used for the PW port.
VC-Id	The Virtual Circuit identifier.
Description	The description string for the PW port.

7 Standards and Protocol Support

Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attrib-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *Multiprotocol Extensions for BGP-4*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 3107, *Carrying Label Information in BGP-4*

RFC 3392, *Capabilities Advertisement with BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*

RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*

RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/ MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*

RFC 4724, *Graceful Restart Mechanism for BGP* (helper mode)

RFC 4760, *Multiprotocol Extensions for BGP-4*

RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)*

RFC 4893, *BGP Support for Four-octet AS Number Space*

RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*

RFC 5065, *Autonomous System Confederations for BGP*

RFC 5291, *Outbound Route Filtering Capability for BGP-4*

RFC 5396, *Textual Representation of Autonomous System (AS) Numbers* (asplain)

RFC 5575, *Dissemination of Flow Specification Rules*

RFC 5668, *4-Octet AS Specific BGP Extended Community*

RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*

RFC 6811, *Prefix Origin Validation*

RFC 6996, *Autonomous System (AS) Reservation for Private Use*

RFC 7311, *The Accumulated IGP Metric Attribute for BGP*

RFC 7607, *Codification of AS 0 Processing*

RFC 7674, *Clarification of the Flowspec Redirect Extended Community*

RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP*

RFC 7911, *Advertisement of Multiple Paths in BGP*

Circuit Emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*
IEEE 802.1aq, *Shortest Path Bridging*
IEEE 802.1ax, *Link Aggregation*
IEEE 802.1D, *MAC Bridges*
IEEE 802.1p, *Traffic Class Expediting*
IEEE 802.1Q, *Virtual LANs*
IEEE 802.1s, *Multiple Spanning Trees*
IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*
IEEE 802.1X, *Port Based Network Access Control*
IEEE 802.3ab, *1000BASE-T*
IEEE 802.3ac, *VLAN Tag*
IEEE 802.3ad, *Link Aggregation*
IEEE 802.3ae, *10 Gb/s Ethernet*
IEEE 802.3ah, *Ethernet in the First Mile*
IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*
IEEE 802.3i, *Ethernet*
IEEE 802.3u, *Fast Ethernet*
IEEE 802.3x, *Ethernet Flow Control*
IEEE 802.3z, *Gigabit Ethernet*
ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*
ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*
ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*
draft-ietf-bess-evpn-etree-11, *E-TREE Support in EVPN & PBB-EVPN*
draft-ietf-bess-evpn-overlay-04, *A Network Virtualization Overlay Solution using EVPN*
draft-ietf-bess-evpn-prefix-advertisement-02, *IP Prefix Advertisement in EVPN*
draft-ietf-bess-evpn-proxy-arp-nd-02, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*
draft-ietf-bess-evpn-vpls-seamless-integ-00, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*
draft-ietf-bess-evpn-vpws-14, *Virtual Private Wire Service support in Ethernet VPN*
draft-rabadan-bess-evpn-pref-df-02, *Preference-based EVPN DF Election*
draft-snr-bess-pbb-evpn-isid-cmacflush-01, *PBB-EVPN ISID-based CMAC-Flush*

RFC 7432, *BGP MPLS-Based Ethernet VPN*

RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

Frame Relay

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*

FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*

FRF.12, *Frame Relay Fragmentation Implementation Agreement*

FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*

FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*

ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 4204, *Link Management Protocol (LMP)*

RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*

RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*

RFC 5063, *Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart* (helper mode)

Intermediate System to Intermediate System (IS-IS)

draft-ginsberg-isis-mi-bis-01, *IS-IS Multi-Instance* (single topology)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

RFC 2973, *IS-IS Mesh Groups*

RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*

RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*

RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*

RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*

RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*

RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*

RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering TE*

RFC 5306, *Restart Signaling for IS-IS* (helper mode)

RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 5308, *Routing IPv6 with IS-IS*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 6213, *IS-IS BFD-Enabled TLV*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*

RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*

RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*

RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*

RFC 7431, *Multicast-Only Fast Reroute*

RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*

RFC 768, *User Datagram Protocol*

RFC 793, *Transmission Control Protocol*

RFC 854, *Telnet Protocol Specifications*

RFC 1350, *The TFTP Protocol (revision 2)*

RFC 2347, *TFTP Option Extension*

RFC 2348, *TFTP Blocksize Option*

RFC 2349, *TFTP Timeout Interval and Transfer Size Options*

RFC 2428, *FTP Extensions for IPv6 and NATs*

RFC 2784, *Generic Routing Encapsulation (GRE)*

RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*

RFC 4251, *The Secure Shell (SSH) Protocol Architecture*

RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)

RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*

RFC 4254, *The Secure Shell (SSH) Connection Protocol*

RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*

RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*

RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer* (ECDSA)

RFC 6398, *IP Router Alert Considerations and Usage* (MLD)

RFC 6528, *Defending against Sequence Number Attacks*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)

draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*

draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*

draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*

RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)* (auto-RP groups)

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*

RFC 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*

RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*

RFC 4607, *Source-Specific Multicast for IP*

RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*

RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*

RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*

RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*

RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*

RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6513, *Multicast in MPLS/BGP IP VPNs*

RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*

RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*

RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*

RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*

RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*

RFC 1534, *Interoperation between DHCP and BOOTP*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 1812, *Requirements for IPv4 Routers*

RFC 1918, *Address Allocation for Private Internets*

RFC 2003, *IP Encapsulation within IP*

RFC 2131, *Dynamic Host Configuration Protocol*

RFC 2132, *DHCP Options and BOOTP Vendor Extensions*

RFC 2401, *Security Architecture for Internet Protocol*

RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

RFC 3046, *DHCP Relay Agent Information Option (Option 82)*

RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*

RFC 4884, *Extended ICMP to Support Multi-Part Messages* (ICMPv4 and ICMPv6 Time Exceeded)

Internet Protocol (IP) — Version 6

RFC 1981, *Path MTU Discovery for IP version 6*

RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*

RFC 2473, *Generic Packet Tunneling in IPv6 Specification*

RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*

RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*

RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3587, *IPv6 Global Unicast Address Format*

RFC 3596, *DNS Extensions to Support IP version 6*

RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*

RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*

RFC 3971, *SEcure Neighbor Discovery (SEND)*

RFC 3972, *Cryptographically Generated Addresses (CGA)*

RFC 4007, *IPv6 Scoped Address Architecture*

RFC 4193, *Unique Local IPv6 Unicast Addresses*

RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*

RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*

RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*

RFC 4862, *IPv6 Stateless Address Autoconfiguration* (router functions)

RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*

RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*

RFC 5007, *DHCPv6 Leasequery*

RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*

RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6* (IPv6)

RFC 5952, *A Recommendation for IPv6 Address Text Representation*

RFC 6092 *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service* (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)

RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*

RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*

draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*

RFC 2401, *Security Architecture for the Internet Protocol*

RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*

RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*

RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*

RFC 2406, *IP Encapsulating Security Payload (ESP)*

RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*

RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*

RFC 2409, *The Internet Key Exchange (IKE)*

RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*

RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*

RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*

RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*

RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*

RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*

RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*

RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*

RFC 4301, *Security Architecture for the Internet Protocol*

RFC 4303, *IP Encapsulating Security Payload*

RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*

RFC 4308, *Cryptographic Suites for IPsec*

RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*

RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPSec*

RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*

RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*

RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*

RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*

RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*

RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*

RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*

draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*

draft-pdutta-mpls-mldp-up-redundancy-00, *Upstream LSR Redundancy for Multi-point LDP Tunnels*

draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*

draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*

RFC 3037, *LDP Applicability*

RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol* (helper mode)

RFC 5036, *LDP Specification*

RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*

RFC 5443, *LDP IGP Synchronization*

RFC 5561, *LDP Capabilities*

RFC 5919, *Signaling LDP Label Advertisement Completion*

RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 7032, *LDP Downstream-on-Demand in Seamless MPLS*

RFC 7552, *Updates to LDP for IPv6*

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*

RFC 2661, *Layer Two Tunneling Protocol "L2TP"*

RFC 2809, *Implementation of L2TP Compulsory Tunneling via RADIUS*

RFC 3438, *Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*

RFC 3931, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*

RFC 4719, *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*

RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

Management

draft-ieft-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6* (IPv6)

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

RFC 1157, *A Simple Network Management Protocol (SNMP)*

RFC 1212, *Concise MIB Definitions*

RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 1901, *Introduction to Community-based SNMPv2*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*

RFC 2515, *Definitions of Managed Objects for ATM Management*

RFC 2570, *SNMP Version 3 Framework*

RFC 2571, *An Architecture for Describing SNMP Management Frameworks*

RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*

RFC 2573, *SNMP Applications*

RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

RFC 2579, *Textual Conventions for SMIv2*

RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416. *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP)* (SNMP over UDP over IPv4)

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*

RFC 4273, *Definitions of Managed Objects for BGP-4*

RFC 4292, *IP Forwarding Table MIB*

RFC 4293, *Management Information Base for the Internet Protocol (IP)*

RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*

RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*

RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms* (TLS)

RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*

RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*

RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*

RFC 5102, *Information Model for IP Flow Information Export*

RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2* (TLS client, RSA public key)

RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*

RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*

RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*

SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching — Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*

RFC 5921, *A Framework for MPLS in Transport Networks*

RFC 5960, *MPLS Transport Profile Data Plane Architecture*

RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*

RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*

RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*

RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*

RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*

RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

Network Configuration Protocol (NETCONF)

RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*

RFC 6241, *Network Configuration Protocol (NETCONF)*

RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*

RFC 6243, *With-defaults Capability for NETCONF*

Open Shortest Path First (OSPF)

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*

RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*

RFC 1765, *OSPF Database Overflow*

RFC 2328, *OSPF Version 2*

RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*

RFC 3509, *Alternative Implementations of OSPF Area Border Routers*

RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart* (helper mode)

RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*

RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*

RFC 4552, *Authentication/Confidentiality for OSPFv3*

RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 5185, *OSPF Multi-Area Adjacency*

RFC 5187, *OSPFv3 Graceful Restart* (helper mode)

RFC 5243, *OSPF Database Exchange Summary List Optimization*

RFC 5250, *The OSPF Opaque LSA Option*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5340, *OSPF for IPv6*

RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*

RFC 5838, *Support of Address Families in OSPFv3*

RFC 6987, *OSPF Stub Router Advertisement*

RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*

RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenConfig

gnmi.proto, *gRPC Network Management Interface (gNMI), version 0.3.1* (Subscribe RPC)

OpenFlow

ONF *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

RFC 1989, *PPP Link Quality Monitoring*

RFC 1990, *The PPP Multilink Protocol (MP)*

RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

RFC 2153, *PPP Vendor Extensions*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 2615, *PPP over SONET/SDH*

RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points* (Gx support as it applies to wireline environment (BNG))

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*

MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*

MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*

MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*

MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*

RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*

RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*

RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*

RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*

RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*

RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*

RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*

RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*

RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*

RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*

RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*

RFC 6073, *Segmented Pseudowire*

RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*

RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*

RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*

RFC 6718, *Pseudowire Redundancy*

RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*

RFC 6870, *Pseudowire Preferential Forwarding Status bit*

RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 2869, *RADIUS Extensions*

RFC 3162, *RADIUS and IPv6*

RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*

RFC 5176, *Dynamic Authorization Extensions to RADIUS*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*

RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*

RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*

RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

Segment Routing (SR)

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*

draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*

draft-ietf-mpls-spring-lsp-ping-02, *Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane*

draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ANSI T1.105.03, *Jitter Network Interfaces*

ANSI T1.105.06, *Physical Layer Specifications*

ANSI T1.105.09, *Network Timing and Synchronization*

ITU-T G.703, *Physical/electrical characteristics of hierarchical digital interfaces*

ITU-T G.707, *Network node interface for the synchronous digital hierarchy (SDH)*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*

ITU-T G.823, *The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy*

ITU-T G.824, *The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy*

ITU-T G.825, *The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)*

ITU-T G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

ITU-T G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*

Time Division Multiplexing (TDM)

ANSI T1.403, *DS1 Metallic Interface Specification*

ANSI T1.404, *DS3 Metallic Interface Specification*

Timing

GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005*

GR-253-CORE, *SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions, issued 09/2008*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003*

ITU-T G.8261, *Timing and synchronization aspects in packet networks, issued 04/2008*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007*

ITU-T G.8264, *Distribution of timing information through packet networks, issued 10/2008*

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization, issued 10/2010*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)* (server, unauthenticated mode)

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications* (estimating the interarrival jitter)

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)

Customer Document and Product Support

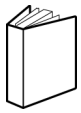

Customer Documentation

Customer Documentation Welcome Page

Technical Support

Product Support Portal

Documentation Feedback

Customer Documentation Feedback

3HE 11970 AAAB TQZZA 01