



Alcatel-Lucent 7950

EXTENSIBLE ROUTING SYSTEM | RELEASE 13.0.R4

LAYER 2 SERVICES AND EVPN GUIDE: VLL, VPLS, PBB, AND EVPN

Alcatel-Lucent Proprietary
This document contains proprietary information of Alcatel-Lucent and is not to be disclosed or used except in
accordance with applicable agreements.
Copyright 2015 © Alcatel-Lucent. All rights reserved.

All specifications, procedures, and information in this document are subject to change and revision at any time without notice. The information contained herein is believed to be accurate as of the date of publication. Alcatel-Lucent provides no warranty, express or implied, regarding its contents. Users are fully responsible for application or use of the documentation.

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

Copyright 2015 Alcatel-Lucent.

All rights reserved.

Disclaimers

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

Table of Contents

Preface	19
About This Guide	19
Audience	19
List of Technical Publications	20
Searching for Information	21
To search for specific information in this guide	21
To search for specific information in multiple documents	21
Technical Support	23
VLL Services	
In This Chapter	25
Ethernet Pipe (Epipe) Services	26
Epipe Service Overview	27
Epipe Service Pseudowire VLAN Tag Processing	28
Epipe Up Operational State Configuration Option	32
Epipe with PBB	33
Epipe over L2TPv3	34
Extension to IP VLL for Discovery of Ethernet CE IP Address	35
VLL Ethernet SAP Procedures	36
IPv6 Support on IP Interworking VLL	38
IPv6 Datapath Operation	38
IPv6 Stack Capability Signaling	40
VCCV BFD support for VLL, Spoke-SDP Termination on IES and VPRN, and VPLS Services	43
VCCV BVD Support	43
VCCV BFD Encapsulation on a Pseudowire	44
BFD Session Operation	44
Configuring VCCV BFD	45
Pseudowire Switching	47
Pseudowire Switching with Protection	48
Pseudowire Switching Behavior	50
Static-to-Dynamic Pseudowire Switching	52
Ingress VLAN Swapping	53
Ingress VLAN Translation	54
Pseudowire Redundancy	55
Dynamic Multi-Segment Pseudowire Routing	56
Overview	56
Pseudowire Routing	61
Configuring VLLs using Dynamic MS-PWs	64
Pseudowire Redundancy	67
VCCV OAM for Dynamic MS-PWs	69
VCCV-Ping on Dynamic MS-PWs	69
VCCV-Trace on Dynamic MS-PWs	70
Example Dynamic MS-PW Configuration	71
VLL Resilience with Two Destination PE Nodes	74
Master-Slave Operation	76

Table of Contents

Pseudowire SAPs	84
Epipe Using BGP-MH Site Support for Ethernet Tunnels	84
Operational Overview	86
Detailed Operation	87
BGP-MH Site Support for Ethernet Tunnels Operational-Group Model	91
BGP-MH Specifics for MH Site Support for Ethernet Tunnels	91
PW Redundancy for BGP MH Site Support for Ethernet Tunnels	91
T-LDP Status Notification Handling Rules of BGP-MH Epipes	92
Access Node Resilience Using MC-LAG and Pseudowire Redundancy	103
VLL Resilience for a Switched Pseudowire Path	105
Pseudowire Redundancy Service Models	107
Redundant VLL Service Model	107
T-LDP Status Notification Handling Rules	109
Processing Endpoint SAP Active/Standby Status Bits	109
Processing and Merging	109
BGP Virtual Private Wire Service (VPWS)	111
Single-Homed BGP VPWS	111
Dual-Homed BGP VPWS	112
BGP VPWS Pseudowire Switching	114
VLL Service Considerations	123
SDPs	123
SDP Statistics for VPLS and VLL Services	124
SAP Encapsulations and Pseudowire Types	125
QoS Policies	126
Filter Policies	126
MAC Resources	126
Configuring a VLL Service with CLI	127
Basic Configurations	128
Common Configuration Tasks	128
Configuring VLL Components	129
Creating an Epipe Service	130
Using Spoke SDP Control Words	141
Pseudowire Configuration Notes	142
Configuring Two VLL Paths Terminating on T-PE2	143
Configuring VLL Resilience	146
Configuring VLL Resilience for a Switched Pseudowire Path	147
Configuring BGP Virtual Private Wire Service (VPWS)	149
Single-Homed BGP VPWS	149
Dual-Homed BGP VPWS	151
Service Management Tasks	157
Modifying Epipe Service Parameters	158
Disabling an Epipe Service	158
Re-Enabling an Epipe Service	159
Deleting an Epipe Service	159
VLL Services Command Reference	161
Command Hierarchies	161
VLL Service Configuration Commands	183
VLL Show Commands	299

Virtual Private LAN Service

In This Chapter	357
VPLS Service Overview	359
VPLS Packet Walkthrough	360
VPLS Features	363
VPLS Enhancements	363
VPLS over MPLS	364
VPLS Service Pseudowire VLAN Tag Processing	365
VPLS MAC Learning and Packet Forwarding	369
MAC Learning Protection	369
DEI in IEEE 802.1ad	371
VPLS Using G.8031 Protected Ethernet Tunnels	372
Pseudowire Control Word	373
Table Management	374
FIB Size	374
FIB Size Alarms	374
Local and Remote Aging Timers	375
Disable MAC Aging	375
Disable MAC Learning	375
Unknown MAC Discard	375
VPLS and Rate Limiting	376
MAC Move	376
Auto-Learn MAC Protect	377
Split Horizon SAP Groups and Split Horizon Spoke SDP Groups	382
VPLS and Spanning Tree Protocol	383
Spanning Tree Operating Modes	383
Multiple Spanning Tree	385
MSTP for QinQ SAPs	387
Provider MSTP	387
Enhancements to the Spanning Tree Protocol	389
Egress Multicast Groups	392
Egress Multicast Group Provisioning	392
VPLS Redundancy	402
Spoke SDP Redundancy for Metro Interconnection	402
Spoke SDP Based Redundant Access	404
Inter-Domain VPLS Resiliency Using Multi-Chassis Endpoints	405
Support for Single Chassis Endpoint Mechanisms	409
Using B-VPLS for Increased Scalability and Reduced Convergence Times	413
MAC Flush Additions for PBB VPLS	415
VPLS Access Redundancy	418
STP-Based Redundant Access to VPLS	418
Redundant Access to VPLS Without STP	419
Object Grouping and State Monitoring	420
VPLS Applicability — Block on VPLS a Failure	420
MAC Flush Message Processing	422
Dual Homing to a VPLS Service	424
MC-Ring and VPLS	425
ACL Next-Hop for VPLS	426
SDP Statistics for VPLS and VLL Services	427
BGP Auto-Discovery for LDP VPLS	428

Table of Contents

BGP AD Overview	428
Information Model	428
FEC Element for T-LDP Signaling	430
BGP-AD and Target LDP (T-LDP) Interaction	431
SDP Usage	433
Automatic Creation of SDPs	433
Manually Provisioned SDP	433
Automatic Instantiation of Pseudowires (SDP Bindings)	434
Mixing Statically Configured and Auto-Discovered Pseudowires in a VPLS	435
Resiliency Schemes	435
BGP VPLS	436
Pseudowire Signaling Details	437
Supported VPLS Features	439
VCCV BFD Support for VPLS Services	441
BGP Multi-Homing for VPLS	442
Information Model and Required Extensions to L2VPN NLRI	443
Supported Services and Multi-Homing Objects	444
Blackhole Avoidance	445
BGP Multi-Homing for VPLS Inter-Domain Resiliency	446
Multicast-Aware VPLS	447
PIM Snooping for VPLS	447
Multicast Listener Discovery (MLD) Snooping and MAC-Based Multicast Forwarding	449
PIM and IGMP Snooping Interaction	450
VPLS Multicast-Aware High Availability Features	450
RSVP and LDP P2MP LSP for Forwarding VPLS/B-VPLS BUM and IP Multicast Packets	451
The allow-ip-int-bind VPLS Flag	452
Routed VPLS SAPs Only Supported on Standard Ethernet Ports	452
LAG Port Membership Constraints	453
Routed VPLS Feature Restrictions	454
Routed I-VPLS Feature Restrictions	454
IES IP Interface VPLS Binding and Chassis Mode Interaction	455
VPRN IP Interface VPLS Binding and Forwarding Plane Constraints	455
Route Leaking Between Routing Contexts	455
Ingress LAG and FP1 to Routed VPLS Discards	456
IPv4 Multicast Routing Support	457
BGP Auto Discovery (BGP-AD) for Routed VPLS Support	460
Routed VPLS Caveats	461
VPLS SAP Ingress IP Filter Override	461
IP Interface Defined Egress QoS Reclassification	461
Remarking for VPLS and Routed Packets	462
IPv4 Multicast Routing	462
Routed VPLS Supported Routing Related Protocols	462
Spanning Tree and Split Horizon	463
VPLS Service Considerations	464
SAP Encapsulations	464
VLAN Processing	464
Ingress VLAN Swapping	465
Service Auto-Discovery using Multiple VLAN Registration Protocol (MVRP)	466
Configure the MVRP Infrastructure using an M-VPLS Context	467
Instantiate Related VLAN FIBs and Trunks in MVRP Scope	467

MVRP Activation of Service Connectivity	470
MVRP Control Plane	473
STP-MVRP Interaction	473
VPLS E-Tree Services	476
VPLS E-Tree Services Overview	476
Leaf-ac and Root-ac SAPs	477
Leaf-ac and Root-ac SDP Binds	478
Root-leaf-tag SAPs	478
Root-leaf-tag SDP Binds	479
Interaction between VPLS E-Tree Services and Other Features	480
Configuring a VPLS Service with CLI	483
Basic Configuration	484
Common Configuration Tasks	486
Configuring VPLS Components	487
Configuring Egress Multicast Groups	488
Creating a VPLS Service	489
Enabling Multiple MAC Registration Protocol (MMRP)	490
Configuring GSMP Parameters	499
Configuring a VPLS SAP	500
Configuring SDP Bindings	510
Configuring Overrides on Service SAPs	511
Configuring VPLS Redundancy	523
Creating a Management VPLS for SAP Protection	523
Creating a Management VPLS for Spoke SDP Protection	524
Configuring Load Balancing with Management VPLS	527
Configuring Selective MAC Flush	532
Configuring Multi-Chassis Endpoints	533
Configuring VPLS E-Tree Services	538
Service Management Tasks	539
Modifying VPLS Service Parameters	539
Modifying Management VPLS Parameters	540
Deleting a Management VPLS	540
Disabling a Management VPLS	541
Deleting a VPLS Service	542
Disabling a VPLS Service	542
Re-Enabling a VPLS Service	543
VPLS Services Command Reference	545
Command Hierarchies	545
VPLS Service Configuration Commands	581
VPLS Show Commands	791

IEEE 802.1ah Provider Backbone Bridging

In This Chapter	887
IEEE 802.1ah Provider Backbone Bridging (PBB) Overview	888
PBB Features	889
Integrated PBB-VPLS Solution	889
PBB Technology	891
PBB Mapping to Existing VPLS Configurations	892
SAP and SDP Support	894
PBB B-VPLS	894

Table of Contents

PBB I-VPLS	894
PBB Packet Walkthrough	896
PBB Control Planes	897
Shortest Path Bridging MAC Mode (SPBM)	898
Flooding and Learning Versus Link State	898
SPB for B-VPLS	899
Control B-VPLS and User B-VPLS	899
Shortest Path and Single Tree	902
Data Path and Forwarding	905
SPB Ethernet OAM	905
SPB Levels	906
SPBM to Non-SPBM Interworking	907
Static MACs and Static ISIDs	907
Epipe Static Configuration	907
SPBM ISID Policies	909
ISID Policy Control	911
Static ISID Advertisement	911
I-VPLS for Unicast Service	911
Default Behaviors	912
Example Network Configuration	913
Sample Configuration for Dut-A	914
IEEE 802.1ak MMRP for Service Aggregation and Zero Touch Provisioning	921
MMRP Support Over B-VPLS SAPs and SDPs	923
I-VPLS Changes and Related MMRP Behavior	923
Limiting the Number of MMRP Entries on a Per B-VPLS Basis	923
Optimization for Improved Convergence Time	924
Controlling MRP Scope using MRP Policies	924
PBB and BGP-AD	928
PBB ELINE Service	928
Non-Redundant PBB Epipe Spoke Termination	928
Support Service and Solution Combinations	929
Periodic MAC Notification	930
MAC Flush	931
PBB Resiliency for B-VPLS Over Pseudowire Infrastructure	931
Access Multi-Homing for Native PBB (B-VPLS over SAP Infrastructure)	936
Solution Description for I-VPLS Over Native PBB Core	937
Solution Description for PBB Epipe over G.8031 Ethernet Tunnels	940
BGP Multi-homing for I-VPLS	944
Access Multi-Homing over MPLS for PBB Epipes	945
PBB and IGMP/MLD Snooping	948
PBB QoS	949
Transparency of Customer QoS Indication through PBB Backbone	950
Egress B-SAP per ISID Shaping	956
B-SAP Egress ISID Shaping Configuration	956
Provisioning Model	958
Egress Queue Scheduling	960
B-SAP per-ISID Shaping Configuration Example	962
PBB OAM	965
Mirroring	966
OAM Commands	966

CFM Support	966
Configuration Examples	967
PBB using G.8031 Protected Ethernet Tunnels	967
MC-LAG Multihoming for Native PBB	970
Access Multi-Homing over MPLS for PBB Epipes	972
PBB Command Reference	975
Command Hierarchies	975
PBB Service Commands	981
PBB Show Commands	1005
PBB Clear Commands	1024
PBB Debug Commands	1026

Ethernet Virtual Private Networks (EVPNs)

In This Chapter	1029
Overview	1031
EVPN for VXLAN Tunnels in a Layer-2 DC GW (EVPN-VXLAN)	1032
EVPN for VXLAN Tunnels in a Layer-2 DC with Integrated Routing Bridging Connectivity on the DC GW	1034
EVPN for VXLAN Tunnels in a Layer 3 DC with Integrated Routing Bridging Connectivity among VPRNs	1035
EVPN for VXLAN Tunnels in a Layer 3 DC with EVPN-Tunnel Connectivity among VPRNs	1037
EVPN for MPLS Tunnels in ELAN Services	1039
EVPN for PBB over MPLS Tunnels (PBB-EVPN)	1041
VXLAN	1042
VXLAN ECMP and LAG	1045
VXLAN VPLS Tag Handling	1045
VXLAN MTU Considerations	1045
VXLAN QoS	1046
VXLAN Ping	1046
IGMP-Snooping on VXLAN	1050
BGP-EVPN Control Plane for VXLAN Overlay Tunnels	1053
EVPN for VXLAN in VPLS Services	1057
Resiliency and BGP Multi-Homing	1059
Use of bgp-evpn, bgp-ad, and Sites in the Same VPLS Service	1059
Use of the unknown-mac-route	1061
EVPN for VXLAN in R-VPLS Services	1062
EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes	1064
EVPN for VXLAN in EVPN Tunnel R-VPLS Services	1068
DC GW integration with the Nuage Virtual Services Directory (VSD)	1073
XMPP Interface on the DC GW	1074
Overview of the Static-Dynamic VSD Integration Model	1078
VSD-Domains and Association to Static-Dynamic Services	1080
VSD-Domain Type L2-DOMAIN	1081
VSD-Domain Type L2-DOMAIN-IRB	1083
VSD-Domain Type VRF-GRE	1083
VSD-Domain Type VRF-VXLAN	1083
Fully-Dynamic VSD Integration Model	1086
Python Script Implementation Details	1089
BGP-EVPN Control Plane for MPLS Tunnels	1095
EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS)	1100

Table of Contents

EVPN and VPLS Integration	1103
Auto-Derived Route-Distinguisher (RD) in Services with Multiple BGP Families	1107
EVPN Multi-Homing in VPLS Services	1108
EVPN All-Active Multi-Homing	1108
All-Active Multi-Homing Service Model	1110
ES Discovery and DF Election Procedures	1112
Aliasing	1119
Network Failures and Convergence for All-Active Multi-Homing	1122
Logical Failures on Ethernet Segments and Black-Holes	1123
Transient Issues Due to MAC Route Delays	1124
EVPN Single-Active Multi-Homing	1125
Single-Active Multi-Homing Service Model	1126
ES and DF Election Procedures	1128
Backup PE Function	1130
Network Failures and Convergence for Single-Active Multi-Homing	1132
BGP-EVPN Control Plane for PBB-EVPN	1134
PBB-EVPN for I-VPLS and PBB Epipe Services	1136
Flood Containment for I-VPLS Services	1139
PBB-EVPN and PBB-VPLS Integration	1141
PBB-EVPN Multi-Homing in I-VPLS and PBB Epipe Services	1142
System BMAC Assignment in PBB-EVPN	1142
PBB-EVPN all-active multi-homing service model	1143
Network failures and convergence for all-active multi-homing	1148
PBB-EVPN Single-Active Multi-Homing Service Model	1150
Network Failures and Convergence for Single-Active Multihoming	1152
PBB-Epipes and EVPN Multi-Homing	1153
ARP/ND Snooping and Proxy Support	1155
Proxy-ARP/ND Periodic Refresh, Unsolicited Refresh and Confirm-Messages	1159
Proxy-ND and the Router Flag in Neighbor Advertisement messages	1160
BGP-EVPN MAC-Mobility	1161
BGP-EVPN MAC-Duplication	1162
Conditional Static MAC and Protection	1164
CFM Interaction with EVPN Services	1165
DC GW Policy Based Forwarding/Routing to an EVPN ESI (Ethernet Segment Identifier)	1167
Policy Based Forwarding in VPLS Services for Nuage Service Chaining Integration in L2-Domains	1167
Policy Based Routing in VPRN Services for Nuage Service Chaining Integration in L2-DOMAIN-IRB Domains	1171
BGP and EVPN Route Selection for EVPN Routes	1175
Interaction of EVPN-VXLAN and EVPN-MPLS with Existing VPLS Features	1177
Interaction of PBB-EVPN with Existing VPLS Features	1179
Interaction of EVPN-VXLAN with Existing VPRN Features	1180
Routing Policies for BGP EVPN IP Prefixes	1181
Configuring an EVPN Service with CLI	1183
EVPN-VXLAN Configuration Examples	1184
Layer 2 PE Example	1184
EVPN for VXLAN in R-VPLS Services Example	1186
EVPN for VXLAN in EVPN Tunnel R-VPLS Services Example	1188
EVPN for VXLAN in R-VPLS Services with IPv6 interfaces and prefixes Example	1189
EVPN-MPLS Configuration Examples	1190

EVPN All-active Multi-homing Example	1190
EVPN Single-active Multi-homing Example	1193
PBB-EVPN Configuration Examples	1195
PBB-EVPN All-active Multi-homing Example	1195
PBB-EVPN Single-active Multi-homing Example	1198
EVPN Command Reference	1201
Command Hierarchies	1201
Show Commands	1241
Clear Commands	1261
Tools Commands	1262
Debug Commands	1268
Common CLI Command Descriptions	
In This Chapter	1271
Common Service Commands	1272
Standards and Protocol Support	1275

Table of Contents

List of Tables

Table 1: List of Technical Publications.	20
--	----

VLL Services

Table 2: Epipe Spoke SDP VLAN Tag Processing: Ingress	29
Table 3: Epipe Spoke SDP VLAN Tag Processing: Egress.	30
Table 4: SAP MEP Signaling	89
Table 5: Supported SAP Types	132
Table 6: Default QinQ and TopQ SAP Dot1P Evaluation	265
Table 7: Bottom Position QinQ and TopQ SAP Dot1P Evaluation	266

Virtual Private LAN Service

Table 8: VPLS Mesh and Spoke SDP VLAN Tag Processing: Ingress	366
Table 9: VPLS Mesh and Spoke SDP VLAN Tag Processing: Egress	367
Table 10: SAP Chain Creation.	399
Table 11: MSTP and MVRP Interaction Table.	473
Table 12: Spoke SDP BPDU Encapsulation States.	520
Table 13: Default QinQ and TopQ SAP Dot1P Evaluation	715
Table 14: Bottom Position QinQ and TopQ SAP Dot1P Evaluation	716

IEEE 802.1ah Provider Backbone Bridging

Table 15: B-VPLS Control Planes	898
Table 16: SPB Ethernet OAM Operation Summary.	906
Table 17: SPBM ISID Policies Table	910

Ethernet Virtual Private Networks (EVPNs)

Table 18: EVPN Routes and Usage	1095
Table 19: Proxy-arp Entry combinations	1158

List of Tables

List of Figures

VLL Services

Figure 1: Epipe/VLL Service	27
Figure 2: L2TPv3 SDP Illustration	34
Figure 3: IP Interworking VLL Datapath	35
Figure 4: Data Path for Ethernet CE to PPP Attached CE	39
Figure 5: Pseudowire Service Switching Node	47
Figure 6: VLL Resilience with Pseudowire Redundancy and Switching	48
Figure 7: Ingress VLAN Swapping	53
Figure 8: Ingress VLAN Translation	54
Figure 9: Dynamic MS-PW Overview	56
Figure 10: MS-PW Addressing using FEC129 All Type 2	57
Figure 11: Advertisement of PE Addresses by PW Routing	58
Figure 12: Signaling of Dynamic MS-PWs using T-LDP	59
Figure 13: Mapping of All to SAP	59
Figure 14: VLL Using Dynamic MS-PWs, Inter-AS Scenario	60
Figure 15: Pseudowire Redundancy	67
Figure 16: Dynamic MS-PW Example	71
Figure 17: VLL Resilience	74
Figure 18: Master-Slave Pseudowire Redundancy	77
Figure 19: Example of SAP OAM Interaction with Master-Slave Pseudowire Redundancy	79
Figure 20: VLL Resilience	81
Figure 21: VLL Resilience with Pseudowire Switching	83
Figure 22: BGP-MH Site Support for Ethernet Tunnels	85
Figure 23: G.8031 for Slave Operation	86
Figure 24: Full Redundancy G.8031 Epipe & BGP-MH	88
Figure 25: Sample Topology Full Redundancy	94
Figure 26: Access Node Resilience	103
Figure 27: VLL Resilience with Pseudowire Redundancy and Switching	105
Figure 28: Redundant VLL Endpoint Objects	107
Figure 29: Single-Homed BGP-VPWS Example	111
Figure 30: Dual-Homed BGP VPWS with Single Pseudowire	112
Figure 31: Dual-homed BGP VPWS with Active/Standby Pseudowires	113
Figure 32: BGP VPWS Update Extended Community Format	115
Figure 33: BGP VPWS NLRI	117
Figure 34: BGP VPWS NLRI TLV Extension Format	117
Figure 35: Circuit Status Vector TLV Type	117
Figure 36: SDP Statistics for VPLS and VLL Services	124
Figure 37: SDPs — Uni-Directional Tunnels	138
Figure 38: VLL Resilience	146
Figure 39: VLL Resilience with Pseudowire Switching	147
Figure 40: Single-Homed BGP VPWS Configuration Example	149
Figure 41: Example of Dual-Homed BGP VPWS with Single Pseudowire	151
Figure 42: Example of Dual-homed BGP VPWS with Active/Standby Pseudowires	154

List of Figures

Virtual Private LAN Service

Figure 43: VPLS Service Architecture	360
Figure 44: Access Port Ingress Packet Format and Lookup	360
Figure 45: Network Port Egress Packet Format and Flooding	361
Figure 46: Access Port Egress Packet Format and Lookup	362
Figure 47: MAC Learning Protection	370
Figure 48: DE Bit in the 802.1ad S-TAG	371
Figure 49: Auto-Learn-Mac-Protect Operation	379
Figure 50: Auto-Learn-Mac-Protect Example	381
Figure 51: Access Resiliency	386
Figure 52: HVPLS with Spoke Redundancy	403
Figure 53: HVPLS Resiliency Based on AS Pseudowires	405
Figure 54: Multi-Chassis Pseudowire Endpoint for VPLS	406
Figure 55: MC-EP in Passive Mode	409
Figure 56: MAC Flush in the MC-EP Solution	411
Figure 57: MC-EP with B-VPLS	414
Figure 58: MC-EP with B-VPLS Failure Scenario	415
Figure 59: MC-EP with B-VPLS Mac Flush Solution	416
Figure 60: Dual Homed MTU-s in Two-Tier Hierarchy H-VPLS	418
Figure 61: Dual Homed CE Connection to VPLS	424
Figure 62: Application 1 Diagram	426
Figure 63: SDP Statistics for VPLS and VLL Services	427
Figure 64: BGP AD NLRI versus IP VPN NLRI	429
Figure 65: Generalized Pseudowire-ID FEC Element	430
Figure 66: BGP-AD and T-LDP Interaction	432
Figure 67: BGP VPLS Solution	436
Figure 68: BGP Multi-Homing for VPLS	442
Figure 69: BGP MH-NLRI for VPLS Multi-Homing	443
Figure 70: BGP MH Used in an HVPLS Topology	446
Figure 71: IPv4 Multicast with a Router VPLS service	459
Figure 72: Ingress VLAN Swapping	465
Figure 73: Infrastructure for MVRP Exchanges	466
Figure 74: Service Instantiation with MVRP - QinQ to PBB Example	470
Figure 75: E-Tree Service	477
Figure 76: Mapping PE Model to 7x50 VPLS Service	478
Figure 77: Leaf and Root Tagging Dot1q	479
Figure 78: Leaf and Root Tagging PW	480
Figure 79: SDPs — Uni-Directional Tunnels	512
Figure 80: Example Configuration for Protected VPLS Spoke SDP	525
Figure 81: Example Configuration for Load Balancing Across Two Protected VPLS Spoke SDPs	527

IEEE 802.1ah Provider Backbone Bridging

Figure 82: Large HVPLS Deployment	889
Figure 83: Large PBB-VPLS Deployment	890
Figure 84: QinQ Payload in Provider Header Example	891
Figure 85: PBB Mapping to VPLS Constructs	892
Figure 86: PBB Packet Walkthrough	896
Figure 87: Control and User B-VPLS with FIDs	900
Figure 88: Sample Partial Mesh network	902

Figure 89: Unicast Paths for Low-path-id and High-path-id	903
Figure 90: Multicast Paths for Low-path-id and High-path-id	904
Figure 91: Static MACs Example	908
Figure 92: Static ISIDs Example	909
Figure 93: ISID Policy Example	912
Figure 94: Sample Network	913
Figure 95: Customer Services Transported in 1 B-VPLS (M:1 Model)	921
Figure 96: Flood Containment Requirement in M:1 Model	922
Figure 97: Inter-Domain Topology	925
Figure 98: Limiting the Scope of MMRP Advertisements	925
Figure 99: TCN Triggered PBB Flush-ALI-But-Mine Procedure	934
Figure 100: Access Dual-Homing into PBB BEBs - Topology View	936
Figure 101: PBB Active Topology and Access Multi-Homing	937
Figure 102: Access Multi-Homing - Link Failure	939
Figure 103: Access Multi-Homing Solution for PBB Epipe	940
Figure 104: Access Dual-Homing for PBB ELINE - BEB Failure	941
Figure 105: Solution for Access Dual-Homing with Local Switching for PBB Eline/Epipe	942
Figure 106: Active/Standby PW into PBB Epipes	945
Figure 107: PCP, DE Bits Transparency in PBB	950
Figure 108: Egress Queue Scheduling	960
Figure 109: PBB OAM View for MPLS Infrastructure	965

Ethernet Virtual Private Networks (EVPNs)

Figure 110: Layer-2 DC PE with VPLS to the WAN	1032
Figure 111: GW IRB on the DC PE for an L2 EVPN/VXLAN DC	1034
Figure 112: GW IRB on the DC PE for an L3 EVPN/VXLAN DC	1035
Figure 113: EVPN-Tunnel GW IRB on the DC PE for an L3 EVPN/VXLAN DC	1037
Figure 114: EVPN for MPLS in VPLS Services	1039
Figure 115: EVPN for PBB over MPLS	1041
Figure 116: VXLAN Frame Format	1043
Figure 117: EVPN-VXLAN Required Routes and Communities	1053
Figure 118: EVPN Route-Type 5	1055
Figure 119: Basic XMPP Architecture	1075
Figure 120: WAN Services Attachment Workflow	1078
Figure 121: Fully-Dynamic VSD Integration Model Workflow	1086
Figure 122: EVPN Routes Type 1 and 4	1097
Figure 123: EVPN-VPLS Integration	1105
Figure 124: DF Election	1109
Figure 125: Split-Horizon	1109
Figure 126: Aliasing	1110
Figure 127: ES Discovery and DF Election	1113
Figure 128: All-Active Multi-Homing ES Failure	1122
Figure 129: Black-hole Caused by SAP/SVC Shutdown	1123
Figure 130: Transient Issues Caused by “slow” MAC Learning	1124
Figure 131: Backup PE	1126
Figure 132: Single-Active Multi-Homing ES Failure	1132
Figure 133: Single-Active Multi-Homing SAP/SDP/Service Shutdown	1133
Figure 134: PBB-EVPN for I-VPLS and PFF Epipe Services	1136
Figure 135: PBB-EVPN and I-VPLS Flooding Containment	1139
Figure 136: PBB-EVPN All-Active Multi-Homing	1143

List of Figures

Figure 137: Source-Bmac Versus Es-Bmac CMAC Flushing	1151
Figure 138: PBB-EVPN MH in a Three-Node Scenario	1153
Figure 139: PBB-EVPN MH in a Two-Node Scenario	1154
Figure 140: Proxy-ARP Example Usage in an EVNP Network	1155
Figure 141: PBF to ESI Function	1168
Figure 142: PBR to ESI Function	1171
Figure 143: IP-VPN Import and EVPN Export BGP Workflow	1181
Figure 144: I-EVPN Import and I-VPN Export BGP Workflow	1182

Preface

About This Guide

This guide describes Layer 2 service and EVPN functionality provided by Alcatel-Lucent's 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.

Audience

This guide is intended for network administrators who are responsible for configuring the 7950 XRS routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Concepts described in this guide include the following:

- Virtual Leased Lines (VLL)
- Virtual Private LAN Service (VPLS)
- Provider Backbone Bridging (PBB)
- Ethernet VPN (EVPN)

List of Technical Publications

The 7950 XRS routers documentation set is composed of the following guides:

Table 1: List of Technical Publications

Guide	Description
7950 XRS Basic System Configuration Guide	This guide describes basic system configurations and operations.
7950 XRS System Management Guide	This guide describes system security and access configurations as well as event logging and accounting logs.
7950 XRS Interface Configuration Guide	This guide describes XMA Control Module (XCM), XRS Media Adaptor (XMA), port and Link Aggregation Group (LAG) provisioning.
7950 XRS Router Configuration Guide	This guide describes logical IP routing interfaces and associated attributes such as an IP address, as well as IP and MAC-based filtering, and VRRP and Cflowd.
7950 XRS Routing Protocols Guide	This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.
7950 XRS MPLS Guide	This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
7950 XRS Services Guide	This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
7950 XRS Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN	This guide describes Virtual Leased Lines (VLL), Virtual Private LAN Service (VPLS), Provider Backbone Bridging (PBB), and Ethernet VPN (EVPN).
7950 XRS Layer 3 Services Guide: Internet Enhanced Services and Virtual Private Routed Network Services	This guide describes Internet Enhanced Services (IES) and Virtual Private Routed Network (VPRN) services.

Table 1: List of Technical Publications

Guide	Description
7950 XRS OAM and Diagnostics Guide	This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
7950 XRS Quality of Service Guide	This guide describes how to configure Quality of Service (QoS) policy management.

Searching for Information

You can use Adobe Reader, Release 6.0 or later, to search one or more PDF files for a term.

To search for specific information in this guide

1. From the Adobe Reader main menu, choose Edit > Search or Advanced Search. The Search panel opens.
2. Click on the In the current document radio button.
3. Enter the term to search for.
4. Select the following search criteria, if required:
 - Whole words only
 - Case-Sensitive
 - Include Bookmarks
 - Include Comments
5. Click on the Search button. Adobe Reader displays the search results.

You can expand the entries by clicking on the + symbol.

To search for specific information in multiple documents

Note: The PDF files that you search must be in the same folder.

1. From the Adobe Reader main menu, choose Edit > Search or Advanced Search. The Search panel opens.

2. Click on the All PDF Documents in radio button.
3. Choose the folder in which to search using the drop-down menu.
4. Enter the term to search for.
5. Select the following search criteria, if required:
 - Whole words only
 - Case-Sensitive
 - Include Bookmarks
 - Include Comments
6. Click on the Search button. Adobe Reader displays the search results.

You can expand the entries for each file by clicking on the + symbol.

Technical Support

If you purchased a service agreement for your 7950 XRS router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, follow this link to contact an Alcatel-Lucent support representative and to access product manuals and documentation updates:

<https://support2.alcatel-lucent.com/portal/olcsHome.do>

VLL Services

In This Chapter

This chapter provides information about Virtual Leased Line (VLL) services and implementation notes.

Topics in this chapter include:

- [Ethernet Pipe \(Epipe\) Services on page 26](#)
- [VCCV BFD support for VLL, Spoke-SDP Termination on IES and VPRN, and VPLS Services on page 43](#)
- [Pseudowire Switching on page 47](#)
- [Pseudowire Redundancy on page 55](#)
- [Dynamic Multi-Segment Pseudowire Routing on page 56](#)
- [Pseudowire SAPs on page 84](#)
- [Epipe Using BGP-MH Site Support for Ethernet Tunnels on page 84](#)
- [BGP Virtual Private Wire Service \(VPWS\) on page 111](#)

Ethernet Pipe (Epipe) Services

This section provides information about the Epipe service and implementation notes.

Topics in this section include:

- [Epipe Service Overview on page 27](#)
 - [SAP Encapsulations and Pseudowire Types on page 125](#)
 - [QoS Policies on page 126](#)
 - [Filter Policies on page 126](#)
 - [MAC Resources on page 126](#)
- [Basic Configurations on page 128](#)
- [Common Configuration Tasks on page 128](#)
 - [Configuring VLL Components on page 129](#)
 - [Creating an Epipe Service on page 130](#)
- [Service Management Tasks on page 157](#)

Epipse Service Overview

An Epipse service is Alcatel-Lucent's implementations 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 Pseudo-wire Draft (draft-so-pwe3-ethernet-00.txt).

An Epipse service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider's IP, MPLS or PBB VPLS network. An Epipse service is completely transparent to the customer's data and protocols. The , , and 7950 XRS Epipse service does not perform any MAC learning. A local Epipse service consists of two SAPs on the same node, whereas a distributed Epipse service consists of two SAPs on different nodes. SDPs are not used in local Epipse services.

Each SAP configuration includes a specific port/channel on which service traffic enters the , , or 7950 XRS 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, then a unique encapsulation value (ID) must be specified.

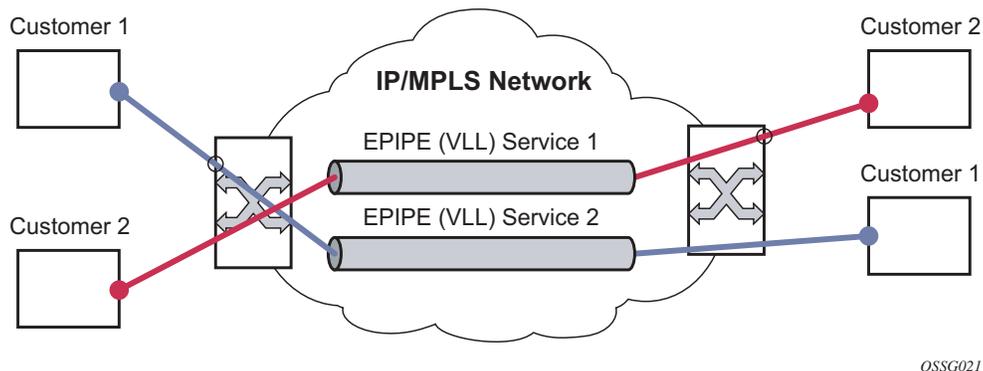


Figure 1: Epipse/VLL Service

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 2](#) and [Table 3](#) for configuration details. 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 (see below for 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 configurations, thus 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 2 and Table 3 describe the VLAN tag processing with respect to the zero, one and two VLAN tag configuration described above for the VLAN identifiers, Ether type, 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 2: 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
Ether type (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 3: 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	<ul style="list-style-type: none"> • the vlan-vc-tag value configured in pw-template or under the spoke SDP or • taken from the inner tag received on a QinQ SAP or QinQ spoke SDP or • taken from the VLAN tag received on a dot1q SAP or 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 spoke SDP 	<p>Both inner and outer VLAN tag:</p> <ul style="list-style-type: none"> • the vlan-vc-tag value configured in pw-template or under the spoke SDP or • taken from the inner tag received on a QinQ SAP or QinQ spoke SDP or • taken from the VLAN tag received on a dot1q SAP or 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 spoke SDP

Table 3: Epipe Spoke SDP VLAN Tag Processing: Egress (Continued)

Egress (sent on mesh or spoke SDP)	Zero VLAN tags	One VLAN tag	Two VLAN tags
Ether type (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)
Egress QoS (dot1p/DE) (set in VLAN tags)	N/A	<p>Taken from the inner most ingress service delimiting tag:</p> <ul style="list-style-type: none"> the inner tag received on a QinQ SAP or QinQ spoke SDP or taken from the VLAN tag received on a dot1q SAP or 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 spoke SDP. <p>Note that neither the inner nor outer dot1p/DE values can be explicitly set.</p>	<p>Both inner and outer dot1p/DE:</p> <p>Taken from the innermost ingress service delimiting tag:</p> <ul style="list-style-type: none"> the inner tag received on a QinQ SAP or QinQ spoke SDP or taken from the VLAN tag received on a dot1q SAP or 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 spoke SDP. <p>Note that neither the inner nor outer dot1p/DE values can be explicitly set.</p>

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

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 does have the ability to configure a single SAP within an Epipe not to affect the operational state of that Epipe using the optional command **ignore-oper-state**. Within an Epipe, if a SAP that includes this optional command becomes operationally down state, the operational state of the Epipe will not transition to down. The operational state of the Epipe will remain up. This does not change the fact that the SAP is down and no traffic will transit an operationally down SAP. Removing and adding this command on the fly will evaluate the service's operational state based on the SAPs and the addition or deletion of this command.

Service OAM (SOAM) designers may consider using this command if an 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 UP, MEP configured on the failed SAP the service must be operationally up. Consider the case where an 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 enter an operational state 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 which has no ports.

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 and 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 then 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.

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.

As 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 and (mixed mode) and 7950 XRS platforms, in mode D with FP2+ (FP3 recommended).

ETH-CFM is supported for OAM services.

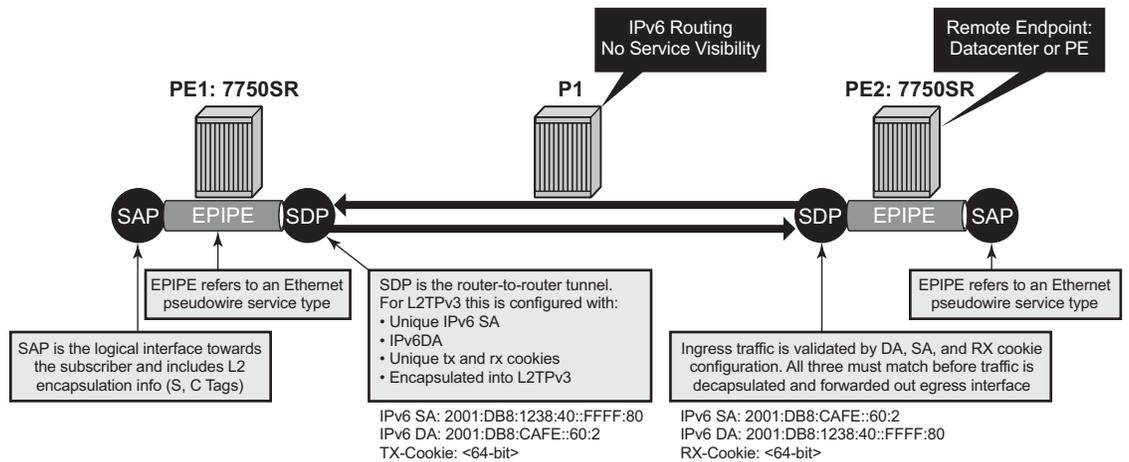


Figure 2: L2TPv3 SDP Illustration

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

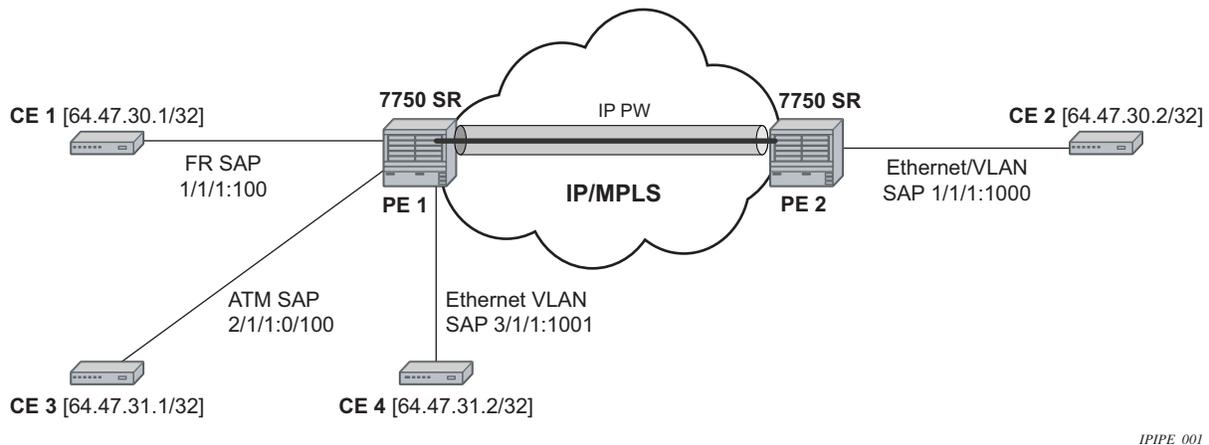


Figure 3: IP Interworking VLL Datapath

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 user through the mac-refresh CLI command.

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.
- 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 IP address ARPed for.
- The LDP address list TLV to signal 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.

VLL Ethernet SAP Procedures

The operator can enable the following CE address discovery procedures 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.
- 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 IP address being ARPed.
- The , , 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 IP address being ARPed 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, it not checks it against the IP address being ARPed for 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 makes use of 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 must be used.
- If the remote PE included the address list TLV in the received label mapping message, the local updates the remote PE node with the most current IP address of the Ethernet CE using a T-LDP notification message with status TLV status code is 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 enters manually the **ce-address** parameter under SAP and spoke SDP.

IPv6 Support on IP Interworking VLL

The , , 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.

IPv6 Datapath Operation

The IPv6 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 7x50 PE device intercepts ICMPv6 Neighbor Discovery (RFC 2461) packets, whether received over the SAP or over the pseudowire, inspects them to learn IPv6 interface addresses and CE link-layer addresses, and modifies these packets as required according to the SAP type, and then forwards them towards the original destination. The 7x50 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 7x50 PE device learns the IPv6 interface addresses for its directly-attached CE and another IPv6 interface addresses for the far-end CE. The 7x50 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 7x50 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 given MAC address, the currently cached address is overwritten with the most recent value.

[Figure 4](#) illustrates the data path operation for IPv6 on an IP Interworking VLL between the Ethernet and PPP (IPv6CP) SAPs.

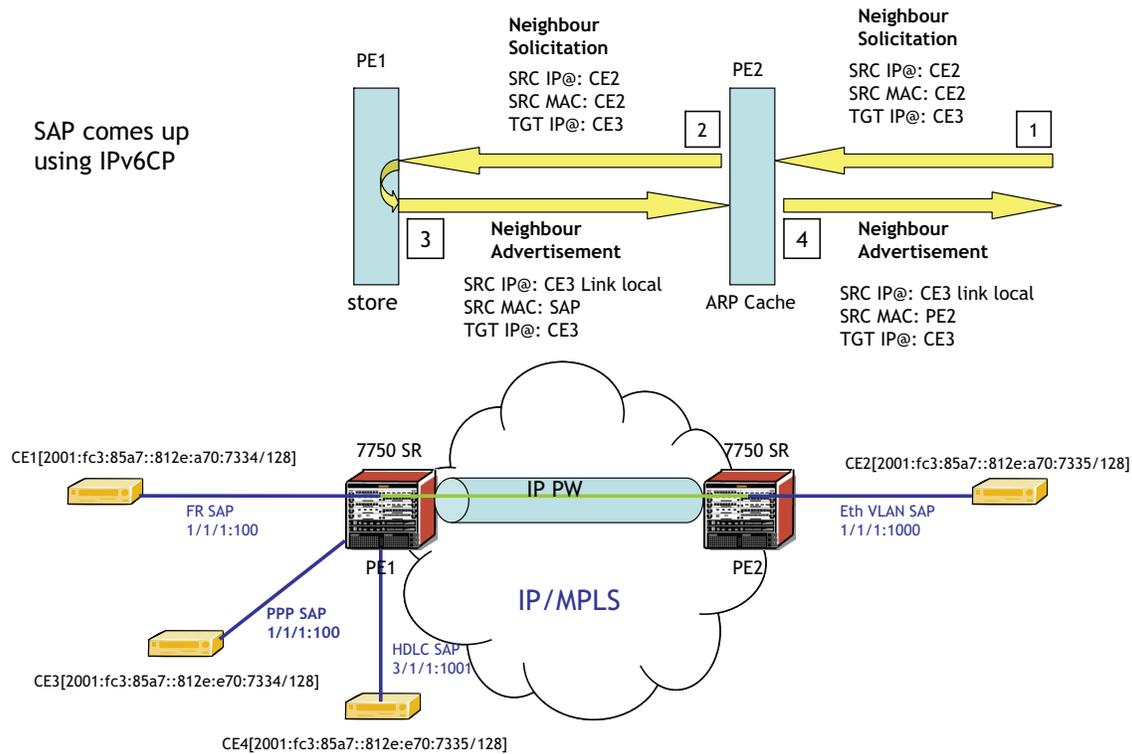


Figure 4: Data Path for Ethernet CE to PPP Attached CE

With reference to neighbor discovery between Ethernet and PPP CEs in [Figure 4](#), the steps are as follows:

1. Ethernet attached CE2 sends a Neighbor Solicitation message towards 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 towards 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.

IPv6 Stack Capability Signaling

The 7x50 supports 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. In which case, 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 7x50 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. Thus 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 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 7x50, 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 this release, 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 7x50 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 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 7x50 with `ce-address-discovery ipv6` enabled withdraws its pseudowire label with the LDP status code "IP Address type mismatch".

If a 7x50 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”.

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

Topics include:

- [VCCV BVD Support on page 43](#)
- [VCCV BFD Encapsulation on a Pseudowire on page 44](#)
- [BFD Session Operation on page 44](#)
- [Configuring VCCV BFD on page 45](#)

VCCV BVD 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, irrespective 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-SPDs 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 signalling 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.

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 7x50, 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.

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. In this sense, 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. Furthermore, mapping the BFD state to SAP OAM is not supported.

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

Note that 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.

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 can not 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.

Note that if the above BFD timer values are changed in a given 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 once the template context has been entered. However, a value will still be stored temporarily in the template-module until the **commit** command is issued. Once 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 (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), a named BFD template is configured on the spoke-SDP using the command **config>service>epipe | cpipe | apipe | fpipe | ipipe>spoke-sdp>bfd-template name** and then enabled using the command **config>service>epipe | cpipe | apipe | fpipe | ipipe>spoke-sdp>bfd-enable**.

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, then a pw-path-id can not 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 command **config>service>ies | vprn>interface>spoke-sdp>bfd-template name** and then enabled using the command **config>service>ies | vprn>interface>spoke-sdp>bfd-enable**.

Configuring VCCV BFD

For H-VPLS where the PW-Template does not apply (i.e LDP-VPLS spoke and mesh-sdps that use the Pwid FEC(FEC128) the bfd template is configured using `config>service>vpls>spoke>sdp>bfd-name name` or `config>service>vpls>mesh-sdp>bfd-name name`. VCCV BFD is then enabled with the `bfd-enable` command under the VPLS spoke-sdp or mesh-sdp context.

Pseudowires where the pw-template does apply and that support VCCV BFD are as follows:

- BGP-AD, which is signaled using the Generalised pseudowire ID FEC (FEC129) with 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 command **`config>service>epipe>bgp>pw-template-binding>bfd-template name`** and then enabled using the command **`config>service>epipe>bgp>pw-template-binding>bfd-enable`**.

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. In addition, 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 signalling 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 signalling 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 towards 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 to 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 [Figure 5](#), the user configures a regular Epipe VLL service PE1 and PE2. These services consist each of a SAP and a spoke SPD. However, the target destination of the SDP is actually not the remote PE but the pseudowire switching node. In addition, the user configures an Epipe VLL service on the pseudowire switching node using the two SDPs.

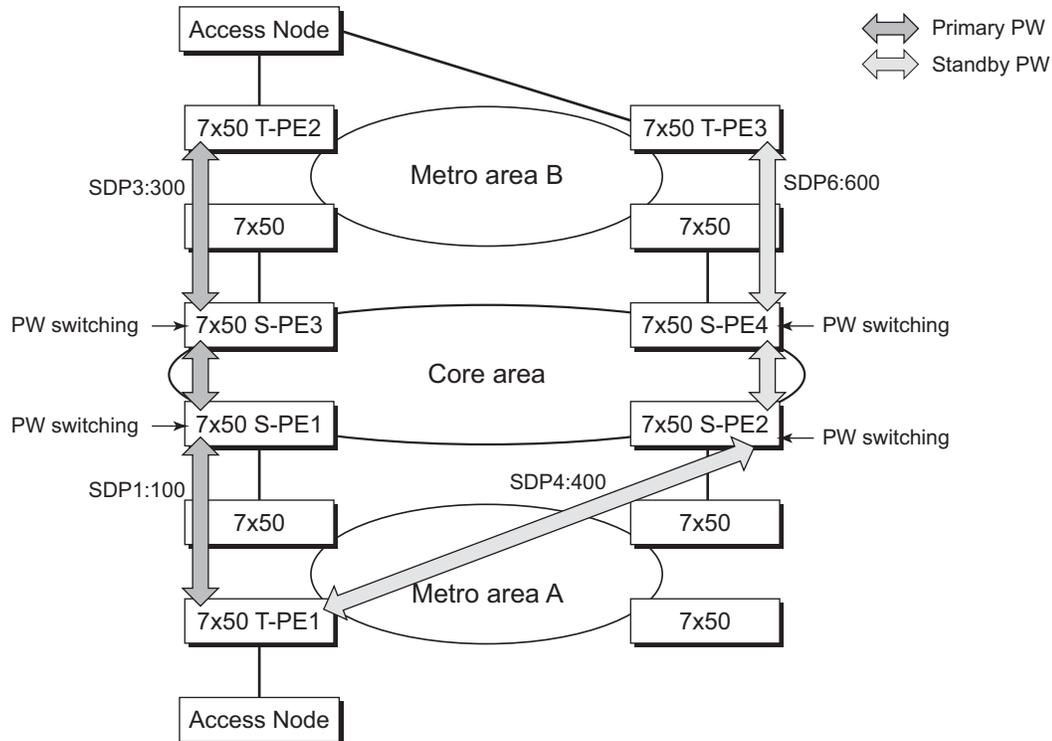
```
|7950 SR PE1 (Epipe) |---sdp 2:10---|7950 SR PW SW (Epipe) |---sdp 7:15---|7950 SR PE2 (Epipe) |
```

Figure 5: Pseudowire Service Switching Node

Configuration examples can be found in [Configuring Two VLL Paths Terminating on T-PE2 on page 143](#).

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 6](#) illustrates the use of pseudowire redundancy to provide a scalable and resilient VLL service across multiple IGP areas in a provider network.



OSSG114

Figure 6: VLL Resilience with Pseudowire Redundancy and Switching

In the network in [Figure 6](#), 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. 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, it 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 effect the spoke SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

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. Secondly, it helps in loop detection of the T-LDP signaling messages where a switching point receives back a label mapping message it already relayed. Finally, it can be inserted in pseudowire status messages when they are sent from a pseudowire switching node towards 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 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, it means the message was originated by a T-PE node and the S-PE should process and pass the message without changes except for the VCID value in the FEC TLV.

Pseudowire Switching Behavior

In the network in [Figure 6](#), 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, it 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.

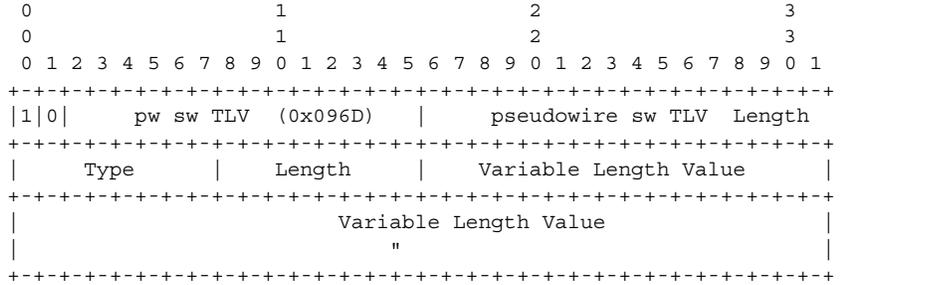
Pseudowire status notification messages can be generated by the T-PE nodes and/or the S-PE nodes. Pseudowire status notification messages received by a switching node are processed and then passed on to the next hop. An S-PE node appends the optional pseudowire switching TLV, with its system address 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, it means 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 spokes 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 SDP-binding down status bits regardless if the received status bits from the remote node indicated SAP up/down or SDP-binding up/down.

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.

Pseudowire Switching Point Sub-TLVs

Below are details specific to pseudowire switching point sub-TLVs:

pseudowire ID of last pseudowire segment traversed — This sub-TLV type contains a pseudowire ID in the format of the pseudowire ID

Pseudowire switching point description string — An optional description string of text up to 80 characters long.

IP address of pseudowire switching point.

The IP V4 or V6 address of the pseudowire switching point. This is an optional sub-TLV.

MH VCCV capability indication.

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.

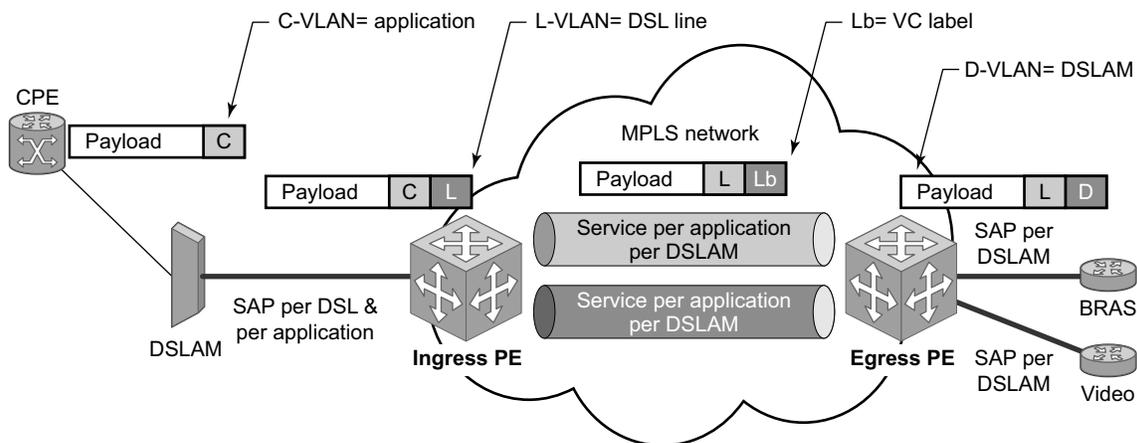
When the label mapping for the egress direction is also received from the T-LDP peer, and the information in the FEC matches that of the local configuration, the static-to-dynamic cross-connect is effected.

Note that it is possible that end nodes of a static pseudowire segment be misconfigured. In this case, an S-PE or T-PE node may be receiving packets with the wrong encapsulation. In this case, it is possible that an invalid payload will be forwarded over the pseudowire or the SAP respectively. Furthermore, 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 this fails the VCCV packet will be discarded.

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.



Fig_36

Figure 7: Ingress VLAN Swapping

The network diagram in [Figure 7](#) 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”.

Ingress VLAN Translation

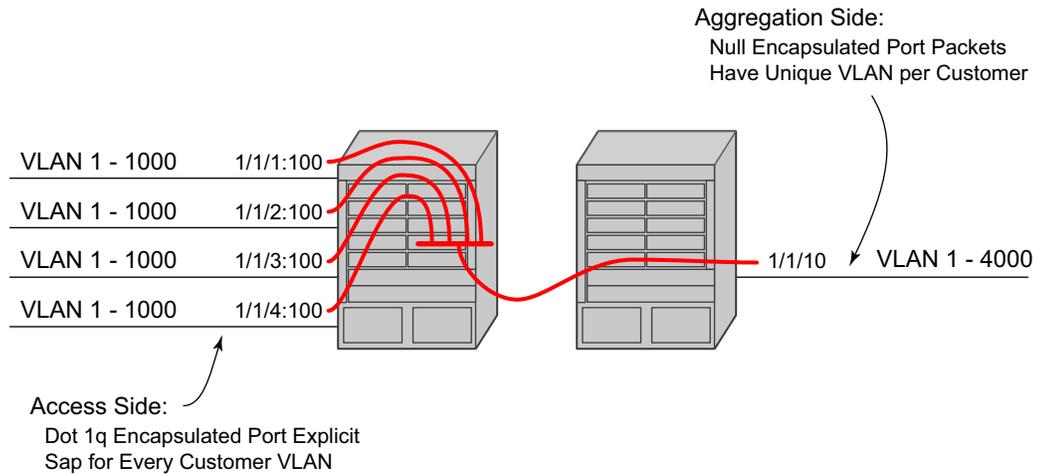


Figure 8: Ingress VLAN Translation

The drawing in [Figure 8](#) indicates an application where different circuits are aggregated in the VPLS-based network. The access side is represented by an explicit dot1q encapsulated SAP. As the VLAN-id is port specific, those connected to different ports might have the same VLAN. The aggregation side (the right side [Figure 8](#)) is aggregated on the same port, and hence, unique a VLAN-id is required.

Pseudowire Redundancy

Pseudowire redundancy provides the ability to protect a pseudowire with a pre-provisioned pseudowire and to switch traffic over to the 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, the pseudowire is also protected. However, there are a couple of 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 SR-Series 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.

Dynamic Multi-Segment Pseudowire Routing

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, so allowing dynamic routing of the MS-PW through the intervening network of S-PEs. Dynamic multi-segment pseudowires are described in the IETF in draft-ietf-pwe3-dynamic-ms-pw-13.txt.

Figure 9 illustrates the operation of dynamic MS-PWs.

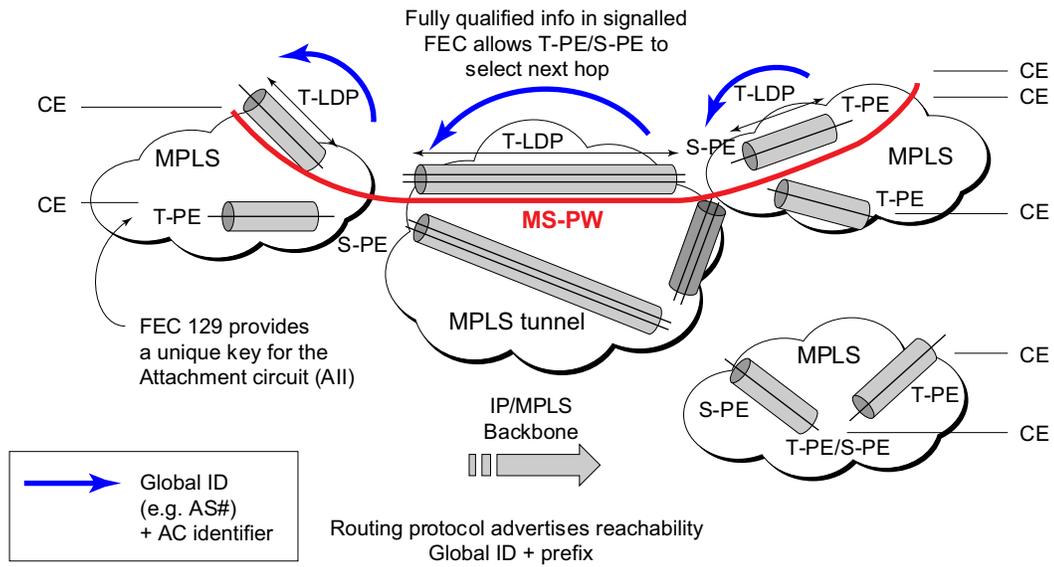


Figure 9: Dynamic MS-PW Overview

The FEC 129 AII Type 2 structure depicted in [Figure 10](#) is used to identify each individual pseudowire endpoint:

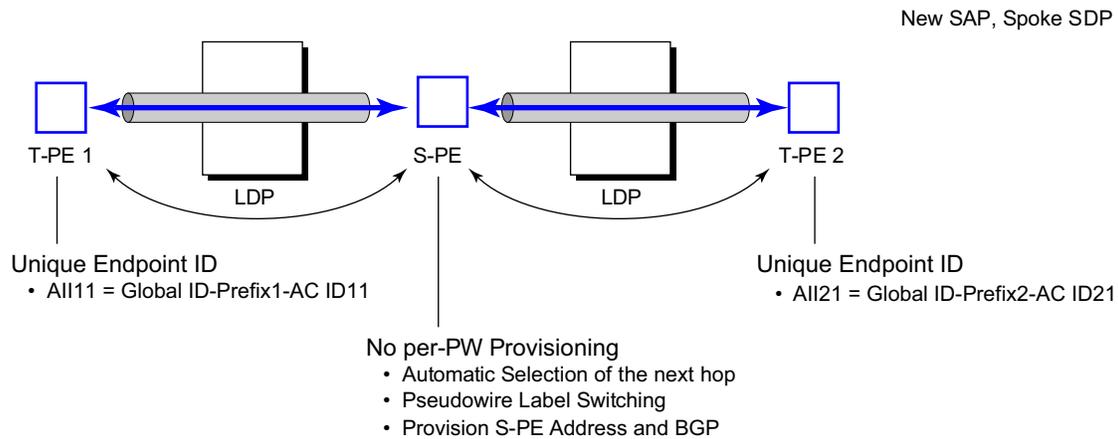


Figure 10: MS-PW Addressing using FEC129 AII Type 2

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.

This new AII type is mapped into the MS-PW BGP NLRI (a new 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 11](#):

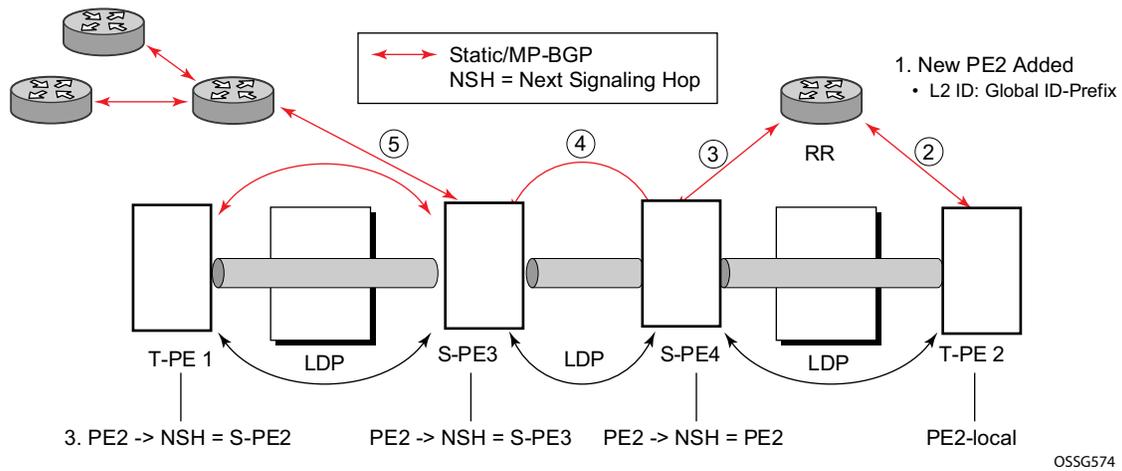


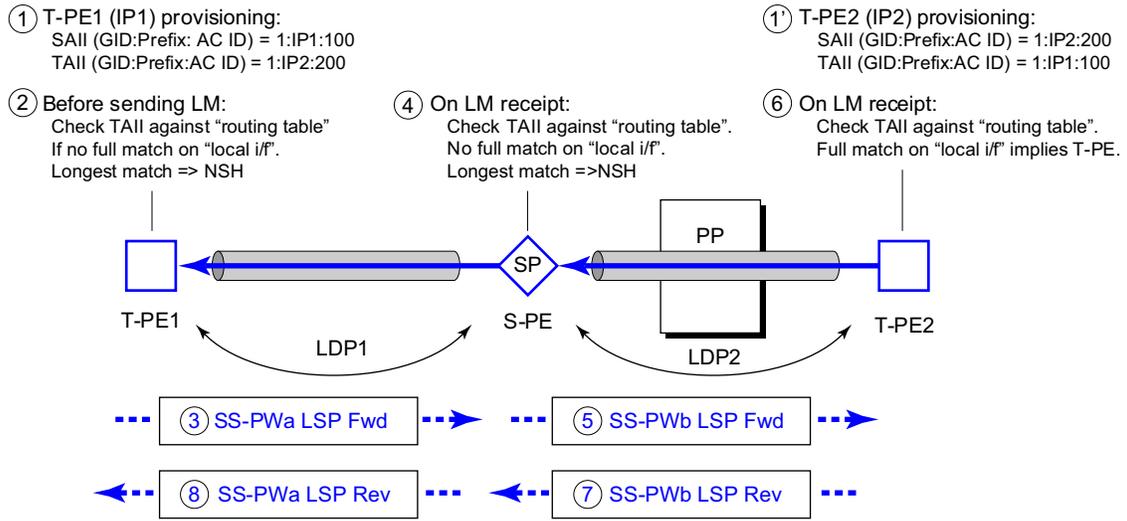
Figure 11: Advertisement of PE Addresses by PW Routing

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

Next, in steps 2-5, 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.

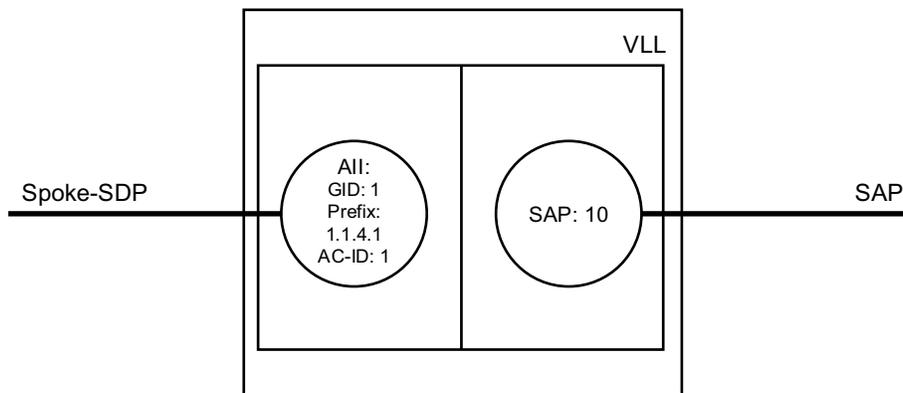
VLL services can then be established, as illustrated in [Figure 12](#).



OSSG575

Figure 12: Signaling of Dynamic MS-PWs using T-LDP

In step 1 and 1' the T-PEs are configured with the local and remote endpoint information, Source AII (SAII), Target AII (TAII). On the 7x50, the AIIs are locally configured for each spoke SDP, according to the model shown in Figure 13. The 7x50 therefore provides for a flexible mapping of 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.



OSSG576

Figure 13: Mapping of All 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 and thus uses a statically configured PW segment, the signaling of individual segments follows the procedures already implemented in the PW Switching feature. Note that 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 once 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 14](#).

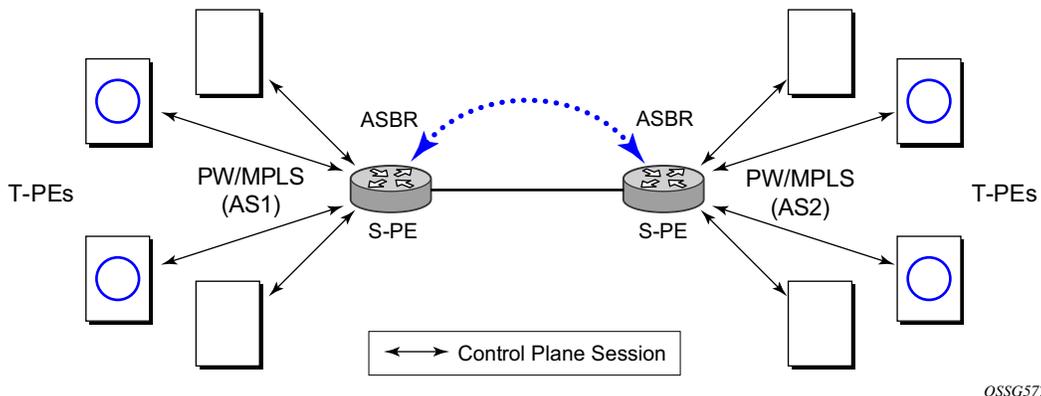


Figure 14: VLL Using Dynamic MS-PWs, Inter-AS Scenario

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

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 given 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 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 **config>router>bgp>family ms-pw**.

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

In order to enable support for dynamic MS-PWs on a 7x50 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 7x50 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 SPDs 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 be configured with the following parameters:

- a. Global ID — This is a 4 byte identifier that uniquely identifies an operator or the local network.
- b. 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.
- c. For each local prefix, at least one 8-byte route distinguisher 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:

Dynamic Multi-Segment Pseudowire Routing

```
*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.

Static Routing

In addition to 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, then this represents the default route. If a static route exists to a given T-PE, then this is used in preference to any BGP route that may exist.

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.

Note that it is possible to configure explicit paths independently of the configuration of BGP or static routing.

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 TAIID identifies the remote end. The SAII and TAIID 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 this end of the pseudowire. This should be locally unique within the scope of the global-id:prefix.

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 and 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 7x50 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 TAI 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 the spoke-sdp-fec. If master is selected at a given T-PE, then it 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 below), 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.

Automatic Endpoint Configuration

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

Automatic endpoint configuration is supported required 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 TAI from the configuration. When auto-configuration is used, the node assumed passive behavior from a point of view of T-LDP signaling (see above). Therefore, the far-end T-PE must be configured for signaling master for that spoke-sdp-fec.

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 simply 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

In order 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, then 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 TAIL.

Pseudowire routing chooses the MS-PW path in terms of the sequence of S-PEs to use to reach a given 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 given 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 a 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.

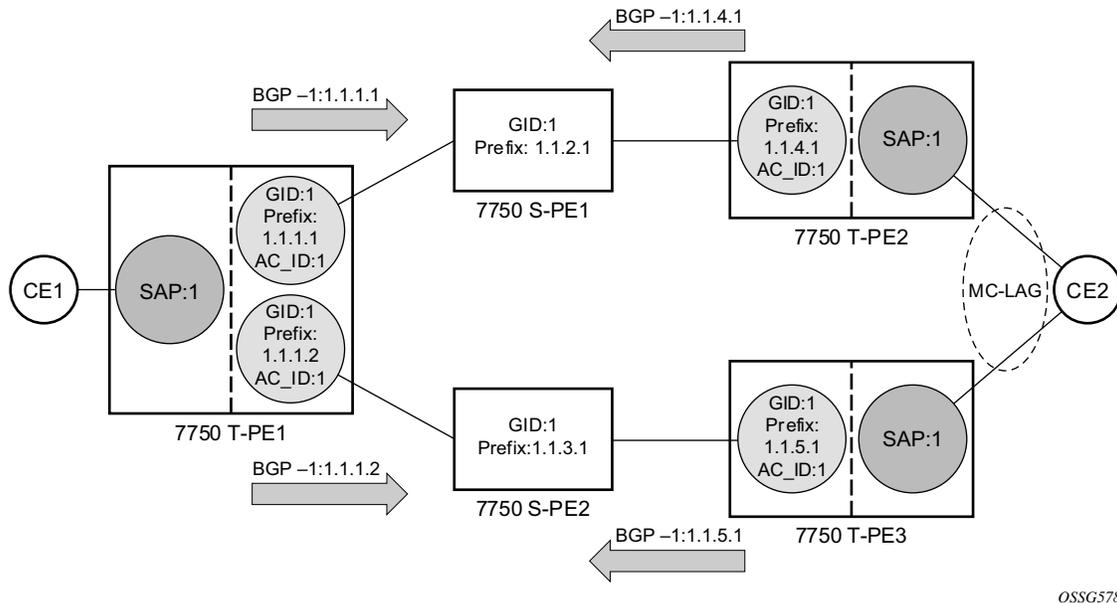
Pseudowire Templates

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.

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 SAI/TAII.

Figure 15 illustrates the use of pseudowire redundancy.



OSSG578

Figure 15: 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 SAI/TAII set and is signalled 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 given 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 route distinguisher.

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

Note that 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.

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.

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, then 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. Note that if SAII:TAII is entered in addition to the *spoke-sdp-fec-id*, then 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.

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

VCCV-Trace on Dynamic MS-PWs

The 7x50 supports 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. Note that the use of *spoke-sdp-fec-id* in vccv-trace is only applicable at T-PE nodes, since it is not configured for a given MS-PW at S-PE nodes.

Example Dynamic MS-PW Configuration

This section presents an example of how to configure Dynamic MS-PWs for a VLL service between a set of 7x50 nodes. The network consists of two 7x50 T-PEs and two 7x50 playing the role of S-PEs, as shown in the following figure. Each 7x50 peers with its neighbor using LDP and BGP.

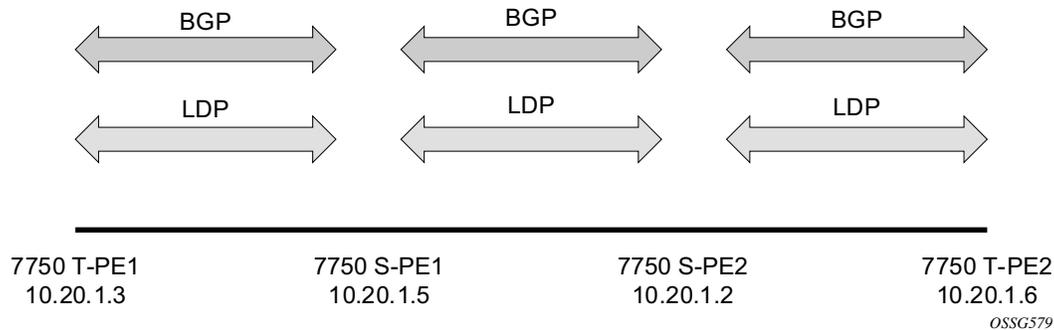


Figure 16: Dynamic MS-PW Example

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.

Once this routing and signaling infrastructure is established, spoke-sdp-fecs can be configured on each of the T-PEs.

Example Dynamic MS-PW Configuration

```
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
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
```

T-PE-1

```
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
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
```

T-PE-2

```

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

```

S-PE-1

```

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

```

S-PE-2

VLL Resilience with Two Destination PE Nodes

Figure 17 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.

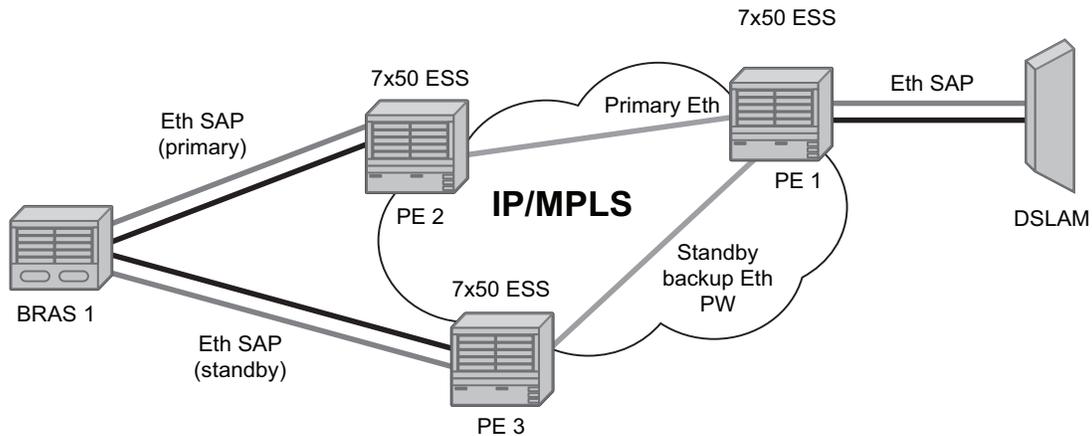


Figure 17: VLL Resilience

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. In order to avoid black holing of in-flight 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 on page 103](#), 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 operational status to DOWN. The following are the events which 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 a LDP or GRE based one. This provides the ability to route the path of the two pseudowires over different areas of the network.

Alcatel-Lucent's 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.

Alcatel-Lucent's 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.

Master-Slave Operation

Master-Slave pseudowire redundancy is discussed in this section. 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 signalled 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 18](#) 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, and thus PE1 is dual-homed to PE2 and PE3 using Epipe spoke SDPs. The objectives 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 towards 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.

transitioning a spoke SDP from standby to active receiving and/or processing a pseudowire preferential forwarding status message before those transitioning a spoke SDP to standby. This transient condition is most likely when a forced switch-over is performed, or the relative preferences of the spoke SDPs is 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, a failure of connectivity on the path of the spoke SDP or the SAP, would not result in such loops in the data path.

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 towards the CE, such as Ethernet CCM, 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 towards a SAP which will drop it.

Figure 19 shows an example for the case of Ethernet LMI.

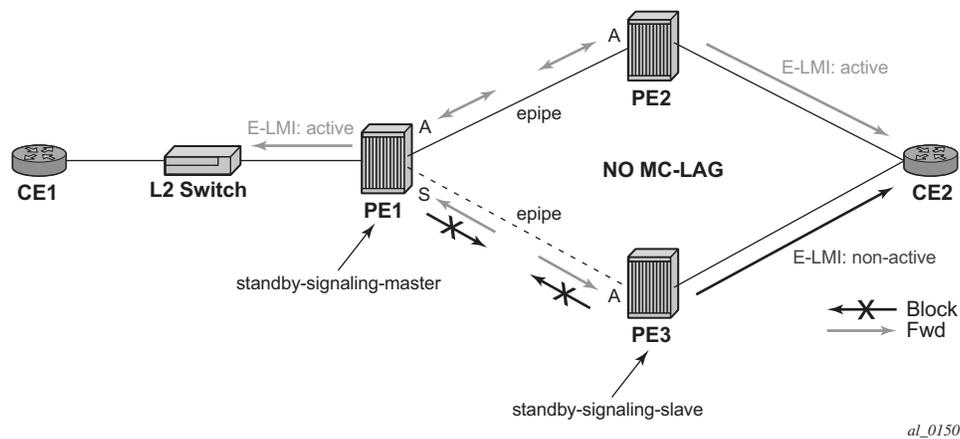


Figure 19: Example of SAP OAM Interaction with Master-Slave Pseudowire Redundancy

Local Rules at Slave VLL PE

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

If **standby-signaling-slave** is configured on a given spoke SDP or explicit endpoint, then the following rules apply. Note that 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 a MC-LAG/MC-APS instance, a forwarding status of ACTIVE is always advertised.
- Rules for processing and merging local and received endpoint object status Up/Down 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 all objects in endpoint 'X' transition locally to Down state, and/or received a "SAP Down" notification via remote T-LDP status bits or via SAP specific OAM signal, and/or received status bits of "SDP-binding down", and/or received status bits of "PW not forwarding", the node must send status bits of "SAP Down" over all 'Y' endpoint spoke SDPs.
 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 received T-LDP "SDP-binding down" status bits, and/or received 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 as per the pseudo-code in Section 5.1.2.
 6. If, all objects in endpoint 'Y', or a single spoke SDP that exists outside of an endpoint (and no endpoint exists), transition locally to down state, and/or received T-LDP "SAP Down" status bits, and/or received T-LDP "SDP-binding down" status bits, and/or received status bits of "PW not forwarding", and/or the received status bits of 'PW FWD standby', the node must send a "SAP down" notification on the 'X' endpoint SAP via the SAP specific OAM signal, if applicable.
 7. If the peer PE for a given 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 given object in endpoint 'Y' does not signal 'PW FWD standby', then spoke SDP is eligible for selection.

Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios

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

VLL Resilience

Figure 20 displays a VLL resilience path example. An sample configuration follows.

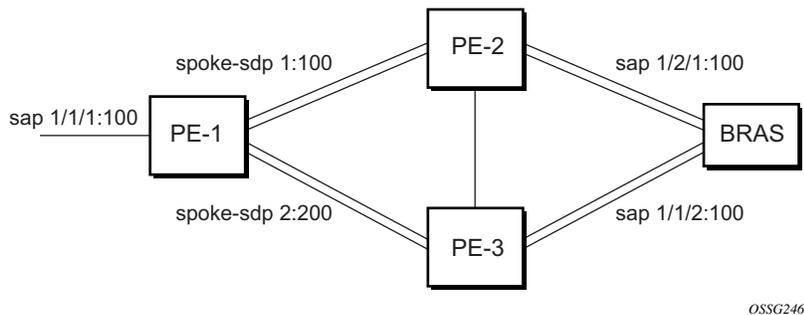


Figure 20: VLL Resilience

Note that 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

```

PE1
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
PE2
configure service epipe 1
  endpoint X
  exit
  sap 2/2/2:200 endpoint X
  spoke-sdp 1:100
  standby-signaling-slave
  
```

VLL Resilience with Two Destination PE Nodes

PE3

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

VLL Resilience for a Switched Pseudowire Path

Figure 21 displays a VLL resilience for a switched pseudowire path example. A sample configuration follows.

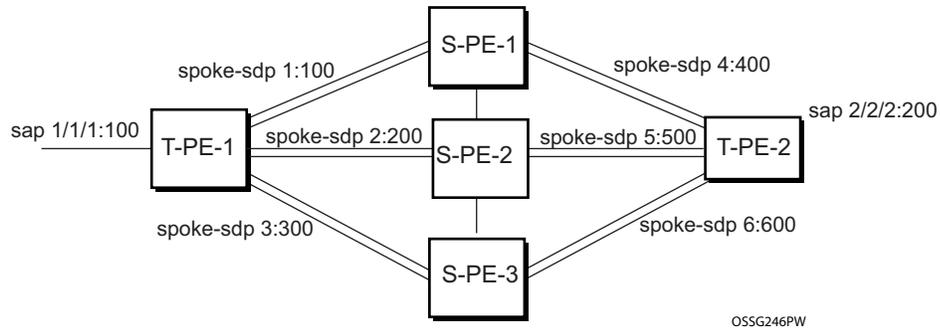


Figure 21: VLL Resilience with Pseudowire Switching

Configuration

```
T-PE1
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-PE2
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
```

Pseudowire SAPs

```
precedence primary
spoke-sdp 5:500 endpoint Y
precedence 1
spoke-sdp 6:600 endpoint Y
precedence 1
```

S-PE1

VC switching indicates a VC cross-connect so that the service manager does not signal the VC label mapping immediately but will put this into passive mode.

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

Pseudowire SAPs

Refer to the *SR OS Layer 3 Services Guide* for details of how to use pseudowire SAPs with Layer-2 services.

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. Note that MC-LAG (see, [Access Node Resilience Using MC-LAG and Pseudowire Redundancy on page 103](#)) offers access node redundancy with active/stand-by links while Ethernet Tunnels offers 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.

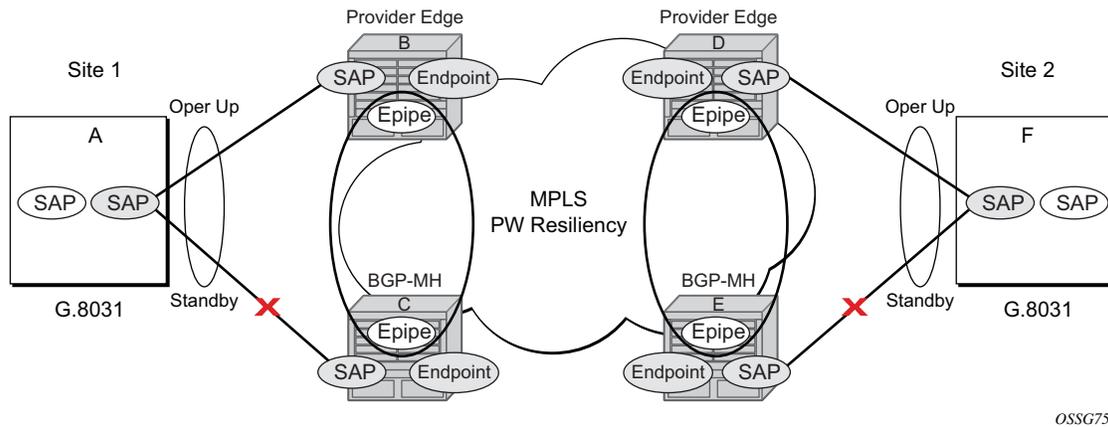
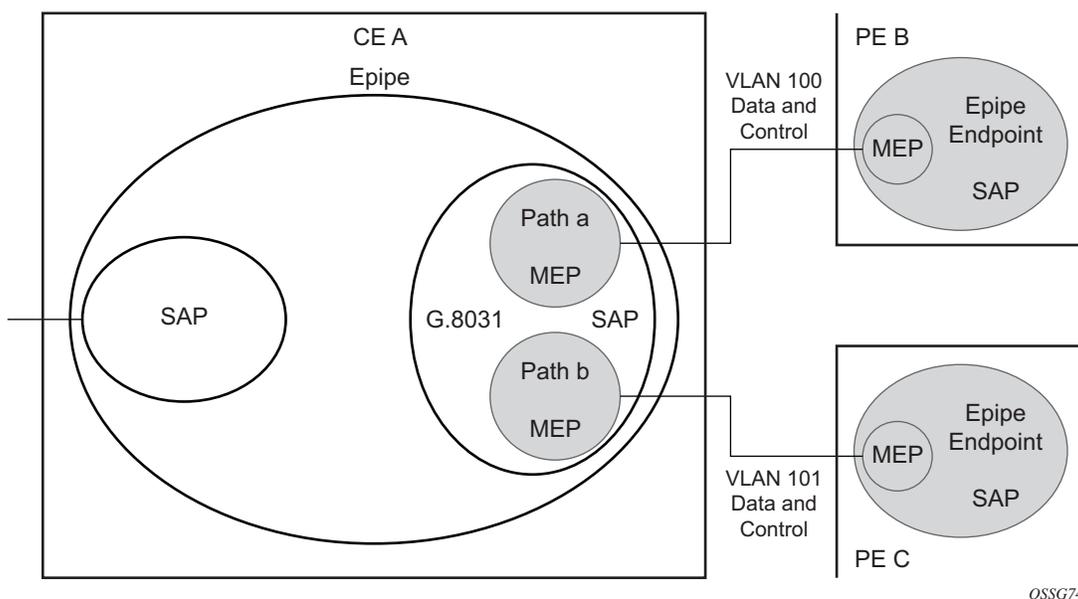


Figure 22: BGP-MH Site Support for Ethernet Tunnels

Figure 22 shows the BGP-MH Site Support for Ethernet Tunnels; where a G.8031 edge device (A) is configured to 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.

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 23 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.



OSSG749

Figure 23: 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 connect to a Epipe SAP on a PE which uses BGP-MH and operational groups to control other G.8031 tunnels. This dedicated G.8031 still has a data control capabilities, but the data Epipe service is not bearing user data packets. On the CE, this dedicated G.8031 is only used for group control. The choice of making this a dedicated Control for a set of G.8031 tunnels is merely to simplify operation and allow individual disabling of services. Using a dedicated G.8031 for 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 slaved 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 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, and Fault to control the G.8031 SAP. However in the 7x50 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.

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

Detailed Operation

For this feature, BGP-MH is used 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. Note with G.8031 by controlling it with BGP-MH, the G.8031 CE is forced to be slaved to the PE BGP-MH election. BGP-MH election is control by the received VPLS preference or BGP local-preference or PE Id (IP address of Provider Edge) if local-preference is equal. 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.

Epipes Using BGP-MH Site Support for Ethernet Tunnels

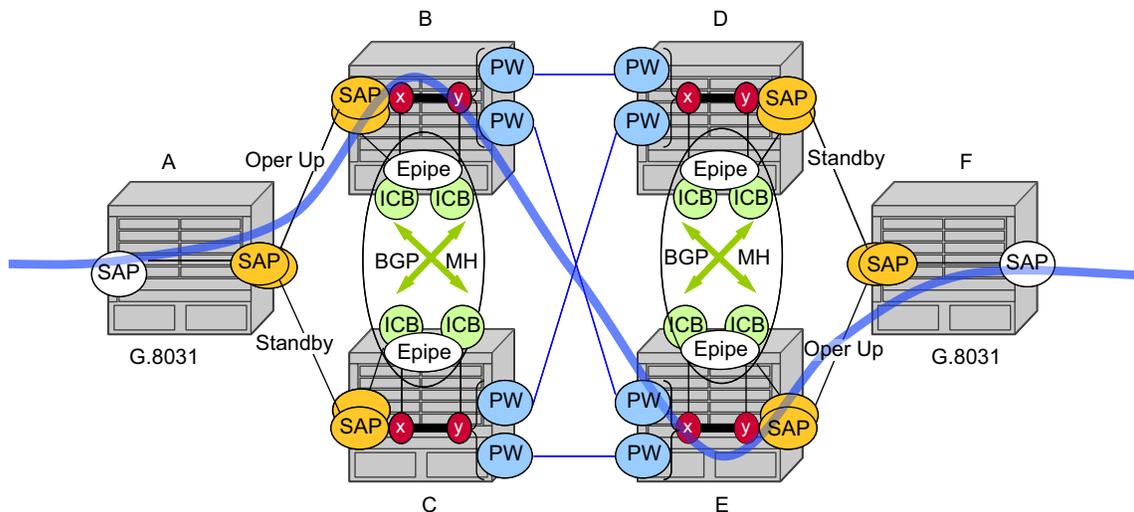
There are two configuration options:

- Every G.8031 service SAP terminates on a single Epipes that has BGP-MH. These Epipes may utilize endpoints with or without ICBs.
- A control Epipes service that monitors a single SAP that is used for group control of fate shared CE services. In this case, the Epipes 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 Epipes 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, the Ethernet Tunnel inactive and is used CCM as an aliveness mechanism. There is no G.8031 logic on the PE, the SAP is simply controlling the correspond CE SAP.

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 but it is suggested to run commands from the BGP-MH side to control these connections.



OSSG751

Figure 24: Full Redundancy G.8031 Epipes & BGP-MH

Table 4 lists the SAP MEP signaling shown in Figure 24. For a description of the events shown in this sample operation, see Events in Sample Operation on page 89.

Table 4: 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 ^a	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 ^b	No Fault	Fault	Fault
5	Down Path A & B fault at A	Fault	No Fault	Fault	Fault
6	Partitioned Network Use Path Precedence Up use Path A	No Fault	No Fault	No Fault	No Fault

a. No Fault = no ifStatusTlv transmit | CCM transmit normally

b. Fault = ifStatusTlv transmit down | no CCM transmit

Events in Sample Operation

The following represents a walk through of the events for switchover in Figure 24. This configuration uses operational groups. The nodes of interest are A, B and C listed in Table 4.

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 desired.
 - An Epipe service is provisioned on each PE node (B,C) purely 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 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 use 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 4](#) row 1 (Fault is signaled in the other direction PE to CE).
 - BGP-MH makes its determination of 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 4](#) row 2.
 - Since fault propagation is enabled under the SAP node C MEP, and ifStatusTlv is operationally Down is 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 time out 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 4](#) 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. Node B on the PE the SAP also detects the failure and the propagation of fault status goes to BGP-MH; see [Table 4](#) 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, Service Manager feeds this to CFM, which then propagates a Fault towards Node A. This is a cyclic fault propagation. However, since path A-B is broken, the situation is stable; see [Table 4](#) 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 4](#) row 3.

13. BGP-MH on node C elects sap and bring it up.

14. ET port transitions to port A-C 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 to the CE Fault.
2. In the case of G.8031 on the CE detecting the failure, it takes the tunnel down and signals a fault to the PE, and then the SAP propagates that to BGP-MH.

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.

BGP-MH Specifics for MH Site Support for Ethernet Tunnels

[BGP Multi-Homing for VPLS on page 442](#) 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.

PW Redundancy for BGP MH Site Support for Ethernet Tunnels

[Pseudowire Redundancy Service Models on page 107](#) and [Figure 27 on page 105](#) are used for the MPLS network resiliency. BGP MH Site Support for Ethernet Tunnels reuses this model.

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

Using [Figure 27](#) as a reference, the following are the rules for generating, processing, and merging T-LDP status notifications in VLL service with endpoints.

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 a 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 a 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.
-

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 SAP 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 a 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 the pseudo-code per Pseudowire Redundancy procedures.
5. If all objects in endpoint X transition locally to Down state due to operator or BGP-MH DF election, or received a SAP Down notification via remote T-LDP status bits or via SAP specific OAM signal (SAP MEP), or received status bits of SDP-binding Down, or received sta-

tus bits of PW not forwarding, the node must send status bits of SAP Down over all Y endpoint spoke-SDPs, including the ICB.

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, received T-LDP SAP Down status bits, or received T-LDP SDP-binding Down status bits, or received 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 as per Pseudowire Redundancy procedures.
9. If all objects in endpoint Y, except the ICB spoke-SDP, transition locally to Down state, or received T-LDP SAP Down status bits, or received T-LDP SDP-binding Down status bits, and/or received status bits of PW not forwarding, the node must send status bits of SDP-binding Down over the X endpoint ICB spoke-SDP only.
10. If all objects in endpoint Y transition locally to Down state, or received T-LDP SAP Down status bits, or received T-LDP SDP-binding Down status bits, or received status bits of PW not forwarding, the 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 via the SAP specific OAM signal in this case the SAP MEP ifStatusTlv operationally-Down and also signal the BGP-MH Site, if this SAP is part of a BGP Site.

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. Note that 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 this release.

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

The following are supported:

- All Ethernet-tunnels G.8031 SAPs on CE:
 - 7x50 G.8031 in load sharing mode (recommended)
 - 7x50 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, namely:
 - MC-LAG
 - MC-EP
 - MC-ring
 - MC-APS
- Master and Slave pseudowire is not supported.

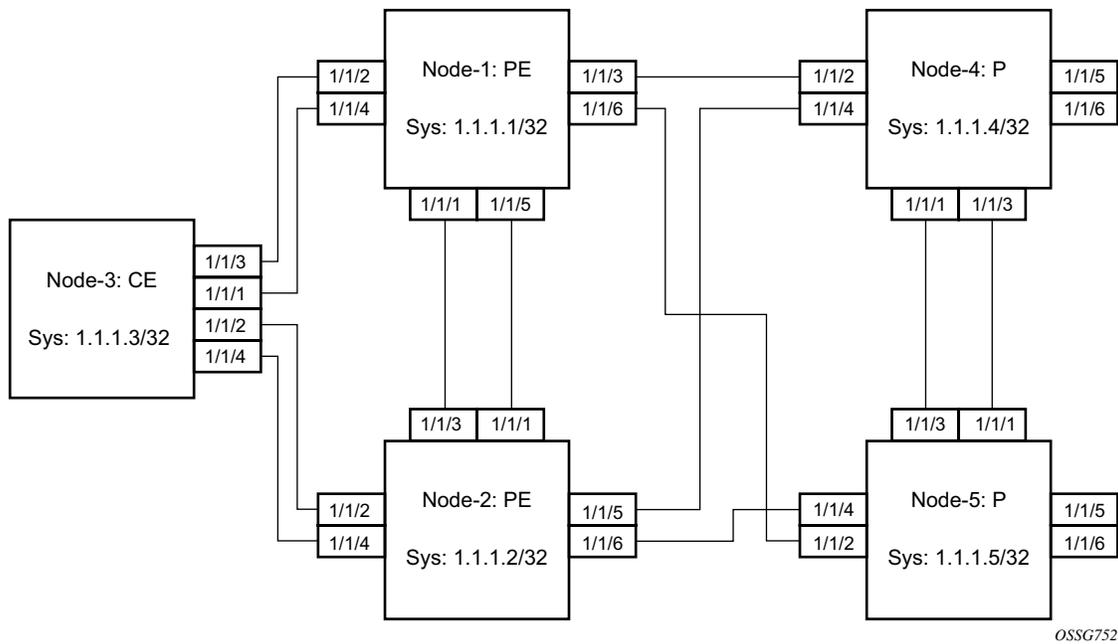


Figure 25: Sample Topology Full Redundancy

Refer to [Configuration Examples on page 95](#) for configuration examples derived from [Figure 25](#).

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
```

Epipe Using BGP-MH Site Support for Ethernet Tunnels

```
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
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

```

Epipe Using BGP-MH Site Support for Ethernet Tunnels

```
#-----
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

```

Epipe Using BGP-MH Site Support for Ethernet Tunnels

```
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

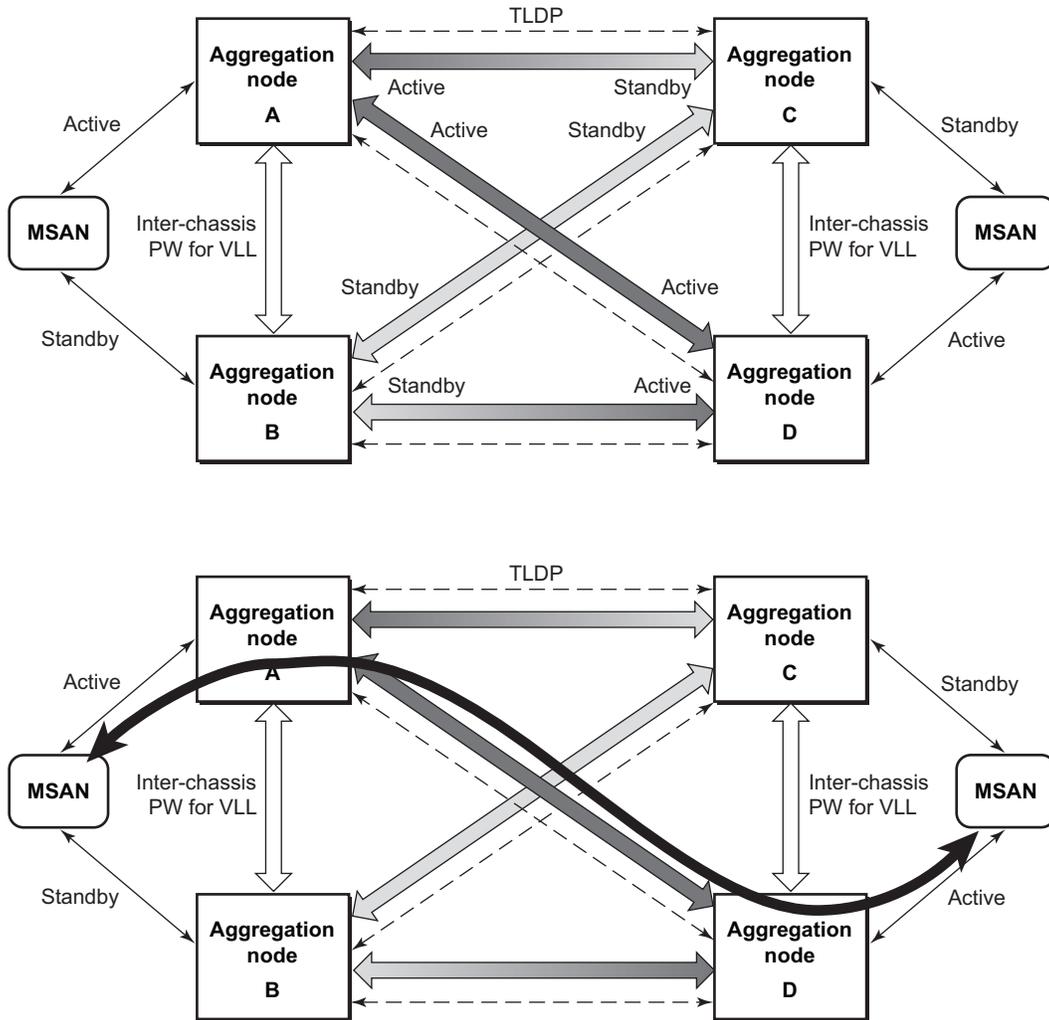
```

Epipe Using BGP-MH Site Support for Ethernet Tunnels

```
#-----  
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
```

Access Node Resilience Using MC-LAG and Pseudowire Redundancy

Figure 26 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.



OSSG116

Figure 26: 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 thus 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 and thus advertises standby status notification messages to both nodes C and D. Node C performs the same operation.

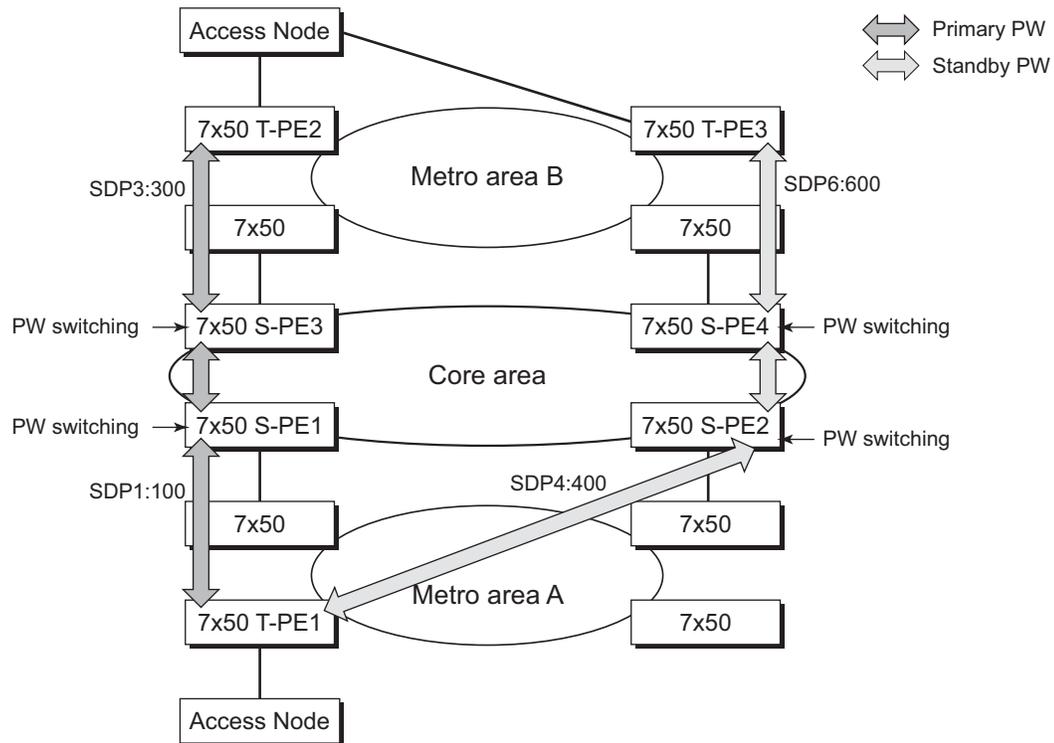
An SR-Series node selects a pseudowire as the active path for forwarding packets when both the local pseudowire status and the received remote pseudowire status indicate active status. However, an SR-Series device in standby status according to the SAP in its local MC-LAG instance is capable of processing packets for a VLL service received over any of the pseudowires which 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 a MC-LAG node to backup a 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.

Note that at configuration time, the user specifies a precedence parameter for each of the pseudowires which are part of the redundancy set as described in the application in [VLL Resilience with Two Destination PE Nodes on page 74](#). An SR-Series node uses this to select which pseudowire to forward packet 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. Furthermore, ICB spoke SDP can only be added to the SAP side of the VLL cross-connect if the SAP is configured on a MC-LAG instance.

VLL Resilience for a Switched Pseudowire Path

Figure 27 illustrates the use of both pseudowire redundancy and pseudowire switching to provide a resilient VLL service across multiple IGP areas in a provider network.



OSSG114

Figure 27: 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 on page 74](#), the T-PE1 node switches the path of a VLL to a secondary standby pseudowire in the case of a network side failure causing the VLL binding status to be 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.

Note that 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 since 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. At that point in time, 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 becomes up even if it has a higher precedence than the currently active secondary path.

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](#).

Redundant VLL Service Model

In order to implement pseudowire redundancy, a VLL service accommodates more than a single object on the SAP side and on the spoke SDP side. [Figure 28](#) illustrates the model for a redundant VLL service based on the concept of endpoints.

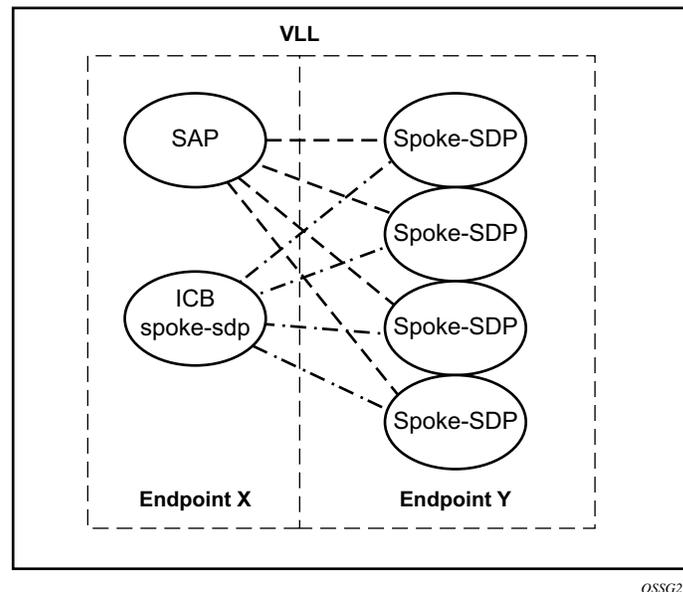


Figure 28: Redundant VLL Endpoint Objects

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

In order to add more objects, up to two (2) 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 28](#).

Note that [Figure 28](#) is merely an example and that the “Y” endpoint can also have a SAP and/or an ICB spoke SDP. The following details 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 it 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 SDP 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 — Special pseudowire used for MC-LAG and pseudowire redundancy application. Forwarding between ICBs is blocked on the same node. The user has to explicitly indicate the spoke SDP is actually an ICB at creation time. There are however a few scenarios below where the user can configure the spoke SDP as ICB or as a regular spoke SDP on a given 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 given time but can receive from all endpoint objects

An explicitly named endpoint can have a maximum of one SAP and one ICB. Once 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 a MC-LAG instance. Conversely, a SAP which is not part of a MC-LAG instance cannot be added to an endpoint which already has an ICB spoke SDP.

An explicitly named endpoint, which 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.

T-LDP Status Notification Handling Rules

Referring to [Figure 28 on page 107](#) as a reference, the following are the rules for generating, processing, and merging T-LDP status notifications in VLL service with endpoints. Note that any allowed combination of objects as specified in [Redundant VLL Service Model on page 107](#) can be used on endpoints “X” and “Y”. The following sections refer to the specific combination objects in [Figure 28](#) as an example to describe the more general rules.

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 a MC-LAG instance, the forwarding status of active is always advertised.

When the SAP in endpoint “X” is part of a MC-LAG 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 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 a MC-LAG instance, then the forwarding status sent over all “Y” endpoint spoke SDP's should always be set to zero (active by default).

Processing and Merging

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.

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. Note that Ethernet SAP does not support SAP OAM protocol. All other SAP types cannot exist on the same endpoint as an ICB spoke SDP since non Ethernet SAP cannot be part of a 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” transition locally to down state, and/or received a SAP down notification by remote T-LDP status bits or by SAP specific OAM signal, and/or received status

T-LDP Status Notification Handling Rules

bits of SDP-binding down, and/or received status bits of pseudowire not forwarding, the node must send status bits of SAP down over all “Y” endpoint spoke SDPs, including the ICB.

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, received T-LDP SAP down status bits, and/or received T-LDP SDP-binding down status bits, and/or received status bits of pseudowire not forwarding, 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 “Y”, except the ICB spoke SDP, transition locally to down state, and/or received T-LDP SAP down status bits, and/or received T-LDP SDP-binding down status bits, and/or received status bits of pseudowire not forwarding, the node must send status bits of SDP-binding down over the “X” endpoint ICB spoke SDP only.

If all objects in endpoint “Y” transition locally to down state, and/or received T-LDP SAP down status bits, and/or received T-LDP SDP-binding down status bits, and/or received status bits of pseudowire not forwarding, the 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. An Ethernet SAP does not support signaling status notifications.

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*.

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 29 shows how this service would be used to provide a virtual lease-line service across an MPLS network between two sites, A and B.

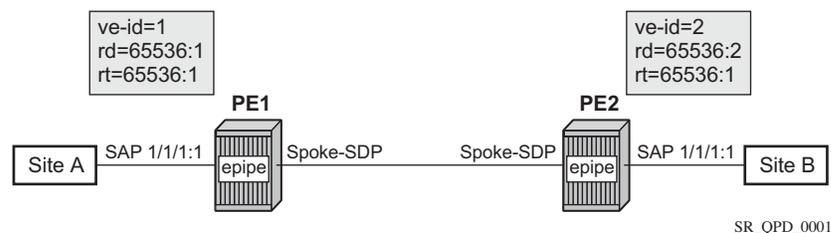


Figure 29: Single-Homed BGP-VPWS Example

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 which is signaled using BGP VPWS updates over a given tunnel LSP.

Dual-Homed BGP VPWS

A BGP-VPWS service can benefit from dual-homing, as described in draft-ietf-l2vpn-vpls-multihoming-03. 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.

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 re-established 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.

An example is shown in [Figure 30](#).

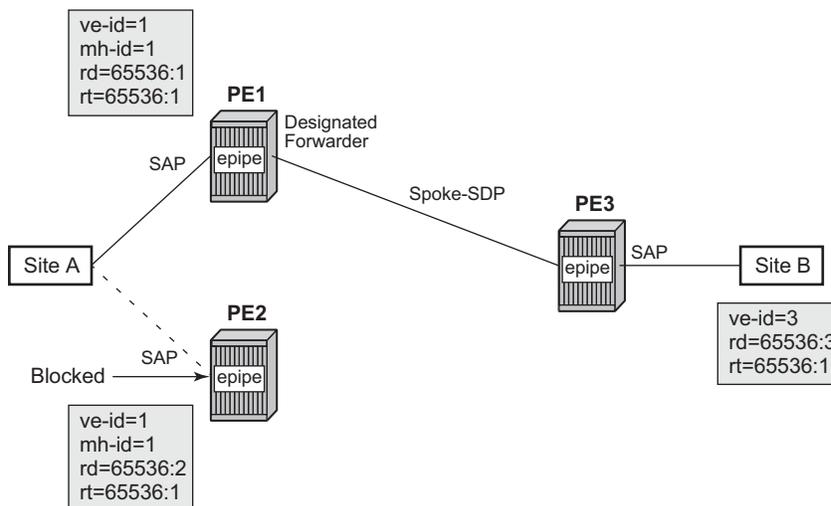


Figure 30: Dual-Homed BGP VPWS with Single Pseudowire

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 (note that 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 (thereby 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 then both the D and CSV bits (see below 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.

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 topology example is shown in [Figure 31](#).

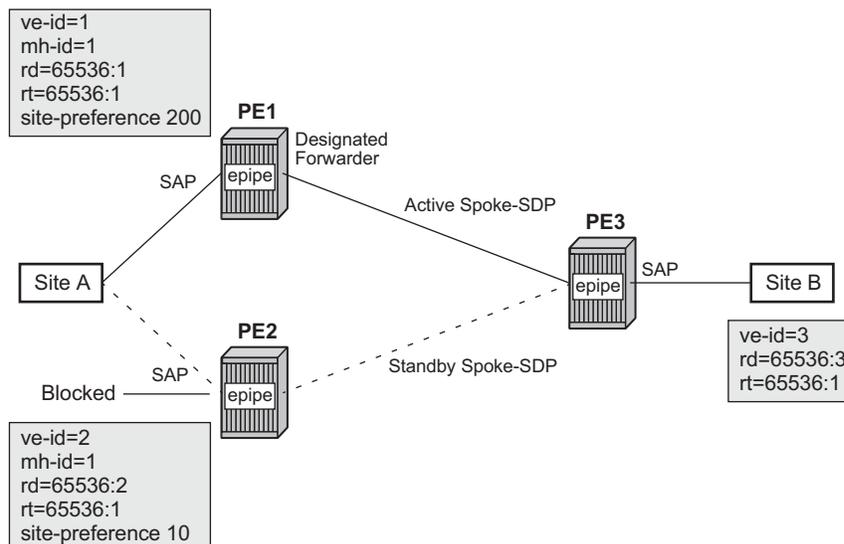


Figure 31: Dual-homed BGP VPWS with Active/Standby Pseudowires

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 traffic on the pseudowire connected to the designated forwarder, PE1.

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's status signaling, specifically using the BGP update D/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 shown below:

```
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

```

Pseudowire Signaling

The BGP signaling mechanism used to establish the pseudowires is described in the BGP VPWS 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 draft-ietf-l2vpn-vpls-multihoming-03.

Figure 32 displays the format of the BGP VPWS update extended community.:

```

+-----+
| Extended community type (2 octets) |
+-----+
| Encaps Type (1 octet) |
+-----+
| Control Flags (1 octet) |
+-----+
| Layer-2 MTU (2 octet) |
+-----+
| VPLS Preference (2 octets) |
+-----+

```

Figure 32: 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 oper state down.
- Control Flags — Control information regarding the pseudowires, see below for details.
- Layer-2 MTU is the 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 oper state down and a 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.

The control flags are described below:

```

0 1 2 3 4 5 6 7
+++++
|D|A|F|Z|Z|Z|C|S| (Z = MUST Be Zero)
+++++

```

The following bits in the Control Flags are defined:

- D — Access circuit down indicator from 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 as 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 send and are expected to be received, with a control word. When this is set to 0, packets will be send and are expected to be received, without a control word. This is the 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 33](#).

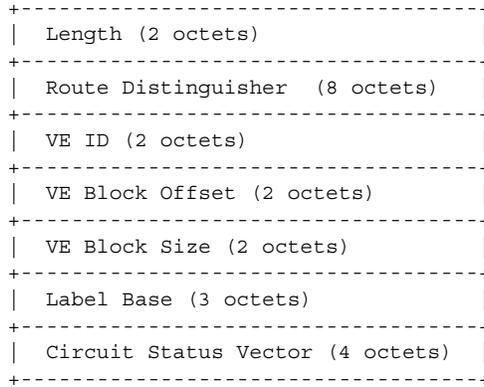


Figure 33: BGP VPWS NLRI

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 as a VE block size of 1 is always used.

The circuit status vector is encoded as a TLV as shown in [Figure 34](#).

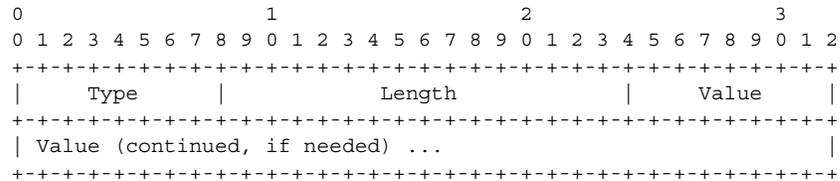


Figure 34: BGP VPWS NLRI TLV Extension Format

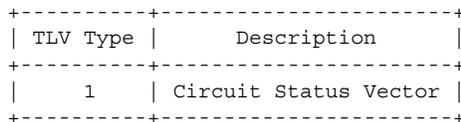


Figure 35: Circuit Status Vector TLV Type

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. As 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 of greater than 800, a notification message will be sent and the BGP session will restart. Also, BGP VPWS services support a single access circuit, consequently only the most significant bit of the CSV is examined on receipt.

A pseudowire will be established when a BGP VPWS update is received which 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 draft-ietf-l2vpn-vpls-multihoming-03.

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

BGP VPWS Configuration Procedure

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

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

Use of Pseudowire Template for BGP VPWS

The pseudowire template concept used for BGP AD is re-used for BGP VPWS to dynamically instantiate pseudowire (SDP-bindings) and the related SDP (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 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 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 oper state down.
 - In order to interoperate with existing implementations if the received MTU value = 0, then MTU negotiation does not take place; the related pseudowire is setup ignoring the MTU.
- If the values 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] [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
```

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

The **tools perform** command can be used similarly as for BGP-AD to force the application of changes in pseudowire-template using the format described below:

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

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.

Note that the command:

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

will have no affect on an automatically created VPWS endpoint.

VLL Service Considerations

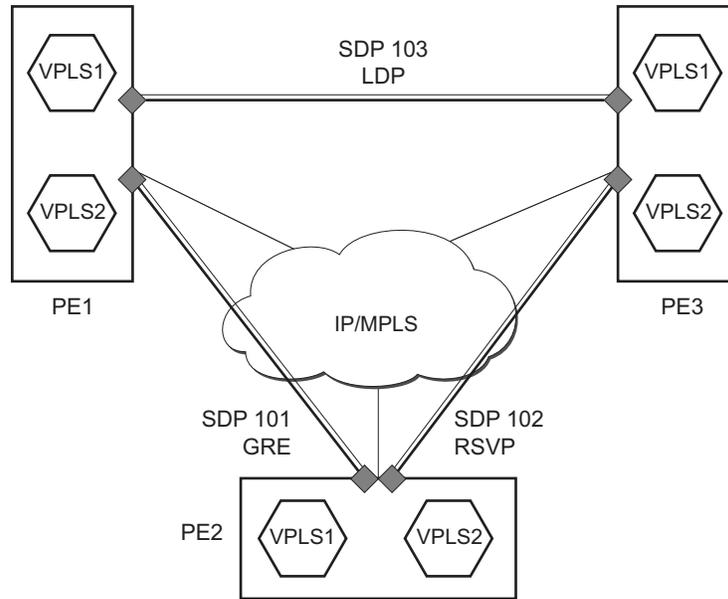
This section describes the general , , and 7950 XRS service features and any special capabilities or considerations as they relate to VLL services.

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.

SDP Statistics for VPLS and VLL Services



OSSG208

Figure 36: SDP Statistics for VPLS and VLL Services

The simple three-node network described in [Figure 36](#) 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.

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 , , and 7950 XRS Epipe service:

- Ethernet null
- Ethernet dot1q
- QinQ
-

Note that while different encapsulation types can be used, encapsulation mismatching can occur if the encapsulation behavior is not understood by connecting devices and 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.

QoS Policies

When applied to , , or 7950 XRS Epipe 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 services, egress QoS policies function as with other services where the class-based queues are created as defined in the policy. Note that both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in a service.

Filter Policies

, , and 7950 XRS Epipe 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.

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 , , and 7950 XRS Epipe implementation does not perform any MAC learning on the service, so Epipe services do not consume any MAC hardware resources.

Configuring a VLL Service with CLI

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

Topics in this section include:

- [Basic Configurations on page 128](#)
- [Common Configuration Tasks on page 128](#)
 - [Configuring VLL Components on page 129](#)
 - [Creating an Epipe Service on page 130](#)
 - [Using Spoke SDP Control Words on page 141](#)
 - [Configuring Pseudowire Scenarios](#)
 - [Pseudowire Configuration Notes on page 142](#)
 - [Configuring Two VLL Paths Terminating on T-PE2 on page 143](#)
 - [Configuring VLL Resilience on page 146](#)
 - [Configuring VLL Resilience for a Switched Pseudowire Path on page 147](#)
 - [Configuring BGP Virtual Private Wire Service \(VPWS\) on page 149](#)
- [Service Management Tasks on page 157](#)
 - Epip: [Epip](#)
 - [Modifying Epipe Service Parameters on page 158](#)
 - [Disabling an Epipe Service on page 158](#)
 - [Re-Enabling an Epipe Service on page 159](#)
 - [Deleting an Epipe Service on page 159](#)

Basic Configurations

- [Creating an Epipe Service on page 130](#)
 - [Using Spoke SDP Control Words on page 141](#)
 - [Pseudowire Configuration Notes on page 142](#)
 - [Configuring Two VLL Paths Terminating on T-PE2 on page 143](#)
 - [Configuring VLL Resilience on page 146](#)
 - [Configuring VLL Resilience for a Switched Pseudowire Path on page 147](#)
-

Common Configuration Tasks

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

- Associate the service with a customer ID.
- Define SAP parameters
 - Optional - select egress and ingress QoS and/or scheduler policies (configured in the **config>qos** context).
 - Optional - select accounting policy (configured in the **config>log** context).
- Define spoke SDP parameters.
- Enable the service.

Configuring VLL Components

This section provides VLL configuration examples for the VLL services:

- [Creating an Epipe Service on page 130](#)
 - [Configuring Epipe SAP Parameters on page 131](#)
 - [Local Epipe SAPs on page 132](#)
 - [Distributed Epipe SAPs on page 134](#)
 - [Configuring Ingress and Egress SAP Parameters on page 137](#)

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 displays an Epipe configuration example:

```
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#
```

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 on page 132](#)
- [Distributed Epipe SAPs on page 134](#)

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`

Local Epipe SAPs

Table 5: 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)

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.

An existing scheduler policy can be applied to ingress and egress SAPs to be used by the SAP queues. The schedulers comprising the policy are created at the time the scheduler policy is applied to the SAP. If any orphaned queues (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.

This example displays 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 displays 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#
```

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. The schedulers comprising the policy are created at the time the scheduler policy is applied to the SAP. If any orphaned queues (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, see the *Services Overview Guide*. For SDP binding information, see [Configuring SDP Bindings on page 138](#).

This example configures 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 displays 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 displays 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. The schedulers comprising the policy are created at the time the scheduler policy is applied to the SAP. If any orphaned queues (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.

This example displays SAP ingress and egress parameters.

```
ALA-1>config>service# epipe 5500
config>service>epipe# sap /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 displays 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 /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#
```

Configuring SDP Bindings

Figure 37 displays 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.

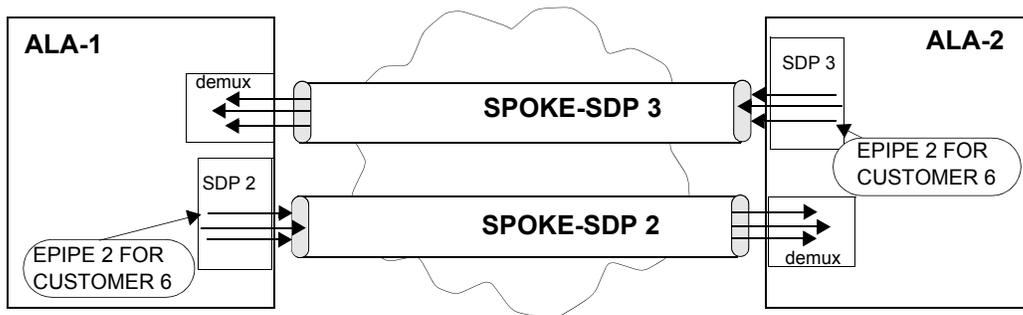


Figure 37: SDPs — Uni-Directional Tunnels

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 displays 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 on page 134](#)).

```
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
```

This example displays 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 /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
```

```
-----
```

Configuring VLL Components

```
...
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#
```

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. The control word might be needed because when ECMP is enabled on the network, packets of a given 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, then the Router Alert label is used. In that case, VCCV CC type 2 is signaled. Note that for a multi-segment pseudowire (MS-PW), the CC type 1 is the only supported and thus the control word must be enabled on the spoke-sdp to be able to use VCCV-ping and VCCV-trace.

The following displays a spoke SDP control word configuration example:

```
-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#
```

Pseudowire Configuration Notes

The **vc-switching** parameter must be specified at the time the VLL service is created. Note that 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 show 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.

CLI Syntax: config>service# epipe service-id [customer customer-id] [vpn
vpn-id] [vc-switching]
description description-string
spoke-sdp sdp-id:vc-id

The following displays an example of the command usage to configure VLL pseudowire switching services:

Configuring Two VLL Paths Terminating on T-PE2

T-PE1

The following displays 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 displays an example of the T-PE2 configuration.

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#
```

S-PE1: Note that 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 displays the configuration:

```
*A:ALA-S-PE1>config>service>epipe# info
```

Configuring Two VLL Paths Terminating on T-PE2

```
-----  
...  
    spoke-sdp 2:200 create  
    exit  
    spoke-sdp 3:300 create  
    exit  
    no shutdown  
-----  
*A:ALA-S-PE1>config>service>epipe#
```

S-PE2: Note that 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 displays 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#
```

Configuring VLL Resilience

Figure 38 displays an example to create VLL resilience. Note that the zero revert-time value means that the VLL path will be switched back to the primary immediately after it comes back up.

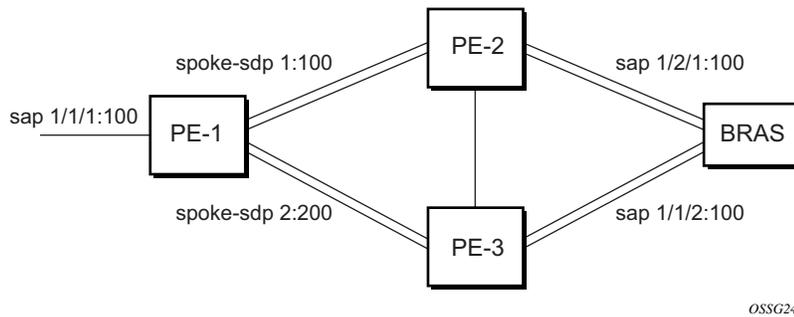


Figure 38: VLL Resilience

PE1:

The following displays an example for the configuration on PE1.

```

*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#
  
```

Configuring VLL Resilience for a Switched Pseudowire Path

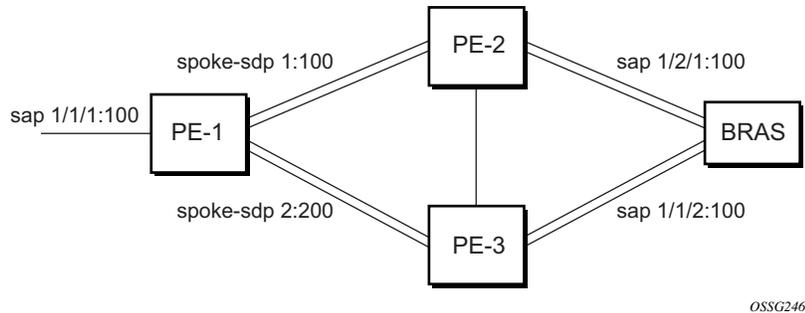


Figure 39: VLL Resilience with Pseudowire Switching

T-PE1

The following displays an example for the configuration on TPE1.

```

*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-PE2

The following displays an example for the configuration on TPE2.

```
*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-PE1

The following displays an example for the configuration on S-PE1.

```
*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#
```

Configuring BGP Virtual Private Wire Service (VPWS)

Single-Homed BGP VPWS

Figure 40 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.

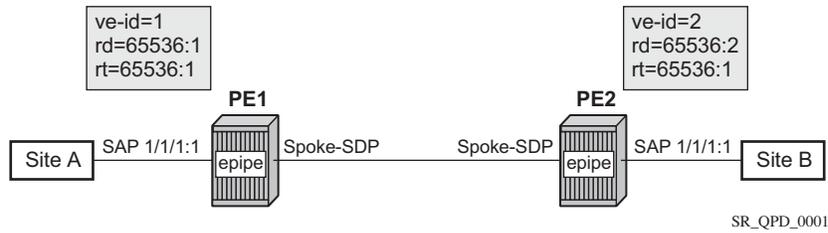


Figure 40: Single-Homed BGP VPWS Configuration Example

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 encaps, 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 displays 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
```

Configuring BGP Virtual Private Wire Service (VPWS)

```
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#
```

Dual-Homed BGP VPWS

Single Pseudowire Example:

Figure 41 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.

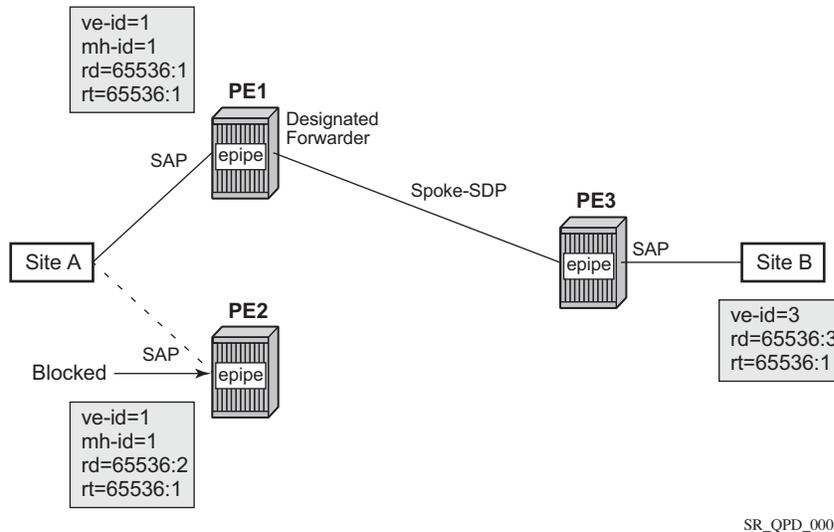


Figure 41: Example of Dual-Homed BGP VPWS with Single Pseudowire

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 LSPs 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 displays 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
exit

```

Configuring BGP Virtual Private Wire Service (VPWS)

```
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
```

Configuring BGP Virtual Private Wire Service (VPWS)

Active/Standby Pseudowire Example:

Figure 42 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.

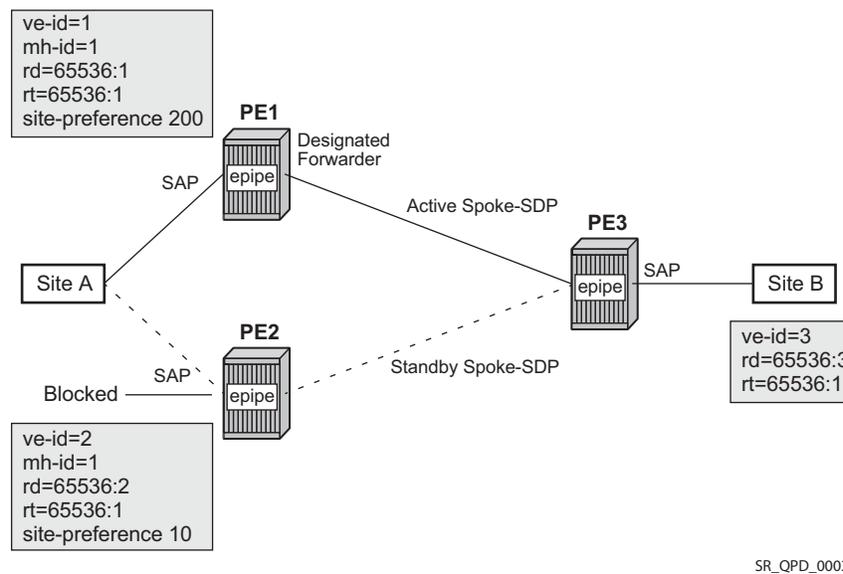


Figure 42: Example of Dual-homed BGP VPWS with Active/Standby Pseudowires

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 LSPs 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 displays 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

```

Configuring BGP Virtual Private Wire Service (VPWS)

```
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
```

Service Management Tasks

This section discusses the following Epipe service management tasks:

- [Modifying Epipe Service Parameters on page 158](#)
- [Disabling an Epipe Service on page 158](#)
- [Re-Enabling an Epipe Service on page 159](#)
- [Deleting an Epipe Service on page 159](#)

Modifying Epipe Service Parameters

The following displays an example of adding an accounting policy to an existing SAP:

```
Example:config>service# epipe 2
        config>service>epipe# sap /1/3:21
        config>service>epipe>sap# accounting-policy 14
        config>service>epipe>sap# exit
```

The following output displays the SAP configuration:

```
ALA-1>config>service# info
-----
      epipe 2 customer 6 vpn 2 create
          description "Distributed Epipe service to east coast"
          sap /1/3:21 create
              accounting-policy 14
          exit
      spoke-sdp 2:6000 create
      exit
      no shutdown
      exit
-----
ALA-1>config>service#
```

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
```

Re-Enabling an Epipe Service

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

Deleting an Epipe Service

Perform the following steps prior to deleting an Epipe service:

1. Shut down the SAP and SDP.
2. Delete the SAP and SDP.
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 /1/3:21
config>service>epipe>sap# shutdown
config>service>epipe>sap# exit
config>service>epipe# no sap /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

VLL Services Command Reference

Command Hierarchies

- [Apipe Service Configuration Commands on page 161](#)
- [Cpipe Service Configuration Commands on page 165](#)
- [Epipe Service Configuration Commands on page 168](#)
- [Ipipe Service Configuration Commands on page 178](#)

Apipe Service Configuration Commands

```

config
  — service
    — apipe service-id [customer customer-id] [vpn vpn-id] [vc-type {atm-vcc | atm-sdu | atm-vpc
      | atm-cell}] [vc-switching] [test] [create]
    — 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**
 - **rate** {*max* | *rate*}
 - **no rate**
 - [no] **limit-unused-bandwidth**
 - [no] **queue-frame-based-accounting**
 - **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** *size-in-kbytes*
 - **no burst-limit**
 - **high-prio-only** *percent*
 - **no high-prio-only**
 - **mbs** {*size-in-kbytes* | **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**
- **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**]
 - **no mbs**
 - **packet-byte-offset** **add** *add-bytes* | **subtract** *sub-bytes*}
 - **percent-rate** *pir-percent* [**cir** *cir-percent*]
 - **b 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**
 - **burst-limit** *size-in-kbytes*
 - **no burst-limit**
 - **high-prio-only** *percent*
 - **no high-prio-only**
 - **mbs** {*size-in-kbytes* | **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**
- **tod-suite** *tod-suite-name*
 - **no tod-suite**
- **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**
 - **qos** *network-policy-id* **port-redirect-group** *queue-group-name* [*instance instance-id*]
 - **no qos**
- **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] **shutdown**

config

- **connection-profile** *conn-prof-id* [**create**]
- **no connection-profile** *conn-prof-id*

Related Apipe Commands

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

```

### Cpipe Service Configuration Commands

- ```

config
  — service
    — cpipe service-id [customer customer-id] [vpn vpn-id] [vc-type {satop-e1 | satop-t1 | satop-e3 | satop-t3 | cesopsn | cesopsn-cas}] [vc-switching] [test] [create]
    — 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>]
        — [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 {max | rate}
            — no rate
            — [no] limit-unused-bandwidth
            — [no] queue-frame-based-accounting
          — [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** *size-in-kbytes*
- **no burst-limit**
- **high-prio-only** *percent*
- **no high-prio-only**
- **mbs** *size-in-kbytes*
- **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**
- **ingress**
 - [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** *size-in-kbytes*
 - **no burst-limit**
 - **high-prio-only** *percent*
 - **no high-prio-only**
 - **mbs** *size-in-kbytes*
 - **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**

- **tod-suite** *tod-suite-name*
- **no tod-suite**
- **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**
 - **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*]
- **precedence** [*precedence-value*] **primary**
- **no precedence**
- **[no] shutdown**

Epipe Service Configuration Commands

- [Epipe Global Commands on page 168](#)
- [Epipe SAP Configuration Commands on page 170](#)
- [Epipe Spoke SDP Configuration Commands on page 174](#)

Epipe Global Commands

```

config
  — service
    — [no] epipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching]
      — [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
      —
      —
      — tunnel service-id backbone-dest-mac mac-name| ieee-address isid ISID
      — no tunnel
      — 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**
 - **hash-label**
 - **no hash-label**
 - [**no**] **standby-signaling-slave**

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
        — accounting-policy acct-policy-id
        — no accounting-policy acct-policy-id
        — [no] cflowd
        — [no] collect-stats
        — description description-string
        — no description
        — 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]
            — 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 policyer-id [create]
            — no policer policyer-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
            — 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**
- **cbs** *size-in-kbytes*
- **no cbs**
- **high-prio-only** *percent*
- **no high-prio-only**
- **mbs** *size* [**bytes**|**kilobytes**]
- **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**
 - [**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**}
 - **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
 - **mac-address** *mac-address*

- **no mac-address**
- **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**]
- **ethernet**
 - **[no] llf**
- **[no] ignore-oper-down**
- **ingress**
 - **agg-rate-limit** *agg-rate*
 - **no agg-rate-limit**
 - **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**
 - **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** *policyer-id* [**create**]
 - **no policer** *policyer-id*
 - **cbs** *size-in-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**
 - **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**
- **high-prio-only** *percent*
- **no high-prio-only**
- **mbs** *size* [*bytes* | *kilobytes*]
- **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..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**
- **tod-suite** *tod-suite-name*
- **no tod-suite**
- **transit-policy** **prefix** *prefix-aasub-policy-id*
- **no transit-policy**

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]
        — accounting-policy acct-policy-id
        — no accounting-policy
        — bandwidth bandwidth
        — no bandwidth
        — [no] bfd-enable
        — bfd-template name
        — no bfd-template
        — [no] collect-stats
        — [no] control-word
        — [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
      — 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
        — [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}]
        — mep mep-id domain md-index association ma-index [direction
          {up | down}]
        — 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*
- **noccm-padding-size** *ccm-padding*
- **fault-propagation-enable** {**use-if-tlv** | **suspend-ccm**}
- **no fault-propagation-enable**
- **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
- **mac-address** *mac-address*
- **no mac-address**
- **[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**
- **[no] use-sdp-bmac**
- **vlan-vc-tag** *0..4094*
- **no vlan-vc-tag** [*0..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*

Command Hierarchies

- **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**

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

```

Ipipe Service Configuration Commands

```

config
  — service
    — 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
        — agg-rate-limit agg-rate
        — no agg-rate-limit
        — [no] agg-rate
          — rate {max | rate}
          — 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] 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** *size-in-kbytes*
 - **no burst-limit**
 - **high-prio-only** *percent*
 - **no high-prio-only**
 - **mbs** {*size-in-kbytes* | **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**
- **eth-cfm**
 - **[no] collect-lmm-stats**
 - **[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**}
 - **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
 - **[no] mac-address** *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**
 - **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**
 - **burst-limit** *size-in-kbytes*
 - **no burst-limit**
 - **high-prio-only** *percent*
 - **no high-prio-only**
 - **mbs** {*size-in-kbytes* | **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..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**
- **tod-suite** *tod-suite-name*
- **no tod-suite**
- **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*
 - **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**
 - **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**
- **hash-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**

VLL Service Configuration Commands

- [Generic Commands on page 183](#)
- [VLL Global Commands on page 188](#)
- [VLL SAP Commands on page 203](#)
- [VLL SDP Commands on page 268](#)

Generic Commands

shutdown

Syntax	<code>[no] shutdown</code>
Context	<pre>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</pre>
Description	<p>This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.</p> <p>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.</p> <p>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.</p> <p>The no form of this command places the entity into an administratively enabled state.</p>
Special Cases	<p>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.</p> <p>Service Operational State — A service is regarded as operational providing that at least one SAP and one SDP are operational or if two SAP's are operational.</p> <p>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.</p> <p>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.</p>

description

Syntax	description <i>description-string</i> no description
Context	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
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	<i>string</i> — 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.

Service Commands

apipe

Syntax	apipe <i>service-id</i> [customer <i>customer-id</i>] [vpn <i>vpn-id</i>] [vc-type { atm-vcc atm-sdu atm-vpc atm-cell }] [vc-switching] no apipe <i>service-id</i>
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	<p><i>service-id</i> — 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 <i>service-id</i> must be the same number used for every or on which this service is defined.</p> <p>Values <i>service-id:</i> 1 — 2147483648 <i>svc-name:</i> 64 characters maximum</p> <p>customer <i>customer-id</i> — 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.</p> <p>Values 1 — 2147483647</p> <p>vpn <i>vpn-id</i> — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.</p> <p>Values 1 — 2147483647</p> <p>Default null (0)</p> <p>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 <i>draft-ietf-pwe3-iana-allocation</i> and it defines both the signaled VC type as well as the resulting datapath encapsulation over the Apipe.</p> <p>Values atm-vcc, atm-sdu, atm-vpc, atm-cell</p> <p>Default atm-sdu</p> <p>vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.</p>

cpipe

Syntax	cpipe <i>service-id</i> [customer <i>customer-id</i>] [vpn <i>vpn-id</i>] [vc-type { satop-e1 satop-t1 [vc-switching] [cesopsn cesopsn-cas]}] [vc-switching] [test] [create] no cpipe <i>service-id</i>
---------------	--

Service Commands

Context	config>service
Description	<p>This command configures a Circuit Emulation Services instance. When a service is created, the customer keyword and <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> 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.</p> <p>Once a service is created, the use of the customer <i>customer-id</i> is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect <i>customer-id</i> specified will result in an error.</p> <p>By default, no services exist until they are explicitly created with this command.</p> <p>The no form of this command deletes the service instance with the specified <i>service-id</i>. The service cannot be deleted until the service has been shutdown.</p>
Parameters	<p><i>service-id</i> — 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 <i>service-id</i> must be the same number used for every router on which this service is defined.</p> <p>Values <i>service-id:</i> 1 — 2147483648 <i>svc-name:</i> Specifies an existing service name up to 64 characters in length.</p> <p>customer <i>customer-id</i> — 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.</p> <p>Values 1 — 2147483647</p> <p>vpn <i>vpn-id</i> — 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.</p> <p>Values 1 — 2147483647</p> <p>Default null (0)</p> <p>vc-type — The vc-type defines the type of unstructured or structured circuit emulation service to be configured.</p> <p>Values satop-e1: unstructured E1 circuit emulation service satop-t1: unstructured DS1 circuit emulation service cesopsn: basic structured n*64 kbps circuit emulation service cesopsn-cas: structured n*64 kbps circuit emulation service with signaling</p>

epipe

Syntax	epipe <i>service-id</i> customer <i>customer-id</i> [vpn <i>vpn-id</i>] [vc-switching] [create] epipe <i>service-id</i> no epipe <i>service-id</i>
Context	config>service

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 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. 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.

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.

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 , , or 7950 XRS on which this service is defined.

Values *service-id:* 1 — 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 — 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 — 2147483647

Default **null** (0)

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

create — Keyword used to create the service instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

or

VLL Global Commands

bgp

Syntax	bgp
Context	config>service>epipe
Description	This command enables 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 <i>policy-id</i> [import-rt { <i>ext-community</i> ,.(upto 5 max)}}] no pw-template-binding <i>policy-id</i>
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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> – 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 oper state down. – 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	<i>policy-id</i> — Specifies an existing policy ID.

Values 1 — 2147483647

import-rt ext-comm — Specify 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-asnumber:comm-val}
 ip-addr a.b.c.d
 comm-val 0 — 65535
 2byte-asnumber 0 — 65535
 ext-comm-val 0 — 4294967295
 4byte-asnumber 0 — 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. <p>An RD value must be configured under BGP node.</p> <p>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.</p> <p>Format: Six bytes, other 2 bytes of type will be automatically generated.</p>
Parameters	<i>ip-addr:comm-val</i> — Specifies the IP address. <p>Values ip-addr a.b.c.d comm-val 0 — 65535 as-number:</p> <p><i>as-number:ext-comm-val</i> — Specifies the AS number. <p>Values as-number 1 — 65535 ext-comm-val 0 — 4294967295</p> <p>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.</p> <p><i>rd</i> — Specifies the route distinguisher. <p>Values <rd> <ip-addr:comm-val> <2byte-asnumber:ext-comm-val> <4byte-asnumber:comm-val> ip-addr a.b.c.d comm-val [0..65535] 2byte-asnumber [1..65535] ext-comm-val [0..4294967295] 4byte-asnumber [0..4294967295]</p></p></p>

route-target

Syntax	route-target { <i>ext-community</i> }{ [export <i>ext-community</i>] [import <i>ext-community</i>]} 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 <i>ext-comm</i> can have two formats: <ul style="list-style-type: none"> • 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	<i>export ext-community</i> — Specifies communities allowed to be sent to remote PE neighbors. <i>import ext-community</i> — Specifies communities allowed to be accepted from remote PE neighbors.

bgp-vpws

Syntax	[no] bgp-vpws
Context	config>service>epipe
Description	This command enables the context to configure BGP-VPWS parameters and addressing.
Default	no bgp-vpws

remote-ve-name

Syntax	[no] remote-ve-name <i>name</i>
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	<i>name</i> — Specifies a site name up to 32 characters in length.

ve-id

Syntax	ve-id <i>value</i>
---------------	---------------------------

no ve-id

Context	config>service>epipe>bgp-vpws>ve-name config>service>epipe>bgp-vpws>remote-ve-name
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>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 L2 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.</p> <p>Each ve-id configured within a service must be unique.</p> <p>The no form of the command removes the configured ve-id. It can be used just when the BGP VPWS status is shutdown. Command “no shutdown” cannot be used if there is no ve-id configured.</p>
Default	no ve-id
Parameters	<p><i>value</i> — A two bytes identifier that represents the local or remote VPWS instance and is advertised through the BGP NLRI.</p> <p>Values 1 — 65535</p>

ve-name

[no] ve-name name

Context	config>service>epipe>bgp-vpws
Description	<p>This command configures the name of the local VPWS instance in this service.</p> <p>The no form of the command removes the ve-name.</p>
Parameters	<i>name</i> — 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 — 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 <i>seconds</i> 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	<i>seconds</i> — Specifies the site activation timer in seconds. Values 0 — 100

site-min-down-timer

Syntax	site-min-down-timer <i>min-down-time</i> 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 above operation is optimized in the following circumstances: <ul style="list-style-type: none"> • 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 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 configure>redundancy>bgp-multi-homing context.
Parameters	<i>min-down-time</i> — Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down. Values 0 — 100 seconds

site-id

Syntax	site-id <i>value</i> 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	<i>value</i> — Specifies the site identifier. Values 1 — 65535

site-preference

Syntax	site-preference <i>preference-value</i> 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	<i>preference-value</i> — Specifies the preference value to advertise in the NLRI L2 extended community for this site. Values 1 — 65535
Parameters	primary — Sets the site-preference to 65535. backup — Sets the site-preference to 1.

endpoint

Syntax	[no] endpoint <i>endpoint-name</i>
Context	config>service>apipe config>service>epipe config>service>ipipe
Description	This command configures a service endpoint.
Parameters	<i>endpoint-name</i> — Specifies an endpoint name.

tunnel

Syntax	tunnel <i>service-id</i> backbone-dest-mac <i>ieee-address</i> isid <i>ISID</i> 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	<i>service-id</i> — Specifies the B-VPLS service for the PBB tunnel associated with this service. Values <i>service-id:</i> 1 — 2147483648 <i>svc-name:</i> 64 characters maximum backbone-dest-mac <i>ieee-address</i> — Specifies the backbone destination MAC-address for PBB packets. isid <i>ISID</i> — 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 — 16777215

active-hold-delay

Syntax	active-hold-delay <i>active-hold-delay</i> no active-hold-delay
Context	config>service>apipe>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 — 60

revert-time

Syntax	revert-time [<i>revert-time</i> infinite] no revert-time
Context	config>service>apipe>endpoint config>service>epipe>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	<i>revert-time</i> — Specify the time, in seconds, to wait before reverting to the primary SDP. Values 0 — 600 Values 0 <i>infinite</i> — Causes the endpoint to be non-revertive.
Syntax	

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 <i>timer-value</i> 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	<i>timer-value</i> — The value of the wait time in tenths of a second (100ms). Granularity is in 500ms increments, starting from 1s and up to 30 seconds.

Values [10 .. 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 given 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.

VLL Global Commands

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 the following table.

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 — Specify Frame Relay to ATM Network Interworking (FRF.5).

service-name

Syntax **service-name service-name**
no service-name

Context config>service>epipe

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given 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 given service once it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

service-mtu

Syntax **service-mtu octets**
no service-mtu

Context
config>service>epipe

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 epipe: 1514

The following table displays MTU values for specific VC types.

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

VLL Global Commands

VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (Q-in-Q with preserved bottom Qtag)	1518	1504

octets — The size of the MTU in octets, expressed as a decimal integer, between 1 — 9194.

signaled-vc-type-override

Syntax	signaled-vc-type-override { atm-vcc } no signaled-vc-type-override
Context	<root>
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of this command returns the Apipe VLL service to signal its default pseudowire type</p>
Default	none
Parameters	atm-vcc — Specifies the pseudowire type to be signaled in the pseudowire establishment.

connection-profile

Syntax	connection-profile <i>conn-prof-id</i> [create] no connection-profile <i>conn-prof-id</i>
Context	<root>
Description	<p>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.</p> <p>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.</p> <p>A maximum of 8000 connection profiles can be created on the system.</p> <p>The no form of this command deletes the profile from the configuration.</p>
Default	none
Parameters	<i>conn-prof-id</i> — Specifies the profile number.
Values	1 — 8000

member

Syntax	member <i>encap-value</i> [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 — 4095; UNI: 0 — 255
vci: 1, 2, 5 — 65535

VLL SAP Commands

sap

Syntax	<pre> sap <i>sap-id</i> [create] [no-endpoint] sap <i>sap-id</i> [create] endpoint <i>endpoint-name</i> no sap <i>sap-id</i> </pre>
Context	<pre> config>service>apipe config>service>epipe </pre>
Description	<p>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.</p> <p>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.</p> <p>Enter an existing SAP without the create keyword to edit SAP parameters. The SAP is owned by the service in which it was created.</p> <p>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.</p> <p>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.</p> <p>The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.</p> <p>The following are supported:</p> <ul style="list-style-type: none"> • Ethernet SAPs support null, dot1q, and qinq <p>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.</p>
Default	No SAPs are defined.
Special Cases	<p>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.</p> <p>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).</p> <p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP. See Common CLI Command Descriptions on page 1271 for command syntax.</p>

port-id — Specifies the physical port ID in the *slot/mda/port* format.

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.

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.

sap

Syntax	[no] sap eth-tunnel-tunnel-id[:eth-tunnel-sap-id] [create]
Context	config>service>epipe config>service>vpls
Description	<p>This command configures an Ethernet tunnel SAP.</p> <p>An Ethernet tunnel control SAP has the format <i>eth-tunnel-tunnel-id</i> 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.</p> <p>An Ethernet tunnel same-fate SAP has the format <i>eth-tunnel-tunnel-id:eth-tunnel-sap-id</i>. 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.</p> <p>Ethernet tunnel SAPs are supported under VPLS, Epipe and Ipipe services only.</p>
Default	no sap
Parameters	<p><i>tunnel-id</i> — Specifies the tunnel ID.</p> <p>Values 1 — 1024</p> <p><i>eth-tunnel-sap-id</i> — Specifies a SAP ID of a same-fate SAP.</p> <p>Values 0 — 4094</p>

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	<i>link-map-profile-id</i> — 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..1024] no per-link-hash
Context	config>service>epipe>sap config>service>ipipe>sap config>service>vpls>sap config>service>ies>if>sap config>service>vprn>if>sap config>service>ies>sub-if>grp-if>sap config>service>vprn>sub-if>grp-if>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)

monitor-oper-group

Syntax	monitor-oper-group group-name no monitor-oper-group
Context	config>service>if config>service>ies>spoke-sdp config>service>ies>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.

agg-rate-limit

Syntax	agg-rate-limit <i>agg-rate</i> no agg-rate-limit
Context	config>service>epipe>sap>ingress
Description	<p>This command defines a maximum total rate for all egress queues on a service SAP or multi-service site. The agg-rate-limit command is mutually exclusive with the egress scheduler policy. When an egress scheduler policy is defined, the agg-rate-limit command will fail. If the agg-rate-limit command is specified, an attempt to bind a scheduler-policy to the SAP or multi-service site will fail.</p> <p>A multi-service site must have a port scope defined that ensures all queues associated with the site are on the same port or channel. If the scope is not set to a port, the agg-rate-limit command will fail. Once an agg-rate-limit has been assigned to a multi-service site, the scope cannot be changed to card level.</p> <p>A port scheduler policy must be applied on the egress port or channel the SAP or multi-service site are bound to in order for the defined agg-rate-limit to take effect. The egress port scheduler enforces the aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.</p> <p>If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.</p> <p>The no form of the command removes the aggregate rate limit from the SAP or multi-service site.</p>
Parameters	<p><i>agg-rate</i> — Defines the rate, in kilobits-per-second, that the maximum aggregate rate the queues on the SAP or MSS can operate.</p> <p>Values 1 — 40000000, max</p>

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	<p>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.</p>

rate

Syntax	rate { max rate } 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 etc.).

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 queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports.

policer-control-override

Syntax `policer-control-override [create]`
`no policer-control-override`

Context `config>service>epipe>sap>egress`

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.

VLL SAP Commands

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 { <i>rate</i> max }
Context	config>service>epipe>sap>egress>policer-control-override
Description	<p>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.</p> <p>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.</p>
Parameters	<i>rate</i> max — Specifies the max rate override in kilobits-per-second or use the maximum.
Values	1 — 2000000000 Kbps, max

priority-mbs-thresholds

Syntax	priority-mbs-thresholds
Context	config>service>epipe>sap>egress>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 <i>size</i> [bytes kilobytes]
Context	config>service>epipe>sap>egress>policer-control-override>priority-mbs-threshold
Description	<p>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.</p> <p>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.</p> <p>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.</p>

Default	no min-thresh-separation
Parameters	<p>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.</p> <p>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.</p> <p>Values 0 – 16777216 or default</p> <p>Default kilobytes</p>

priority

Syntax	[no] priority level
Context	config>service>epipe>sap>egress>policer-control-override>priority-mbs-thresholds
Description	<p>The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.</p> <p>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.</p>
Parameters	<p>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.</p> <p>Values 1 — 8</p>

mbs-contribution

Syntax	mbs-contribution size [bytes kilobytes]
Context	config>service>epipe>sap>egress>policer-control-override>priority-mbs-threshold>priority
Description	<p>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.</p> <p>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.</p> <p>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.</p>
Default	no mbs-contribution
Parameters	bytes — This keyword signifies that size is expressed in bytes.

kilobytes — The optional kilobytes keyword signifies that size is expressed in kilobytes.

Values 0 – 16777216 or default

policer-control-policy

Syntax	policer-control-policy <i>policy-name</i> [create] no policer-control-policy
Context	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.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP). 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.

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 thus 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 (in-profile / out-of-profile) 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.

In order 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 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>epipe>sap>egress

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 <i>policer-id</i> [create] no policer <i>policer-id</i>
Context	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	<i>policer-id</i> — 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 <i>size-in-kbytes</i> no cbs
Context	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	<i>size-in-kbytes</i> — 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 – 16777216 or default <i>kilobytes</i> — When kilobytes is defined, the value given for size is interpreted as the policer's CBS value given in kilobytes.

mbs

Syntax	mbs <i>size</i> [bytes kilobytes] no mbs
Context	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? 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 – 16777216 or default kilobytes — When kilobytes is defined, the value given for size is interpreted as the policer’s MBS value given in kilobytes.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> }
Context	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. The no packet-byte-offset command is used to restore the policer? packet-byte-offset setting to the policy defined value.
Default	no packet-byte-offset
Parameters	add <i>add-bytes</i> — 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 — 31 subtract <i>sub-bytes</i> — 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 — 64

percent-rate

Syntax	percent-rate <i>pir-percent</i> [<i>cir cir-percent</i>] no percent-rate
Context	config>service>epipe>sap>egress>policer-override>policer
Description	This command configures the percent rates (CIR and PIR) override.
Parameters	<i>pir-rate</i> — The <i>pir-percent</i> parameter is used to express the policer's PIR as a percentage of the policer's parent arbiter rate. Values Percentage ranging from 0.01 to 100.00. The default is 100.00.
	<i>cir cir-rate</i> — Configures the administrative CIR specified by the user. Values 0 — 20000000, max

percent-rate

Syntax	percent-rate <i>pir-percent</i> [<i>cir cir-percent</i>] [no percent-rate
Context	config>service>epipe>sap>egress>queue-override>queue
Description	The <code>percent-rate</code> 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-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit 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. Values 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 (kbps) 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* — The pir-percent parameter is used to express the queue's PIR as a percentage dependant on the use of the port-limit or local-limit.
- Values** Percentage ranging from 0.01 to 100.00. The default is 100.00.
- pir-percent* — The pir-percent parameter is used to express the queue's PIR as a percentage dependant on the use of the port-limit or local-limit.
- cir** *cir-percent* — The cir keyword is optional and when defined the required cir-percent CIR parameter expresses the queue's CIR as a percentage dependant on the use of the port-limit or local-limit.
- Percentage ranging from 0.00 to 100.00. The default is 100.00

rate

- Syntax** `rate {rate | max} [cir {max | rate}]`
- Context** 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 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** | **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 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 — 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 — 2000000000, **max**

stat-mode

Syntax	stat-mode <i>stat-mode</i> no stat-mode
Context	config>service>epipe>sap>egress>policer-override>policer
Description	<p>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. An egress policer has multiple types of offered packets (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.</p> <p>While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires at 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. Once a policer has been made a child to a parent policer, the stat-mode cannot be changed to no-stats unless the policer parenting is first removed.</p> <p>Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero.</p> <p>Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. 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.</p> <p>The default stat-mode when a policer is created within the policy is no-stats.</p> <p>The stat-mode setting defined for the policer in the QoS policy may be overridden on an sla-profile or 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.</p> <p>The no stat-mode command attempts to return the policer's stat-mode setting to no-stats. The command will fail if the policer is currently configured as a child policer using the policer's parent command. The no parent command must first be executed for the no stat-mode command to succeed.</p>
Parameters	<p><i>stat-mode</i> — Specifies the mode of statistics collected by this policer.</p> <p>Values no-stats, minimal, offered-profile-no-cir, offered-profile-cir, offered-total-cir</p> <p>no-stats — Counter resource allocation: 0</p> <p>The no-stats mode is the default stat-mode for the policer. The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policers parent command will fail.</p>

When collect-stats is enabled, the lack of counters causes the system to generate the following statistics:

a. offered-in	= 0
b. offered-out	= 0
c. discard-in	= 0
d. discard-out	= 0
e. forward-in	= 0
f. forward-out	= 0

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation: 1 The minimal mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (soft or hard profile) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

1. offered	<= soft-in-profile-out-of-profile, profile in/out
2. discarded	<= Same as 1
3. forwarded	<= Derived from 1 – 2

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a. offered-in	= 1
b. offered-out	= 0
c. discard-in	= 2
d. discard-out	= 0
e. forward-in	= 3
f. forward-out	= 0

Counter 0 indicates that the accounting statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation: 2

The offered-profile-no-cir mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The offered-profile-no-cir mode is most useful when profile based offered, discard and forwarding stats are required from the ingress policer, but a CIR is not being used to recolor the soft in-profile and out-of-profile packets. This mode does not prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-in	<= soft-in-profile, profile in
2. offered-out	<= soft-out-of-profile, profile out
3. dropped-in	<= Same as 1
4. dropped-out	<= Same as 2
5. forwarded-in	<= Derived from 1 – 3
6. forwarded-out	<= Derived from 2 – 4

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

offered-profile-cir — Counter resource allocation: 3

The offered-profile-cir mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The offered-profile-cir mode is most useful when profile based offered, discard and forwarding stats are required from the ingress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

- 1. offered-in-that-stayed-green-or-turned-red <= profile in
- 2. offered-soft-that-turned-green <= soft-in-profile-out-of-profile
- 3. offered-soft-or-out-that-turned-yellow-or-red <= soft-in-profile-out-of-profile, profile out
- 4. dropped-in-that-stayed-green-or-turned-red <= Same as 1
- 5. dropped-soft-that-turned-green <= Same as 2
- 6. dropped-soft-or-out-that-turned-yellow-or-red <= Same as 3
- 7. forwarded-in-that-stayed-green <= Derived from 1 – 4
- 8. forwarded-soft-that-turned-green <= Derived from 2 – 5
- 9. forwarded-soft-or-out-that-turned-yellow <= Derived from 3 – 6

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2 + 3
- c. discard-in = 4
- d. discard-out = 5 + 6
- e. forward-in = 7 + 8
- f. forward-out = 9

offered-total-cir — Counter resource allocation: 2

The offered-total-cir mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The offered-total-cir mode is most useful when profile based offered stats are not required from the ingress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

1. offered-that-turned-green <= soft-in-profile-out-of-profile, profile in/out
2. offered- that-turned-yellow-or-red<= soft-in-profile-out-of-profile, profile in/out
3. dropped-offered-that-turned-green<= Same as 1
4. dropped-offered-that-turned-yellow-or-red<= Same as 2
5. forwarded-offered-that-turned-green<= Derived from 1 – 3
6. forwarded-offered-that-turned-yellow<= Derived from 2 – 4

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- | | |
|----------------|---|
| a. offered-in | = 1 + 2 (Or 1 and 2 could be summed on b) |
| b. offered-out | = 0 |
| c. discard-in | = 3 |
| d. discard-out | = 4 |
| e. forward-in | = 5 |
| f. forward-out | = 6 |

Counter 0 indicates that the accounting statistic returns a value of zero.

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.

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>epipe>sap>egress

VLL SAP Commands

- 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>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
- 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

tod-suite

- Syntax** **tod-suite** *tod-suite-name*
no tod-suite

Context	config>service>epipe>sap
Description	This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the config>system>cron context.
Default	no tod-suite
Parameters	<i>tod-suite-name</i> — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

transit-policy

Syntax	transit-policy prefix <i>prefix-aasub-policy-id</i> no transit-policy
Context	config>service>epipe>sap
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	<i>prefix-aasub-policy-id</i> — Specifies the transit policy ID. Values 1 — 65535

mac

Syntax	[no] mac <i>ieee-address</i>
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	<i>ieee-address</i> — 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 <i>refresh interval</i> no mac-refresh
Context	config>service>ipipe>sap

VLL SAP Commands

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	<i>refresh interval</i> — Specifies the interval, in seconds, between ARP requests sent on this Ipipe SAP. Values 0 — 65535

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>service>apipe>sap config>service>epipe>sap 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 <i>policy-id</i> 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	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context. Values 1 — 99

bandwidth

Syntax	bandwidth <i>bandwidth</i> no bandwidth
Context	config>service>epipe>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 an 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 operational UP state until the user performs a shutdown/no-shutdown of the spoke-sdp and the bandwidth check succeeds. Thus, the service manager will not automatically audit spoke-sdp’s subsequently to their creation to check if bandwidth is available.

If the VLL service contains an endpoint with multiple redundant spoke-sdp’s, 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.

Note 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 LSP’s 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 selected LSP(s) available bandwidth 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.

Values 0 — 100000000, max in units of kilobits/sec.

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 **bfd-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>spoke-sdp
config>service>epipe>spoke-sdp
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

VLL SAP Commands

```
config>service>epipe>spoke-sdp
config>service>epipe>sap
```

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 <i>cpu-protection-policy-id</i> 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 as shown above. If no CPU protection policy is assigned to a service SAP then a the default policy is used to limit the overall-rate.
Parameters	<i>policy-id</i> — Specifies an existing CPU protection policy. Values 1 — 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 <i>policy-name</i> no dist-cpu-protection
Context	config>service>epipe>sap config>service>apipe>sap config>service>cpipe>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 created DCP policy can be assigned to a SAP or a network interface (note that this rule does not apply to templates such as an msap-policy)
Default	no dist-cup-protection

ethernet

Syntax	ethernet
Context	config>service>epipe>sap
Description	Use this command to configure Ethernet properties in this SAP.

llf**Syntax** [no] llf**Context** 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. It brings down the Ethernet port (Ethernet LLF) towards 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.

ETH-CFM Service Commands

eth-cfm

Syntax	eth-cfm
Context	config>service>epipe>spoke-sdp config>service>epipe>sap
Description	This command enables 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-mep-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>ies>interface>sap>>eth-cfm config>service>ies>interface>spoke-sdp>>eth-cfm config>service>ies>subscriber-interface>group-interface>sap>eth-cfm config>service>vprn>interface>sap>eth-cfm config>service>vprn>interface>spoke-sdp>eth-cfm config>service>vprn>subscriber-interface>group-interface>sap>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

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 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 DOWN MEP is configured.

client-meg-level

Syntax	client-meg-level <i>[[/level [/level ...]]</i> 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(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	<i>level</i> — Specifies the client MEG level.
	Values 1 — 7
	Default 1

interval

Syntax	interval {1 60} no interval
Context	config>service>epipe>sap>eth-cfm>mep config>service>epipe>spoke-sdp>eth-cfm>aid-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 <i>priority-value</i> 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	<i>priority-value</i> — Specify the priority value of the AIS messages originated by the node.
	Values 0 — 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 <i>path-index</i> tag <i>qtag</i> [<i>.qtag</i>] no path <i>path-index</i>
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	<i>path-index</i> — Specifies the path index value. Values 1 — 16 <i>tag</i> <i>qtag</i> [<i>.qtag</i>] — Specifies the qtag value. Values 0 — 4094, *

mep

Syntax	mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [direction { down }] no mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> primary-valn-enable [vlan <i>vlan-id</i>]
Context	config>service>epipe>sap>eth-cfm config>service>epipe>spoke-sdp>eth-cfm
Description	This command provisions the maintenance endpoint (MEP). The no form of the command reverts to the default values.
Parameters	<i>mep-id</i> — Specifies the maintenance association end point identifier. Values 1 — 81921 <i>md-index</i> — Specifies the maintenance domain (MD) index value. Values 1 — 4294967295 <i>ma-index</i> — Specifies the MA index value. Values 1 — 4294967295

direction down — Indicates the direction in which the maintenance association (MEP) faces on the bridge port. The UP direction is not supported for all Fpipe services.

down — Sends ETH-CFM messages away from the MAC relay entity.

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 can not 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 — 4094

ccm-enable

Syntax [no] **ccm-enable**

Context

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

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 — 7

ccm-padding-size

Syntax **ccm-padding-size** *ccm-padding*
no ccm-padding-size *ccm-padding*

Context 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>service>ies>if>sap>eth-cfm>mep>
config>service>ies>if>spoke-sdp>eth-cfm>mep
config>service>vprn>if>sap>eth-cfm>mep
config>service>vprn>if>spoke-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	<i>ccm-padding</i> — specifies the byte size of the Optional Data TLV
Values	3 — 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 <i>multiplier-value</i> 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	<i>multiplier-value</i> — Specifies the multiplier used for timing out CSF.
Values	0.0, 2.0 .. 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>interface>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 TVL on reception.

eth-test-enable

Syntax	[no] eth-test-enable
Context	config>service>epipe>spoke-sdp>eth-cfm>mep config>service>epipe>sap>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 <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [<i>priority priority</i>] [<i>data-length data-length</i>] 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 <i>errors</i> 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.

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>mep>eth-test-enable This command configures the test pattern for eth-test frames. Description 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 This command configures the fault propagation for the MEP. Description
Parameters	use-if-tlv — Specifies to use the interface TLV. suspend-ccm — Specifies to suspend the continuity check messages.

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 This command specifies the lowest priority defect that is allowed to generate a fault alarm.								
Default	macRemErrXcon								
Values	<table> <tr> <td>allDef</td> <td>DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM</td> </tr> <tr> <td>macRemErrXcon</td> <td>Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM</td> </tr> <tr> <td>remErrXcon</td> <td>Only DefRemoteCCM, DefErrorCCM, and DefXconCCM</td> </tr> <tr> <td>errXcon</td> <td>Only DefErrorCCM and DefXconCCM</td> </tr> </table>	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
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>epipe>spoke-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 or the bridge (if the MEP is on a spoke SDP).
 Mac address can not be set on SAPs/Bindings under non-VPLS service.

Parameters *mac-address* — Specifies the MAC address of the MEP.

Values 6-byte unicast mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) of the MEP.
 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 0-600

Default 3

mip

Syntax **mip** [*mac mac-address*] **primary-vlan-enable** [*vlan vlan-id*]
mip default-mac
no mip

Context config>service>epipe>sap>eth-cfm

Description This command allows Maintenance Intermediate Points (MIPs). The creation rules of the MIP are dependant on the mhf-creation configuration for the MA. This MIP option is only available for default and static mhf-creation methods.

Parameters **mac** — provides a method for manually configuring the MIP MAC.
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 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 can not 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 — 4094

Default no mip

snmp-ingress-levels

Syntax **snmp-ingress-levels** [*md-level* [*md-level*...]]
no snmp-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>ies>interface>sap>eth-cfm
config>service>ies>interface>spoke-sdp>eth-cfm
config>service>ies>subscriber-interface>group-interface>sap>eth-cfm
config>service>vprn>interface>sap>eth-cfm
config>service>vprn>interface>spoke-sdp>eth-cfm
config>service>vprn>subscriber-interface>group-interface>sap>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.

VLL SAP Commands

Default	no squelch-ingress-levels
Parameters	<i>md-level</i> — Identifies the level.
Values	[0..7]

tunnel-fault

Syntax	tunnel-fault {accept ignore}
Context	config>service>epipe>eth-cfm config>service>epipe>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.
Parameters	accept — Share fate with the facility tunnel MEP ignore — Do not share fate with the facility tunnel MEP
Default	ignore (Service Level) accept (SAP Level for Epipe and VPLS)

Service Filter and QoS Policy Commands

egress

Syntax	egress
Context	config>service>epipe>spoke-sdp config>service>epipe>sap
Description	This command enables the context to configure egress SAP parameters.

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	<p>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.</p> <p>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/spoke SDP configuration.</p> <p>The Ether type used for both VLAN tags is 0x8100. A different Ether type can be used for the outer VLAN tag by configuring the PW template with use-provisioned-sdps and setting the Ether type using the SDP vlan-vc-etype parameter (this Ether type value is then used for all mesh/spoke SDPs using that SDP).</p> <p>The no version of this command sets default behavior.</p>

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

VLL SAP Commands

- 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** Per default this feature is disabled

ingress

- Syntax** **ingress**
- Context** config>service>epipe>spoke-sdp
config>service>epipe>sap
config>service>epipe>sap
- Description** This command enables 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 [**mac** *mac-filter-id*]
no filter [**ip** *ip-filter-id*]
no filter [**mac** *mac-filter-id*]
- Context** 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.
- Special Cases** **Epipe** — Both MAC and IP filters are supported on an Epipe service SAP.
- Parameters** **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.
Values 1 — 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 — 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 an RX/TX cookie for L2TPv3 spoke-SDPs for EPipe services.

cookie

Syntax	cookie [<i>cookie1</i>] [<i>cookie2</i>] no cookie
Context	config>service>epipe>spoke-sdp>egress>l2tpv3 config>service>epipe>spoke-sdp>ingress>l2tpv3
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>A cookie is not mandatory. An operator may delete an egress cookie or either or both ingress cookies.</p>
Default	no cookie1 cookie2
Parameters	<i>cookie</i> — Specify 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 HSM DA egress and ingress queue overrides.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> } no packet-byte-offset
Context	config>service>epipe>sap>egress>hsmda-queue-over config>service>ipipe>sap>egress>hsmda-queue-over
Description	<p>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.</p> <p>The accounting functions affected include:</p> <ul style="list-style-type: none"> • 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 <p>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.</p> <p>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.</p> <p>As mentioned above, 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.</p> <p>The packet-byte-offset value can be overridden for the HSMDA queue at the SAP or subscriber profile level.</p>

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** *add-bytes* — The **add** keyword is mutually exclusive with the subtract keyword. Either the add or subtract keyword must be specified. The add keyword is used to indicate that the following byte value should be added to the packet for queue and queue group level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 0 — 31

subtract *sub-bytes* — The **subtract** keyword is mutually exclusive with the add keyword. Either the add or subtract keyword must be specified. The subtract keyword is used to indicate that the following byte value should be subtracted from the packet for queue and queue group level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 1 — 64

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>service>epipe>sap>egress>hsmdda-queue-over
config>service>ipipe>sap>egress>hsmdda-queue-over

Description This command, within the QoS policy hsmdda-queue context, is a container for the configuration parameters controlling the behavior of an HSMDDA 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 HSMDDA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDDA 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. HSMDDA 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 weighed group assumes its highest member classis 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 HSMDDA 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 in 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 dependant 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 — 8

rate

Syntax `rate pir-rate`
`no rate`

Context `config>service>epipe>sap>egress>hsmda-queue-over`
`config>service>ipipe>sap>egress>hsmda-queue-over`

Description This command specifies the administrative PIR by the user.

Parameters *pir-rate* — Configures the administrative PIR specified by the user.

Values 1 — 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 *percentage* — Specifies the weight for the HSMDA queue.

Values 1— 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 exist prior to applying the policy name to an HSMDA queue.

secondary-shaper

Syntax	secondary-shaper <i>secondary-shaper-name</i> 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	<i>secondary-shaper-name</i> — Specifies a secondary shaper name up to 32 characters in length.

qos

Syntax	qos <i>policy-id</i> [shared-queuing] [fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i>] no qos
Context	config>service>apipe>sap>ingress config>service>fpipe>sap>ingress config>service>ipipe>sap>ingress config>service>epipe>sap>ingress
Description	<p>This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).</p> <p>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.</p> <p>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.</p> <p>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 given type will return an error.</p> <p>By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.</p> <p>The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.</p>
Default	none
Parameters	<p><i>policy-id</i> — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.</p> <p>Values 1 — 65535</p> <p>shared-queuing — This keyword can only be specified on SAP ingress. The shared-queueing 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.</p>

multipoint-shared — This keyword specifies that this queue-id is for multipoint forwarded traffic only. This queue-id can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint queue generates an error; no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit queue-id parameters.

Default Present (the queue is created as non-multipoint).

Values Multipoint or not present.

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 given 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 *instance-id* — Specifies the instance of the named queue group on the IOM/IMM/XMA ingress forwarding plane.

qos

Syntax	qos <i>policy-id</i> [port-redirect-group <i>queue-group-name</i> instance <i>instance-id</i>] no qos
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	<p>This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).</p> <p>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 <i>policy-id</i> does not exist, an error will be returned.</p> <p>The qos command, when used under the egress context, is used to associate egress QoS policies.</p> <p>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.</p> <p>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 given type will return an error.</p> <p>By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.</p>

VLL SAP Commands

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	<i>policy-id</i> — The egress policy ID to associate with SAP on egress. The policy ID must already exist.
Values	1 — 65535
port-redirect-group	— This keyword associates a SAP egress with an instance of a named queue group template on the egress port of a given IOM/IMM/XMA. The queue-group-name and 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 <i>config>port>ethernet>access>egress</i> .
instance instance-id	— Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.
Values	1 — 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 enables 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 — 32

adaptation-rule

- Syntax** **adaptation-rule** [*pir adaptation-rule*]] [*cir adaptation-rule*]]
no adaptation-rule
- Context** 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 <i>percent</i> no avg-frame-overhead
Context	config>service>epipe>sap>egress>queue-override>queue
Description	<p>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).</p> <p>When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <ul style="list-style-type: none"> • 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.

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	<i>percent</i> — 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 — 100.00

burst-limit

Syntax	burst-limit { default <i>size</i> [byte kilobyte]} no burst-limit
Context	config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue
Description	<p>The <code>queue burst-limit</code> command is used to define 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.</p> <p>The <code>burst-limit</code> command is supported under the <code>sap-ingress</code> and <code>sap-egress</code> QoS policy queues. The command is also supported under the <code>ingress</code> and <code>egress queue-group-templates</code> queues.</p> <p>The no form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying <code>burst-limit default</code> within the QoS policies or queue group templates. When specified within a <code>queue-override</code> queue context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.</p>
Parameters	<p>default — The default parameter is mutually exclusive to specifying an explicit size value. When <code>burst-limit default</code> is executed, the queue is returned to the system default value.</p> <p><i>size</i> — When a numeric value is specified (<i>size</i>), 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 byte qualifier must be added following size.</p> <p>Values 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)</p> <p>Default No default for size, use the default keyword to specify default burst limit</p> <p>byte — The bytes qualifier is used to specify that the value given for size must be interpreted as the burst limit in bytes. The byte qualifier is optional and mutually exclusive with the kilobytes qualifier.</p> <p>kilobyte — The kilobyte qualifier is used to specify that the value given for size must be interpreted as the burst limit in Kilobytes. The kilobyte qualifier is optional and mutually exclusive with the bytes qualifier. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.</p>

cbs

Syntax	cbs <i>size-in-kbytes</i> no cbs
Context	config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue
Description	<p>This command can be used to override specific attributes of the specified queue's CBS parameters.</p> <p>It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given 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.</p>

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. If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

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 desired, 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 — 131072, default

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context 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 high-prio-only parameters. The **high-prio-only** command configures the percentage of buffer space for the queue, used exclusively by high priority packets.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The defined **high-prio-only** value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the **high-prio-only** value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command restores the default high priority reserved size.

Parameters *percent* — The *percent* parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

Values 0 — 100, default

mbs

Syntax **mbs** *size* [bytes|kilobytes]

no mbs

Context	config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue
Description	<p>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 enqueueing 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.</p> <p>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 packet's RED slope will not force the discard of the packet. Setting proper 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. If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.</p> <p>The no form of this command returns the MBS size assigned to the queue to the default value.</p>
Default	default
Parameters	<p><i>size-in-kbytes</i> — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.</p> <p>Values 0 — 1073741824 or default in expressed bytes or kilobytes</p>

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 config>port>eth>access>ing>qgrp>qover>q config>port>eth>access>egr>qgrp>qover>q config>port>ethernet>network>egr>qgrp>qover>q
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

This command defines an optional parent scheduler that further divides 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 weight** — These optional keywords are mutually exclusive to the keyword **level**. *weight* defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.

All **weight** values from all weighted active 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 queue or scheduler after the strict children are serviced. A weight is considered to be active when the pertaining queue or scheduler

has not reached its maximum rate and still has packets to transmit. All child queues and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all strict and non-zero weighted queues and schedulers are operating at the maximum bandwidth or are idle.

Values 0 — 100

Default 1

cir-weight *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 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 — 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-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit 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* — 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.

Values Percentage ranging from 0.01 to 100.00. The default is 100.00.

cir *cir-percent* — The **cir** keyword is optional and when defined the required *percent-of-line-rate* CIR parameter expresses 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.

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue
Description	<p>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.</p> <p>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.</p> <p>The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-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 as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.</p> <p>The CIR can be used by the queue's parent commands <i>cir-level</i> and <i>cir-weight</i> parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.</p> <p>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 <i>queue-id</i>.</p> <p>The no form of the command returns all queues created with the <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters (max, 0).</p>
Default	rate max cir 0 — The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.
Parameters	<p><i>pir-rate</i> — 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.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queue's adaptation-rule parameters and the actual hardware where the queue is provisioned.</p> <p>Values 1 — 3200000000 or max kbps</p> <p>Default max</p>

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 or queues.

Values 0 — 3200000000, **max**, **sum kbps**

Default 0

scheduler-override

Syntax [no] **scheduler-override**

Context 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**

Context 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 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 queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). 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 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 criteria below, 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, etc.), 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	<p>weight <i>weight</i> — 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 queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.</p> <p>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.</p> <p>Values 0 to 100</p> <p>Default 1</p> <p>cir-weight <i>cir-weight</i> — The cir-weight keyword defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same <i>cir-level</i> 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 queue or scheduler that the cir-weight pertains to has not reached the CIR and still has packets to transmit.</p> <p>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.</p> <p>Values 0 — 100</p> <p>Default 1</p>

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	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	<p>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 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.</p> <p>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 queues on the parent based on higher priority. The children of the scheduler may not need the</p>

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 queues and schedulers to operate at their maximum rates.

The **no** form of this command returns all queues created with this *queue-id* by association with the QoS policy to the default PIR and CIR parameters.

Parameters

pir-rate — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** or **sum** is accepted. Any other value will result in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue's **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue will be allowed to forward packets in a given second, thus shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queues PIR rate to the default value, that value must be specified as the PIR value.

Values 1 — 100000000, **max**

Default max

cir cir-rate — The **cir** parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue will operate. A value of 0 to 250 or the keyword **max** is accepted. Any other value will result in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir pir-rate** is multiplied by the *cir cir-rate*. If the **cir** is set to **max**, then the CIR rate is set to infinity.

The SAP ingress context for CIR is dependent on the defined forwarding class (fc) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

Values 0 — 10000000, **max**, **sum**

Default sum

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context 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 once 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. 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 ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

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 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 Per 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 — 4094

outer-copy — 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>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 6](#) defines the default behavior for Dot1P evaluation.

Table 6: 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

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.

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

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	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.

Table 7: 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

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

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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.

Note that 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**.

VLL SDP Commands

spoke-sdp

Syntax	spoke-sdp <i>sdp-id[:vc-id]</i> [vc-type {ether vlan}] [no-endpoint] spoke-sdp <i>sdp-id[:vc-id]</i> [vc-type {ether vlan}] endpoint <i>endpoint-name</i> [icb] no spoke-sdp <i>sdp-id[:vc-id]</i>
Context	config>service>cpipe config>service>epipe
Description	<p>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.</p> <p>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.</p> <p>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 <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> 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.</p> <p>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.</p>
Default	No <i>sdp-id</i> is bound to a service.
Special Cases	<p>Epipe — At most, only one <i>sdp-id</i> 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.</p> <p>L2TPv3 SDP types are only supported on EPipe services and not other xPipe services.</p>
Parameters	<p><i>sdp-id</i> — The SDP identifier. Allowed values are integers in the range of 1 to 17407 for existing SDPs.</p> <p><i>vc-id</i> — The virtual circuit identifier. The VC-ID is not used with L2TPv3 SDPs, however it must be configured.</p> <p>Values 1 — 4294967295</p> <p>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</p>

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 **ether** 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 **ether** 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 *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 <i>sdp-id[:vc-id]</i> [no-endpoint] spoke-sdp <i>sdp-id[:vc-id]</i> endpoint <i>endpoint-name</i> [icb] no spoke-sdp <i>sdp-id[:vc-id]</i>
Context	config>service>cpipe
Description	<p>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.</p> <p>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.</p> <p>The SDP must already be defined in the config>service>sdp context in order to associate an SDP with a service. If the sdp <i>sdp-id</i> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</p>

VLL SDP Commands

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 <i>sdp-id</i> is bound to a service.
Parameters	<i>sdp-id</i> — The SDP identifier. Allowed values are integers in the range of 1 to 17407 for existing SDPs. <i>vc-id</i> — The virtual circuit identifier. Values 1 — 4294967295 no endpoint — Adds or removes a spoke SDP association. endpoint <i>endpoint-name</i> — Specifies the name of the service endpoint. icb — Configures the spoke SDP as an inter-chassis backup SDP binding.

hash-label

Syntax	hash-label [signal-capability] no hash-label
Context	config>service>epipe>spoke-sdp config>service>pw-template config>service>vprn config>service>vprn>interface>spoke-sdp config>service>ies>interface>spoke-sdp
Description	<p>This command enables the use of the hash label on a VLL or VPLS service bound to LDP or RSVP SDP. This feature is not supported on a service bound to a GRE SDP. This feature is also not supported on multicast packets forwarded using RSVP P2MP LPS 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 spoke-interface.</p> <p>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).</p> <p>In order 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.</p> <p>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</p>

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 or 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 or 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.

control-word

Syntax **[no] control-word**
 config>service>cpipe>spoke-sdp
 config>service>epipe>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).

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.

hash-label

Syntax	hash-label [signal-capability] no hash label
Context	config>service>epipe>spoke-sdp
Description	<p>This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the with the ldp, rsvp-te, or mpls options. 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..</p> <p>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.</p> <p>In order 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.</p> <p>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.</p> <p>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.</p> <p>The TTL of the Hash Label is set to a value of 0.</p> <p>The no form of this command disables the use of the hash label.</p> <p>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</p>

solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The or 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 or 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 ePipe service will not transition to Oper State: Down when a SAP fails and when this optional command configured under that specific SAP. Only a single SAP in an ePipe may have this optional command included.

Default no ignore-oper-down

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
config>service>vprn>interface>spoke-sdp>egress
config>service>ies>interface>spoke-sdp>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/dot1.p/DSCP and the tunnel's 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 PW. This is true regardless if 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/dot1.p and the tunnel's DEI/dot1.p/EXP, but the DSCP is not modified by the policer's operation.

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/dot1.p/DSCP and the tunnel's 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 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—65535

port-redirect-group *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

instance *instance-id* — Specifies the optional identification of a specific instance of the queue-group.

Values 1—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
config>service>vprn>interface>spoke-sdp>ingress

```
config>service>ies>interface>spoke-sdp>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/VP RN 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 if 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 VP RN 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	<p><i>network-policy-id</i> — Specifies the network policy identification on the system.</p> <p>Values 1—65535</p> <p>fp-redirect-group <i>queue-group-name</i> — Specifies the name of the queue group template up to 32 characters in length.</p> <p>instance <i>instance-id</i> — Specifies the identification of a specific instance of the queue-group.</p> <p>Values 1—16384</p>
Description	This command configures the egress VC label.
Parameters	<p><i>vc-label</i> — A VC egress value that indicates a specific connection.</p> <p>Values 16 — 1048575</p>

vc-label

Syntax [no] **vc-label** *vc-label*

Context

monitor-oper-group

Syntax	monitor-oper-group <i>group-name</i> 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 <i>name</i> 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	<i>group-name</i> — Specifies an oper group name.

oper-group

Syntax	oper-group <i>group-name</i> 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	<i>group-name</i> — Specifies the Operational-Group identifier up to 32 characters in length.

precedence

Syntax	precedence [<i>precedence-value</i> primary] no precedence
Context	config>service>cpipe>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 — 4

primary — Specifies to make this the primary 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

Description This command configures the egress VC label.

Parameters *vc-label* — A VC egress value that indicates a specific connection.

Values 16 — 1048575

vc-label

Syntax	[no] vc-label <i>vc-label</i>
Context	config>service>cpipe>spoke-sdp>ingress config>service>epipe>spoke-sdp>ingress
Description	This command configures the ingress VC label.
Parameters	<i>vc-label</i> — A VC ingress value that indicates a specific connection.
	Values 2048 — 18431

vlan-vc-tag

Syntax	vlan-vc-tag <i>0..4094</i> no vlan-vc-tag [<i>0..4094</i>]
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	<i>0..4094</i> — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

spoke-sdp-fec

Syntax	spoke-sdp-fec spoke-sdp-fec <i>spoke-sdp-fec-id</i> [fec <i>fec-type</i>] [aii-type <i>aii-type</i>] [create] spoke-sdp-fec <i>spoke-sdp-fec-id</i> no-endpoint spoke-sdp-fec <i>spoke-sdp-fec-id</i> [fec <i>fec-type</i>] [aii-type <i>aii-type</i>] [create] endpoint <i>name</i> [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 7x50 of the MS-PW. The 7x50 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. In Release 9.0, only the Generalised 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	<p><i>spoke-sdp-fec-id</i> — An unsigned integer value identifying the spoke-SDP.</p> <p>Values 1 — 4294967295</p> <p><i>fec fec-type</i> — An unsigned integer value for the type of the FEC used by the MS-PW.</p> <p>Values 129 — 130</p> <p><i>aai-type aii-type</i> — An unsigned integer value for the Attachment Individual Identifier (AII) type used to identify the MS-PW endpoints.</p> <p>Values 1 — 2</p> <p>endpoint endpoint-name — Specifies the name of the service endpoint</p> <p>no endpoint — Adds or removes a spoke SDP association.</p> <p>icb — Configures the spoke-SDP as an inter-chassis backup SDP binding.</p>

auto-config

Syntax	[no] auto-config
Context	<code>config>service>epipe>spoke-sdp-fec</code>
Description	<p>This command enables single sided automatic endpoint configuration of the spoke-SDP. The 7x50 acts as the passive T-PE for signaling this MS-PW.</p> <p>Automatic Endpoint Configuration allows the configuration of a spoke-SDP endpoint without specifying the TAIID associated with that spoke-SDP. It allows a single-sided provisioning model where an incoming label mapping message with a TAIID that matches the SAIID 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 7x50 T-PE for which auto-config is specified will act as the passive T-PE.</p>

VLL SDP Commands

The **auto-config** command is blocked in CLI if signaling active has been enabled for this spoke-SDP. 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 7x50 T-PE either acts as the active T-PE (if signaling active is configured) or automatically determines which 7x50 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 7x50 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 7x50 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 *path-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 *precedence-value* — Specifies the spoke SDP precedence.

Values 1 — 4

primary — Specifies to make this the primary spoke SDP.

pw-template-bind

Syntax	pw-template-bind <i>policy-id</i> 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	<i>policy-id</i> — Specifies the existing policy ID Values 1 — 2147483647

retry-count

Syntax	retry-count <i>retry-count</i> no retry-count
Context	config>service>epipe>spoke-sdp-fec
Description	This optional command specifies the number of attempts software should make to re-establish 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	<i>retry-count</i> — The maximum number of retries before putting the spoke-sdp into the shutdown state. Values 10 — 10000

retry-timer

Syntax	retry-timer <i>retry-timer</i> 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 re-establish a spoke-SDP if it fails and a label withdraw message is received with the status code “All unreachable”.

VLL SDP Commands

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 — 480

saii-type2

Syntax **saii-type2** *global-id:prefix:ac-id*
no saii-type2

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 7x50 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 — 4294967295

prefix — The prefix on this 7x50 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 — 4294967295

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 — 4294967295

signaling

Syntax **signaling** *signaling*

Context config>service>epipe>spoke-sdp-fec

Description This command enables a user to configure this 7x50 as the active or 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 7x50 will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAI has the greater value prefix, then the 7x50 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 7x50 T-PE automatically selects the which 7x50 will initiate MS-PW signaling based on the prefix values configured in the SAII and TAI of the spoke-SDP, as described above.

Default auto

Parameters *signaling* — Configures this 7x50 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

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 7x50 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 — 4294967295

prefix — The prefix on this 7x50 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 — 4294967295

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 — 4294967295

ATM Commands

atm

Syntax	atm
Context	config>service>epipe>sap config>service>apipe>sap config>service>ipipe>sap config>service>epipe>sap
Description	<p>This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:</p> <ul style="list-style-type: none">• 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. <p>If ATM functionality is not supported for a given context, the command returns an error.</p>

egress

Syntax	egress
Context	config>service>epipe>sap config>service>epipe>sap>atm config>service>apipe>sap>atm config>service>fpipe>sap
	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 <i>atm-encap-type</i>
Context	config>service>epipe>sap>atm config>service>ipipe>sap>atm
Description	This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684, <i>Multiprotocol Encapsulation over ATM AAL5</i> , and to the ATM Forum LAN Emulation specification. Ingress traffic that does not match the configured encapsulation will be dropped.
Default	The encapsulation is driven by the services for which the SAP is configured. For IES and VPRN service SAPs, the default is aal5snap-routed .
Parameters	<i>atm-encap-type</i> — Specify 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 <i>traffic-desc-profile-id</i> 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 given 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	<i>traffic-desc-profile-id</i> — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

OAM Commands

oam

Syntax	oam
Context	config>service>epipe>sap config>service>apipe>sap>atm
Description	<p>This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.</p> <ul style="list-style-type: none">• The ATM-capable MDAs support end-to-end and segment OAM functionality (AIS, RDI, Loop-back) 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	<p>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.</p> <p>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 DOWN, or enters a fault state and comes back 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).</p> <p>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 (Note that when alarm-cells is disabled, a PVCC will change operational status to 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 described above.</p>
Default	Enabled for PVCCs delimiting IES SAPs

terminate

Syntax	[no] terminate
Context	config>service>apipe>sap>atm>oam
Description	<p>This command specifies whether this SAP will act as an OAM termination point. ATM SAPs can be configured to tunnel or terminate OAM cells.</p> <p>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 towards the directly-attached ATM equipment and see the results of the reply.</p> <p>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 towards the directly-attached ATM equipment and see the results of the reply.</p> <p>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.</p> <p>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/VP RN. This is because the ATM and AAL5 layers are terminated.</p> <p>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).</p> <p>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/VP RN. This is because the ATM and AAL5 layers are terminated.</p>
Default	no terminate

CPipe SDP Commands

spoke-sdp

Syntax	spoke-sdp <i>sdp-id[:vc-id]</i> [no-endpoint] [create] spoke-sdp <i>sdp-id:vc-id</i> [create] endpoint <i>endpoint-name</i> [icb] no spoke-sdp <i>sdp-id[:vc-id]</i>
Context	config>service>cpipe
Description	<p>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.</p> <p>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.</p> <p>The SDP must already be defined in the config>service>sdp context. If the sdp <i>sdp-id</i> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</p> <p>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.</p> <p>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.</p>
Default	No <i>sdp-id</i> is bound to a service.
Parameters	<p><i>sdp-id</i> — The SDP identifier. Allowed values are integers in the range of 1 to 17407 for existing SDPs.</p> <p><i>vc-id</i> — The virtual circuit identifier.</p> <p>Values 1 — 4294967295</p> <p>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.</p> <p>VC types are derived according to IETF <i>draft-martini-l2circuit-trans-mpls</i>.</p> <ul style="list-style-type: none"> • 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. <p>Values ethernet</p> <p>no endpoint — removes the association of a spoke SDP with an explicit endpoint name.</p>

endpoint *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 enables the context to configure egress spoke-SDP context.

ingress

Syntax	ingress
Context	config>service>cpipe>spoke-sdp
Description	This command enables the context to configure ingress spoke-SDP context.

vc-label

Syntax	vc-label <i>egress-vc-label</i> no vc-label [<i>egress-vc-label</i>]
Context	config>service>cpipe>spoke-sdp>egress
Description	This command configures the spoke-SDP egress VC label.
Parameters	<i>egress-vc-label</i> — A VC egress value that indicates a specific connection. Values 16 — 1048575

vc-label

Syntax	vc-label <i>ingress-vc-label</i> no vc-label [<i>ingress-vc-label</i>]
Context	config>service>cpipe>spoke-sdp>ingress
Description	This command configures the spoke-SDP ingress VC label.
Parameters	<i>ingress-vc-label</i> — A VC ingress value that indicates a specific connection. Values 2048 — 18431

precedence

Syntax	precedence [<i>precedence-value</i>] primary no precedence
Context	config>service>cpipe>spoke-sdp
Description	<p>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.</p> <p>The no form of the command returns the precedence value to the default.</p>
Default	4
Parameters	<i>precedence-value</i> — Specifies the spoke SDP precedence. Values 1 — 4 primary — Specifies to make this the primary spoke SDP.

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 <i>name</i> [create] no epipe-sap-template <i>name</i>
Context	config>service>template
Description	This command specifies which SAP parameter template should be applied to the I2-ap SAP. This can only be changed when the I2-ap is shutdown. The no form of the command removes the template, the SAP will use default parameters.
Default	None
Parameters	<i>name</i> — Specifies the SAP template name associated with this template.

egress

Syntax	egress
Context	config>service>template
Description	This command enables the context to configure egress filter policies.

ingress

Syntax	ingress
Context	config>service>template>epipe-sap-template
Description	This command enables 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 enables 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	ip ip-filter-id — Specifies IP filter policy. The filter ID must already exist within the created IP filters. Values 1 — 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 ipv6-filter-id — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters. Values 1 — 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. 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 — 65535

qos

Syntax	qos <i>policy-id</i> no qos
Context	config>service>template>epipe-sap-template>egress
Description	This command associates an existing QoS policy with the template.
Parameters	<i>policy-id</i> —
	Values 1 — 65535, or a name up to 64 characters in length

qos

Syntax	qos <i>policy-id</i> { shared-queuing multipoint-shared } qos <i>policy-id</i> 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	<i>policy-id</i> — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.
	Values 1 — 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, in addition to 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.

Epipe SAP Template Commands

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).

Show Commands

```

show
  — service
    — id service-id
      — all
      — authentication
      — base
      — bgp-vpws
      — endpoint [endpoint-name]
      — labels
      — sap sap-id [detail]
      — sdp [[sdp-id:vc-id] | far-end ip-address] [mrp] [detail]
      — stp [sap-id] [detail]
      — spoke-sdp-fec [[1..4294967295]
      — vccv-bfd session [detail]
    — sap-using [sap sap-id]
    — sap-using interface [ip-address | ip-int-name]
    — sap-using [ingress | egress] filter filter-id
    — sap-using [ingress | egress] qos-policy qos-policy-id
    — sap-using authentication-policy policy-name
    — sdp[[sdp-id:vc-id] | far-end ip-address] [mrp] [detail]
    — sdp-using [sdp-id:vc-id] | far-end ip-address]
    — sdp
      — 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
    — service-using [epipe] [ies] [vpls] [mirror] [ipipe] [sdp sdp-id] [customer customer-id]
    — 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]
  — pw-port
    — pw-port [pw-port-id] [detail]
    — pw-port sdp [sdp-id]
    — pw-port sdp none

```

Clear Commands

```
clear
  — service
    — id service-id
      — mesh-sdp sdp-id[:vc-id] ingress-vc-label
      — spoke-sdp sdp-id:vc-id [ingress-vc-label] [l2tpv3]
    — statistics
      — id service-id
        — counters
        — spoke-sdp sdp-id:vc-id {all | counters | stp}
      — sap sap-id {all | counters | stp}
      — sdp sdp-id keep-alive
```

Debug Commands

```
debug
  — service
    — id service-id
      — [no] sap sap-id
        — [no] event-type {arp | config-change | oper-status-change}
      — [no] sdp sdp-id:vc-id
```

Tools Commands

```
tools
  — dump
    — epipe-map-access-to-egress-port {service target-svc-id [end-service end-svc-id]} | lag lag-id
```

VLL Show Commands

sap-using

Syntax **sap-using** [**msap**] [**dyn-script**] [**description**]
sap-using [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**] [**description**]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** *qos-policy-id*

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.
ingress — Specifies matching an ingress policy.
ingress — Specifies matching an ingress policy.
egress — Specifies matching an egress policy.
qos-policy *qos-policy-id* — The ingress QoS Policy ID for which to display matching SAPs.
Values 1 — 65535
filter *filter-id* — The ingress or egress filter policy ID for which to display matching SAPs.
Values 1 — 65535
sap *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1271 for command syntax.

dyn-script — Displays dynamic service SAPs information.

Output **Show Service SAP** — The following table describes 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.
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.

Label	Description (Continued)
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.

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
=====

```

sdp

- Syntax** `sdp [[sdp-id[:vc-id] | far-end ip-address] [detail]]`
`sdp sdp-id:vc-id mrp`
- Context** show>service
- Description** This command displays SDP information.
 If no optional parameters are specified, a summary SDP output for all SDPs is displayed.
- Parameters** *sdp-id* — Specifies the SDP ID.
Values 1 — 17407
- vc-id* — The virtual circuit ID on the SDP ID to be reset.
Values 1 — 4294967295
- far-end ip-address** — Displays only SDPs matching with the specified far-end IP address.
Default SDPs with any far-end IP address.
- mrp** — Specifies to display Multiple Registration Protocol (MRP) information.
- detail** — Displays detailed SDP information.
Default SDP summary output.
- **Show Service SDP** — The following table describes show service SDP output fields:

Label	Description
SDP Id	The SDP identifier.
Adm MTU	Specifies 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.
Opr MTU	Specifies 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.
IP address	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the desired state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
Deliver Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Flags	Specifies all the conditions that affect the operating status of this SDP.
Signal Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.

Label	Description (Continued)
Last Status Change	Specifies the time of the most recent operating status change to this SDP.
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SDP.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
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.
Hello Timeout	Specifies the number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	Specifies the number of SDP unmatched message replies.
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.
TX Hello Msgs	Specifies the number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	Specifies the number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
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.

Label	Description (Continued)
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

Sample Output

```
*A:ALA-12# show service sdp
=====
Services: Service Destination Points
=====
SdpId    Adm MTU   Opr MTU   IP address    Adm  Opr        Deliver Signal
-----
10       4462     4462     10.20.1.3    Up   Dn NotReady MPLS   TLDP
40       4462     1534     10.20.1.20   Up   Up          MPLS   TLDP
60       4462     1514     10.20.1.21   Up   Up          GRE    TLDP
100      4462     4462     180.0.0.2    Down Down       GRE    TLDP
500      4462     4462     10.20.1.50   Up   Dn NotReady GRE    TLDP
-----
Number of SDPs : 5
=====
*A:ALA-12#

*A:ALA-12# show service sdp 2 detail
=====
Service Destination Point (Sdp Id : 2) Details
=====
Sdp Id 2 - (10.10.10.104)
-----
Description          : GRE-10.10.10.104
SDP Id               : 2
Admin Path MTU       : 0
Far End              : 10.10.10.104
Admin State          : Up
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP
Last Status Change   : 02/01/2007 09:11:39
Last Mgmt Change     : 02/01/2007 09:11:46
Oper Path MTU        : 0
Delivery             : GRE
Oper State           : Down
VLAN VC Etype        : 0x8100
Adv. MTU Over.       : No

KeepAlive Information :
Admin State          : Disabled
Hello Time           : 10
Hello Timeout        : 5
Max Drop Count       : 3
Tx Hello Msgs        : 0
Oper State           : Disabled
Hello Msg Len        : 0
Unmatched Replies    : 0
Hold Down Time       : 10
Rx Hello Msgs        : 0

Statistics           :
I. Fwd. Pkts.        : 0
I. Fwd. Octs.        : 0
E. Fwd. Pkts.        : 0
I. Dro. Pkts.        : 0
I. Dro. Octs.        : 0
E. Fwd. Octets       : 0

L2TPv3 Information
-----
Ingress Cookie       : AB:BA:BA:BB:A0:00:00:00
```

Show, Clear, Debug Commands

```
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
```

```
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
```

```
=====
*A:ALA-12#
*A:ALA-12# show service sdp 8
=====
Service Destination Point (Sdp Id : 8)
=====
SdpId      Adm MTU   Opr MTU   IP address      Adm  Opr      Deliver Signal
-----
8          4462     4462     10.10.10.104   Up   Dn NotReady MPLS   TLDP
=====
*A:ALA-12#
```

```
*A:ALA-12# show service sdp 8 detail
=====
Service Destination Point (Sdp Id : 8) Details
=====
Sdp Id 8 - (10.10.10.104)
-----
Description          : MPLS-10.10.10.104
SDP Id              : 8
Admin Path MTU      : 0
Oper Path MTU       : 0
Far End             : 10.10.10.104
Delivery            : MPLS
Admin State         : Up
Oper State          : Down
Flags              : SignalingSessDown TransportTunnDown
Signaling           : TLDP
VLAN VC Etype       : 0x8100
Last Status Change  : 02/01/2007 09:11:39
Adv. MTU Over.      : No
Last Mgmt Change    : 02/01/2007 09:11:46

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

Associated LSP LIST :
Lsp Name            : to-104
Admin State         : Up
Oper State          : Down
Time Since Last Tran* : 01d07h36m
=====
```

```
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#
```

When network domains are configured, the SDP egress interface state can be verified by using the following command:

```
*A:Dut-T# show service sdp egressifs
=====
SDP Egress Ifs State Table
=====
Sdp Id      Network Domain      State
```

```

-----
100                               net1                               consistent
-----
SDPs : 1
=====
*A:Dut-Tr#

```

sdp-using

Syntax **sdp-using** [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]

Context show>service

Description Display services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

far-end *ip-address* — Displays only services matching with the specified far-end IP address.

Default Services with any far-end IP address.

Output **Show Service SDP Using** — The following table describes 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.

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

```

Show, Clear, Debug Commands

```

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#

```

service-using

Syntax `service-using`

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 `[service]` — Displays information for the specified service type.

- 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 (ELAN) 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** `sdp-id` — Displays only services bound to the specified SDP ID.
 - Default** Services bound to any SDP ID.
 - Values** 1 — 17407
- customer** `customer-id` — Displays services only associated with the specified customer ID.
 - Default** Services associated with any customer.
 - Values** 1 — 2147483647

Output **Show service-using output** — The following table describes the command 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.

Label	Description
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

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
300        Epipe    Up    Up        10           09/05/2006 13:24:15
-----
Matching Services :
=====
*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#
```

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]`
- Context** `show>service`
- Description** Displays the SDPs used by spoke-sdp-fecs at this node.

Sample Output

```
*A:Dut-C# show service spoke-sdp-fec-using
=====
Service Spoke-SDP-Fec Information
=====
```

Show, Clear, Debug Commands

SvcId Path	SpokeSdpFec	Oper-SdpBind	SAIL-Type2 TAII-Type2
1 n/a	1	17407:4294967245	3:10.20.1.3:1 6:10.20.1.6:1
2 n/a	2	17407:4294967247	3:10.20.1.3:2 6:10.20.1.6:2
3 n/a	3	17407:4294967248	3:10.20.1.3:3 6:10.20.1.6:3
4 n/a	4	17407:4294967249	3:10.20.1.3:4 6:10.20.1.6:4
5 n/a	5	17407:4294967250	3:10.20.1.3:5 6:10.20.1.6:5
6 n/a	6	17407:4294967251	3:10.20.1.3:6 6:10.20.1.6:6
7 n/a	7	17407:4294967252	3:10.20.1.3:7 6:10.20.1.6:7
8 n/a	8	17407:4294967253	3:10.20.1.3:8 6:10.20.1.6:8
9 n/a	9	17407:4294967254	3:10.20.1.3:9 6:10.20.1.6:9
10 n/a	10	17407:4294967255	3:10.20.1.3:10 6:10.20.1.6:10

Entries found: 10

vccv-bfd

- Syntax** **vccv-bfd session [detail]**
[sdp sdp-id[:vc-id]] vccv-bfd session [detail]
- 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** *service-id* — The identification number of the specific service.
- Values** service-id: 1 — 214748364
- sdp-id* — Specifies the SDP ID.
- Values** 1 — 17407
- vc-id* — The virtual circuit ID on the SDP.
- Values** 1 — 4294967295

Sample Output

```
*A:Dut-C# show service id 1000 vccv-bfd session
Forwarding Information

Local Discr      : 4002                Local State      : Up (3)
Local Diag       : 0 (None)           Local Mode       : Async
Local Min Tx     : 1000                Local Mult       : 3
Last Sent        : 12/06/2013 19:38:13  Local Min Rx    : 1000
Type             : central
Remote Discr     : 4001                Remote State     : Up (3)
Remote Diag      : 0 (None)           Remote Mode      : Async
Remote Min Tx    : 1000                Remote Mult      : 3
Last Recv        : 12/06/2013 19:38:12  Remote Min Rx   : 1000
=====
```

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 — 214748364
 svc-name: A string up to 64 characters in length.
 - all** — Display detailed information about the service.
 - arp** — Display ARP entries for the service.
 - base** — Display basic service information.
 - endpoint** — Display service endpoint information.
 - interface** — Display service interfaces.
 - labels** — Display labels being used by this service.
 - sap** — Display SAPs associated to the service.
 - sdp** — Display SDPs associated with the service.
 - split-horizon-group** — Display split horizon group information.
 - stp** — Display STP information.

Sample Output

```
A:bkxsim1611>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)
```

VLL Show Commands

```

Split Horiz Grp      : (Not Specified)
VC Type              : Ipipe
Admin Path MTU       : 0
Far End              : 5.5.5.5
Tunnel Far End       : n/a
Hash Label           : Disabled

Admin State          : Up
Acct. Pol            : None
Ingress Label        : 131048
Ingr Mac Fltr-Id    : n/a
Ingr IP Fltr-Id     : n/a
Ingr IPv6 Fltr-Id   : n/a
Admin ControlWord    : Not Preferred
Admin BW(Kbps)       : 0
Last Status Change   : 09/15/2010 13:06:46
Last Mgmt Change     : 09/15/2010 13:06:02
Endpoint             : N/A
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

VC Tag               : 0
Oper Path MTU        : 1568
Delivery             : MPLS

Oper State           : Up
Collect Stats        : Disabled
Egress Label         : 131053
Egr Mac Fltr-Id     : n/a
Egr IP Fltr-Id      : n/a
Egr IPv6 Fltr-Id    : n/a
Oper ControlWord     : False
Oper BW(Kbps)        : 0
Signaling            : TLDP
Precedence           : 4

KeepAlive Information :
Admin State          : Disabled
Hello Time           : 10
Max Drop Count       : 3
Oper State           : Disabled
Hello Msg Len        : 0
Hold Down Time       : 10

Statistics           :
I. Fwd. Pkts.        : 15
I. Fwd. Octs.        : 1460
E. Fwd. Pkts.        : 17
I. Dro. Pkts.        : 0
I. Dro. Octs.        : 0
E. Fwd. Octets       : 1604

-----
RSVP/Static LSPs
-----
Associated LSP LIST :
Lsp Name             : to-bksim180-1
Admin State          : Up
Time Since Last Tr* : 16h07m44s
Oper State           : Up

Lsp Name             : to-bksim180-2
Admin State          : Up
Time Since Last Tr* : 16h07m45s
Oper State           : Up

-----
Class-based forwarding :
-----
Class forwarding     : Enabled
Default LSP          : to-bksim180-1
EnforceDSTELspFc    : Disabled
Multicast LSP        : None

=====
FC Mapping Table
=====
FC Name              LSP Name
-----
ef                   to-bksim180-2
=====

```

Show, Clear, Debug Commands

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.
-----
```

Show, Clear, Debug Commands

```
=====
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)

```

Show, Clear, Debug Commands

```
-----  
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** *policy name* — Specifies the subscriber authentication policy statistics to display.
sap sap-id — Specifies the SAP ID statistics to display. See Common CLI Command Descriptions on page 1271 for command syntax.

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics  
=====  
Authentication statistics  
=====
```

Interface / SAP	Authentication Successful	Authentication 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** **Show service ID Output** — The following table describes the output fields when the **all** option is specified:

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.
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.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, L2OperDown, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
SAP Count	The number of SAPs specified for this service.

Sample Output

```
*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
```

Show, Clear, Debug Commands

```
MTU : 1514
Vc Switching : False
SAP Count : 1
Per Svc Hashing : Disabled
Force QTag Fwd : Disabled
SDP Bind Count : 1
```

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)
Ingress FP QGrp   : (none)
Ing FP QGrp Inst  : (none)
Egress Qos Policy : (none)
Egress Port QGrp : (none)
Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State          : Disabled
Hello Time           : 10
Max Drop Count       : 3
Oper State           : Disabled
Hello Msg Len        : 0
Hold Down Time       : 10

Statistics           :
I. Fwd. Pkts.        : 0
I. Fwd. Octs.        : 0
E. Fwd. Pkts.        : 0
I. Dro. Pkts.        : 0
I. Dro. Octs.        : 0
E. Fwd. Octets       : 0

-----
Control Channel Status
-----
PW Status            : disabled
Peer Status Expire   : false
Request Timer        : <none>
Acknowledgement      : false
Refresh Timer        : <none>

-----
ETH-CFM SDP-Bind specifics
-----
Squelch Levels       : None

-----
RSVP/Static LSPs
-----
Associated LSP List :
No LSPs Associated

-----
Class-based forwarding :
-----
Class forwarding     : Disabled
Default LSP          : Uknwn
EnforcedSTELspFc    : Disabled
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
QinQ Dot1p         : Default
Description         : (Not Specified)
Encap               : qinq

```

Show, Clear, Debug Commands

```
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
```

VLL Show Commands

```

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

```

Show, Clear, Debug Commands

```
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. Octets    : 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

```

Show, Clear, Debug Commands

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#

```

base

- Syntax** **base**
- Context** show>service>id
- Description** Displays basic information about the service ID including service type, description, SAPs.
- Output** **Show Service-ID Base** — The following table describes 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, 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.
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.

Label	Description (Continued)
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.

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 & 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#
```

bgp-vpws

Syntax	bgp-vpws
Context	show>service>id
Description	This command displays BGP VPWS related information for the service.

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
=====
```

Show, Clear, Debug Commands

```
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.

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 **Show Service-ID Labels** — The following table describes 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.

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#
```

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*. See Common CLI Command Descriptions on page 1271 for command syntax.
- interface interface-name* — Displays information for the specified IP interface.
- ip-address ip-address* — Displays information associated with the specified IP address.
- detail** — Displays detailed information.
- detail** — Displays detailed information for the SAP.
- Output** **Show Service-ID SAP** — The following table describes show service 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 Ether type 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, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCelpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, 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.

Label	Description (Continued)
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.
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

```
*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|bpggrp-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
Admin State    : Up                   Oper State       : Down
Flags          : ServiceAdminDown
                        PortOperDown
```

Show, Clear, Debug Commands

```
Last Status Change : 02/06/2007 12:01:14
Last Mgmt Change   : 02/06/2007 12:01:17
Admin MTU          : 1522
Ingress qos-policy : 1
Shared Q plcy      : n/a
Ingress Filter-Id  : n/a
tod-suite          : None
Oper MTU           : 1522
Egress qos-policy  : 1
Multipoint shared  : Disabled
Egress Filter-Id   : n/a
```

```
Multi Svc Site     : None
Acct. Pol          : None
Collect Stats      : Disabled
```

```
=====  
A:ALA-48>config>service>epipe#  
A:ALA-48>config>service>epipe# show service id 8 sap 881/1/2:4094 detail
```

Service Access Points(SAP)

```
=====  
Service Id         : 8  
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
Ingress qos-policy : 1
Shared Q plcy      : n/a
Ingress Filter-Id  : n/a
tod-suite          : None
Oper MTU           : 1522
Egress qos-policy  : 1
Multipoint shared  : Disabled
Egress Filter-Id   : n/a
```

```
Multi Svc Site     : None
Acct. Pol          : None
Collect Stats      : Disabled
```

Sap Statistics

```
-----  
Packets          Octets  
Forwarding Engine Stats  
Dropped          : 0
Off. HiPrio       : 0
Off. LowPrio      : 0
Off. Uncolor      : 0
Queueing Stats(Ingress QoS Policy 1)  
Dro. HiPrio       : 0
Dro. LowPrio      : 0
For. InProf       : 0
For. OutProf      : 0
Queueing Stats(Egress QoS Policy 1)  
Dro. InProf       : 0
Dro. OutProf      : 0
For. InProf       : 0
For. OutProf      : 0
```

Sap per Queue stats

```
-----  
Packets          Octets  
Ingress Queue 1 (Unicast) (Priority)  
Off. HiPrio       : 0
Off. LoPrio       : 0
Dro. HiPrio       : 0
Dro. LoPrio       : 0
For. InProf       : 0
For. OutProf      : 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#

```

If a TOD Suite is configured on a SAP, the name of the suite is shown in the show command output. The values of the policies may be different from those configured on the SAP, because the configured policy assignments may have been overruled by policy assignments of the TOD Suite.

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
Learned MAC Addr   : 0                    Total MAC Addr  : 0
Admin MTU          : 1518                  Static MAC Addr : 0
Ingress qos-policy : 1190                  Oper MTU        : 1518
Shared Q plcy     : n/a                    Egress qos-policy : 1190
Ingr IP Fltr-Id   : n/a                    Multipoint shared : Disabled
Ingr Mac Fltr-Id  : n/a                    Egr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a                    Egr Mac Fltr-Id : n/a
tod-suite         : suite_sixteen          Egr IPv6 Fltr-Id : n/a
Egr Agg Rate Limit : max                    qinq-pbit-marking : both
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
...

```

Show, Clear, Debug Commands

A:kerckhot_4#

sdp

- Syntax** **sdp** [[*sdp-id*[:*vc-id*] | **far-end** *ip-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.
- Default** All SDPs.
- Values** 1 — 17407
- far-end** *ip-address* — Displays only SDPs matching the specified far-end IP address.
- Default** SDPs with any far-end IP address.
- detail** — Displays detailed SDP information.
- Output** **Show Service-ID SDP** — The following table describes 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.

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 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.
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.

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

```

Show, Clear, Debug Commands

```

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#
```

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#
```

Show, Clear, Debug Commands

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
VC Type         : Ether
Admin Path MTU  : 1600
Far End         : 3.3.3.3
Admin State     : Up
Acct. Pol      : None
Ingress Label   : 119066
Ing mac Fltr   : n/a
Ing ip Fltr     : n/a
Ing ipv6 Fltr  : n/a
Admin ControlWord : Preferred
Last Status Change : 02/04/2007 22:52:43
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
Learned MAC Addr : 0
MAC Learning    : Enabled
MAC Aging       : Enabled
L2PT Termination : Disabled
MAC Pinning     : Disabled
KeepAlive Information :
Admin State     : Disabled
Hello Time      : 10
Max Drop Count  : 3

Type            : Spoke
VC Tag         : n/a
Oper Path MTU  : 1600
Delivery       : GRE
Oper State     : Down
Collect Stats  : Disabled
Egress Label   : 0
Egr mac Fltr  : n/a
Egr ip Fltr   : n/a
Egr ipv6 Fltr : n/a
Oper ControlWord : True
Signaling      : TLDP

Total MAC Addr : 0
Static MAC Addr : 0
Discard Unkwn Srce: Disabled
BPDU Translation : Disabled
Oper State      : Disabled
Hello Msg Len   : 0
Hold Down Time : 10

Statistics      :
I. Fwd. Pkts.  : 0
E. Fwd. Pkts.  : 0
I. Dro. 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
VC Type         : Ether
Admin Path MTU  : 1600
Far End         : 1.1.1.1
Admin State     : Up
Acct. Pol      : None
Ingress Label   : 115066
Ing mac Fltr   : n/a
Ing ip Fltr     : n/a
Ing ipv6 Fltr  : n/a
Admin ControlWord : Not Preferred

Type            : Spoke
VC Tag         : n/a
Oper Path MTU  : 1600
Delivery       : GRE
Oper State     : Up
Collect Stats  : Disabled
Egress Label   : 119068
Egr mac Fltr  : n/a
Egr ip Fltr   : n/a
Egr ipv6 Fltr : n/a
Oper ControlWord : False

```

Show, Clear, Debug Commands

```
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
Learned MAC Addr  : 0
MAC Learning      : Enabled
MAC Aging         : Enabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
KeepAlive Information :
Admin State       : Disabled
Hello Time        : 10
Max Drop Count    : 3
Statistics        :
I. Fwd. Pkts.    : 0
E. Fwd. Pkts.    : 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
Spoke Descr      : (Not Specified)
VC Type          : Ether
Admin Path MTU   : 1500
Far End          : 101.101.101.101
Hash Label       : Enabled
Admin State      : Up
Acct. Pol        : None
Ingress Label    : 0
Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred
Admin BW(Kbps)   : 0
Last Status Change : 10/08/2009 06:55:54
Last Mgmt Change  : 10/08/2009 07:04:27
Endpoint         : N/A
Class Fwding State : Down
Flags            : SvcAdminDown SdpOperDown
                  NoIngrVCLLabel NoEgrVCLLabel
                  PathMTUTooSmall
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
Description      : (Not Specified)
SDP Id           : 2000:1
Type             : Spoke
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
Admin ControlWord : False
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
                  NoIngrVCLLabel NoEgrVCLLabel
                  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

This 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		
Admin State	: Up	Oper State	: Up
Acct. Pol	: None	Collect Stats	: Disabled
Ingress Label	: 0	Egress Label	: 0

Show, Clear, Debug Commands

```
Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred
Admin BW(Kbps) : 0
BFD Template : None
BFD-Enabled : no
Last Status Change : 06/19/2014 17:31:16
Last Mgmt Change : 06/19/2014 17:23:47
Endpoint : N/A
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

Egr Mac Fltr-Id : n/a
Egr IP Fltr-Id : n/a
Egr IPv6 Fltr-Id : n/a
Oper ControlWord : False
Oper BW(Kbps) : 0
BFD-Encap : ipv4
Signaling : None
Force Vlan-Vc : Disabled
Precedence : 4

Application Profile: None
Transit Policy : None
Standby Sig Slave : False
Block On Peer Fault: False
Use SDP B-MAC : False

Ingress Qos Policy : (none)
Ingress FP QGrp : (none)
Ing FP QGrp Inst : (none)

Egress Qos Policy : (none)
Egress Port QGrp : (none)
Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State : Disabled
Hello Time : 10
Max Drop Count : 3

Oper State : Disabled
Hello Msg Len : 0
Hold Down Time : 10

Statistics :
I. Fwd. Pkts. : 0
I. Fwd. Octs. : 0
E. Fwd. Pkts. : 0
I. Dro. Pkts. : 0
I. Dro. Octs. : 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
Peer Status Expire : false
Request Timer : <none>
Acknowledgement : false
Refresh Timer : <none>

-----
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..4294967295]`
- Context** `show>service>id`
- Description** This command displays spoke-SDP FEC information.
- Parameters** `detail` — Displays detailed 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
=====

```

Show, Clear, Debug Commands

stp

Syntax	stp [detail]
Context	show>service>id
Description	This command displays information for the spanning tree protocol instance for the service.
Parameters	detail — Displays detailed information.

spoke-sdp-fec

Syntax	spoke-sdp-fec [[1..4294967295]
Context	show>service>id
Description	This command displays the details of a spoke-sdp-fec spoke-sdp.

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>sdp

- Description** 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.
 - Default** All SDPs.
 - Values** 1 — 17407
 - pw-port-id* — Specifies the pseudo-wire port identifier.
 - Values** 1 — 10239
 - far-end ip-address* — Displays only SDPs matching with the specified far-end IP address.
 - Default** SDPs with any far-end IP address.
 - 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.

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)
=====
```

Show, Clear, Debug Commands

```

SDP Binding port      : lag-1
VC-Id                 : 3
Encap                 : dot1q
VC Type               : ether
Oper Flags            : (Not Specified)
Monitor Oper-Group   : (Not Specified)

Admin Status         : up
Oper Status          : up

Statistics           :
I. Fwd. Pkts.       : 0
I. Fwd. Octs.       : 0
E. Fwd. Pkts.       : 0
I. Dro. Pkts.       : 0
I. Dro. Octs.       : 0
E. Fwd. Octets      : 0
=====

```

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 — 10239

detail — Displays detailed port information that includes all the **pw-port** output fields.

sdp *sdp-id* — The SDP ID for which to display matching PW port information.

Values 1 — 17407

Output **Show PW-Port** — The following table describes **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.

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      1        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
=====

```

VLL Clear Commands

id

Syntax	id <i>service-id</i>
Context	clear>service clear>service>statistics
Description	This command clears commands for a specific service.
Parameters	<i>service-id</i> — The ID that uniquely identifies a service.
Values	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

mesh-sdp

Syntax	mesh-sdp <i>sdp-id[:vc-id]</i> ingress-vc-label
Context	clear>service>id
Description	This command clears and reset sthe mesh SDP binding.
Parameters	<i>sdp-id</i> — The spoke SDP ID for which to clear statistics.
Values	1 — 17407
	<i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset.
Values	1 — 4294967295
	ingress-vc-label — Specifies to clear the ingress VC label.

spoke-sdp

Syntax	spoke-sdp <i>sdp-id:vc-id</i> [ingress-vc-label] [lt2pv3]
Context	clear>service>id
Description	This command clears and resets the spoke SDP bindings for the service.
Parameters	<i>sdp-id</i> — The spoke SDP ID to be reset.
Values	1 — 17407
	<i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset.
Values	1 — 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.

sap

Syntax	sap <i>sap-id</i> { all counters stp }
Context	clear>service>statistics
Description	This command clears SAP statistics for a SAP.
Parameters	<p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1271 for command syntax.</p> <p>all — Clears all SAP queue statistics and STP statistics.</p> <p>counters — Clears all queue statistics associated with the SAP.</p> <p>stp — Clears all STP statistics associated with the SAP.</p>

sdp

Syntax	sdp <i>sdp-id</i> keep-alive
Context	clear>service>statistics
Description	This command clears keepalive statistics associated with the SDP ID.
Parameters	<p><i>sdp-id</i> — The SDP ID for which to clear keepalive statistics.</p> <p>Values 1 — 17407</p>

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 <i>sdp-id[:vc-id]</i> { all counters stp }
Context	clear>service>statistics>id
Description	This command clears statistics for the spoke SDP bound to the service.

Show, Clear, Debug Commands

- Parameters** *sdp-id* — The spoke SDP ID for which to clear statistics.
- Values** 1 — 17407
- vc-id* — The virtual circuit ID on the SDP ID to be reset.
- Values** 1 — 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.

VLL Debug Commands

id

Syntax	id <i>service-id</i>
Context	debug>service
Description	This command debugs commands for a specific service.
Parameters	<i>service-id</i> — The ID that uniquely identifies a service.

Values

sap

Syntax	[no] sap <i>sap-id</i>
Context	debug>service>id
Description	This command enables debugging for a particular SAP.
Parameters	<i>sap-id</i> — Specifies the SAP ID.

event-type

Syntax	[no] event-type { arp config-change 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	arp — Displays ARP events. config-change — Debugs configuration change events. svc-oper-status-change — Debugs service operational status changes.

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
```

Show, Clear, Debug Commands

```
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 <i>sdp-id:vc-id</i>
Context	debug>service>id
Description	This command enables debugging for a particular SDP.
Parameters	<i>sdp-id</i> — Specifies the SDP ID.

VLL Tools Commands

epipe-map-access-to-egress-port

Syntax	epipe-map-access-to-egress-port { service <i>target-svc-id</i> [end-service <i>end-svc-id</i>]} lag <i>lag-id</i>
Context	tools>dump
Description	<p>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 travelling from SAP to egress SDP or SAP.</p> <p>This command will support Epipe services with the following combinations:</p> <ul style="list-style-type: none"> - 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 <p>The command can be executed by specifying either a service ID, service-ID range or an ingress LAG ID.</p> <p>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.</p> <p>This command replaces the command tools dump epipe-map-to-network, which has been deprecated.</p>
Parameters	<p><i>target-svc-id</i> — 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.</p> <p><i>end-svc-id</i> — This parameter is used to identify the end of the service ID range for which the command will be executed against.</p> <p><i>lag-id</i> — This parameter caused the command to return the egress port for all service with SAPs associated with the specified LAG ID.</p>

Show, Clear, Debug Commands

Virtual Private LAN Service

In This Chapter

This chapter provides information about Virtual Private LAN Service (VPLS), process overview, and implementation notes.

Topics in this chapter include:

- [VPLS Service Overview on page 359](#)
- [VPLS Features on page 363](#)
 - [VPLS Packet Walkthrough on page 360](#)
 - [VPLS Enhancements on page 363](#)
 - [VPLS over MPLS on page 364](#)
 - [VPLS Service Pseudowire VLAN Tag Processing on page 365](#)
 - [VPLS MAC Learning and Packet Forwarding on page 369](#)
 - [Pseudowire Control Word on page 373](#)
 - [Table Management on page 374](#)
 - [VPLS and Spanning Tree Protocol on page 383](#)
 - [Multiple Spanning Tree on page 385](#)
 - [Egress Multicast Groups on page 392](#)
 - [VPLS Redundancy on page 402](#)
 - [Object Grouping and State Monitoring on page 420](#)
 - [MAC Flush Message Processing on page 422](#)
 - [ACL Next-Hop for VPLS on page 426](#)
 - [SDP Statistics for VPLS and VLL Services on page 427](#)
 - [BGP Auto-Discovery for LDP VPLS on page 428](#)
 - [Multicast-Aware VPLS on page 447](#)

In This Chapter

- [RSVP and LDP P2MP LSP for Forwarding VPLS/B-VPLS BUM and IP Multicast Packets on page 451](#)
- [VPLS Service Considerations on page 464](#)
 - [SAP Encapsulations on page 464](#)
 - [VLAN Processing on page 464](#)
 - [Ingress VLAN Swapping on page 465](#)
 - [Service Auto-Discovery using Multiple VLAN Registration Protocol \(MVRP\) on page 466](#)
 - [VPLS E-Tree Services on page 476](#)

VPLS Service Overview

Virtual Private LAN Service (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. 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.

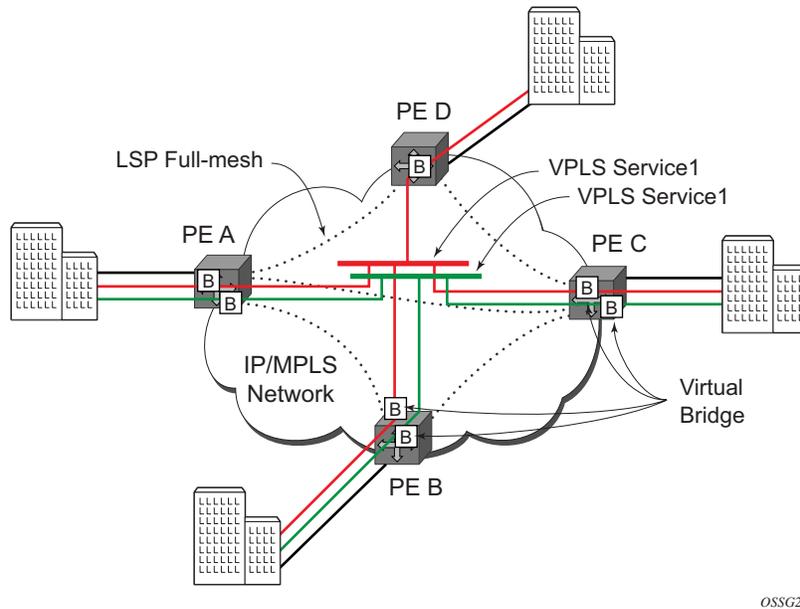
A VPLS service provides connectivity between two or more SAPs on one (which is considered a local service) or more (which is considered a distributed service) service routers. The connection appears to be a bridged domain to the customer sites so protocols, including routing protocols, can traverse the 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, thus, eliminating the need to train personnel on WAN technologies such as Frame Relay.

VPLS Packet Walkthrough

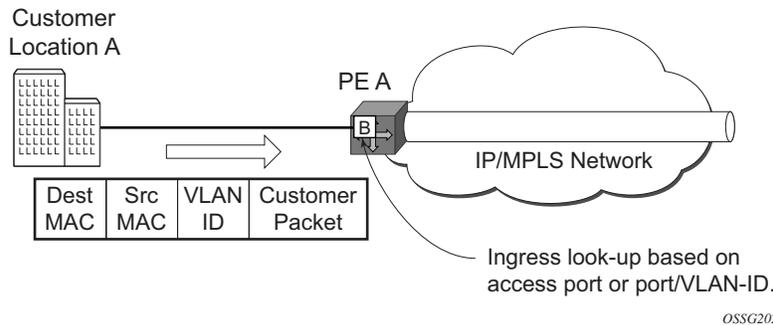
This section provides an example of VPLS processing of a customer packet sent across the network (Figure 43) from site-A, which is connected to PE-Router-A, to site-B, which is connected to PE-Router-C (Figure 44).



OSSG201

Figure 43: VPLS Service Architecture

1. PE-Router-A (Figure 44)
 - 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



OSSG202

Figure 44: 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 FIB 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 FIB 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 45):

- d. If the destination MAC address has already been learned by PE-Router-A, an existing entry in the FIB 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 once 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 45):

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 signaled for the VPLS instance. Note that the packet is not sent to PE-Router-D since this VPLS service does not exist on that PE-router.

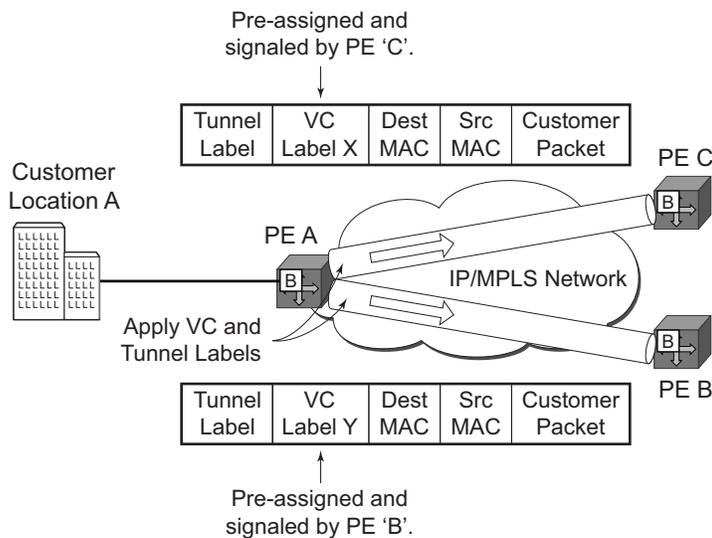


Figure 45: Network Port Egress Packet Format and Flooding

2. Core Router Switching

- a. 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 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 FIB table that associates the MAC address to PE-Router-A and the VC-label that PE-Router-A signaled it for the VPLS service on which the packet was received.
- c. The destination MAC address in the packet is looked up in the FIB 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).

Known MAC address (Figure 46)

- d. If the destination MAC address has been learned by PE-Router-C, an existing entry in the FIB 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.

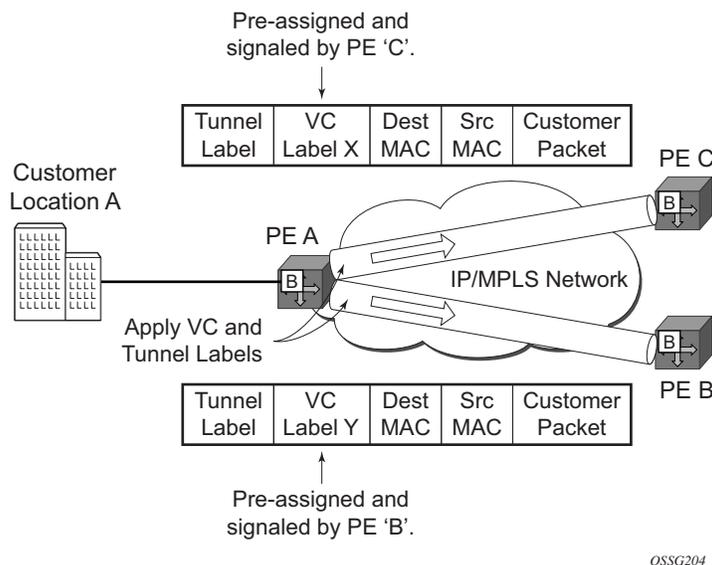


Figure 46: Access Port Egress Packet Format and Lookup

VPLS Features

This section features:

- [VPLS Enhancements on page 363](#)
 - [Pseudowire Control Word on page 373](#)
 - [Split Horizon SAP Groups and Split Horizon Spoke SDP Groups on page 382](#)
 - [VPLS and Spanning Tree Protocol on page 383](#)
 - [VPLS Redundancy on page 402](#)
 - [VPLS Access Redundancy on page 418](#)
 - [VCCV BFD Support for VPLS Services on page 441](#)
-

VPLS Enhancements

Alcatel-Lucent'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 Information Base (FIB) management features on a per service level including:
 - Configurable FIB size limit
 - FIB size alarms
 - MAC learning disable
 - Discard unknown
 - 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 Spanning Tree Protocol (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.
- IGMP snooping on a per-SAP and per-SDP basis.
- Optional SAP and/or spoke SDP redundancy to protect against node failure.

VPLS over MPLS

The VPLS architecture proposed in RFC 4762, *Virtual Private LAN Services Using LDP Signalling* 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 de-multiplexing 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 (label switched path) 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 forwarding information base (FIB) per VPLS service.

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 configurations 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 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 or spoke SDP, or in the related **pw-template**.
 - **vc-type ether** and **force-vlan-vc-forwarding** under the mesh 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 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 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 or spoke SDP, or in the related **pw-template**:

VPLS Service Pseudowire VLAN Tag Processing

- Routed, Etree or PBB VPLS services.
- L2PT termination on QinQ mesh 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 8](#) and [Table 9](#) describe the VLAN tag processing with respect to the zero, one and two VLAN tag configuration described above for the VLAN identifiers, Ether type, 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 8: 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
Ether type (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 9: 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	<ul style="list-style-type: none"> • the vlan-vc-tag value configured in pw-template or under the mesh/spoke SDP or • taken from the inner tag received on a QinQ SAP or QinQ mesh/spoke SDP or • taken 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 	Both inner and outer VLAN tag: <ul style="list-style-type: none"> • the vlan-vc-tag value configured in pw-template or under the mesh/spoke SDP or • taken from the inner tag received on a QinQ SAP or QinQ mesh/spoke SDP or • taken 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

Table 9: 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
Ether type (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)
Egress QoS (dot1p/DE) (set in VLAN tags)	N/A	<p>Taken from the innermost ingress service delimiting tag:</p> <ul style="list-style-type: none"> • the inner tag received on a QinQ SAP or QinQ mesh/spoke SDP or • taken 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. <p>Note that neither the inner nor outer dot1p/DE values can be explicitly set.</p>	<p>Both inner and outer dot1p/DE:</p> <p>Taken from the innermost ingress service delimiting tag:</p> <ul style="list-style-type: none"> • the inner tag received on a QinQ SAP or QinQ mesh/spoke SDP or • taken 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. <p>Note that neither the inner nor outer dot1p/DE values can be explicitly set.</p>

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

VPLS MAC Learning and Packet Forwarding

The 7950 XRS edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the to reduce the amount of unknown destination MAC address flooding.

7950 XRS routers learn the source MAC addresses of the traffic arriving on their access and network ports.

Each 7950 XRS maintains a Forwarding Information Base (FIB) for each VPLS service instance and learned MAC addresses are populated in the FIB table of the service. All traffic is switched based on MAC addresses and forwarded between all participating nodes using the LSP tunnels. Unknown destination packets (for example, the destination MAC address has not been learned) are forwarded on all LSPs 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.

MAC Learning Protection

In a Layer 2 environment, 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 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 customers is also disallowed but split-horizon will not be sufficient in the topology depicted in [Figure 47](#).

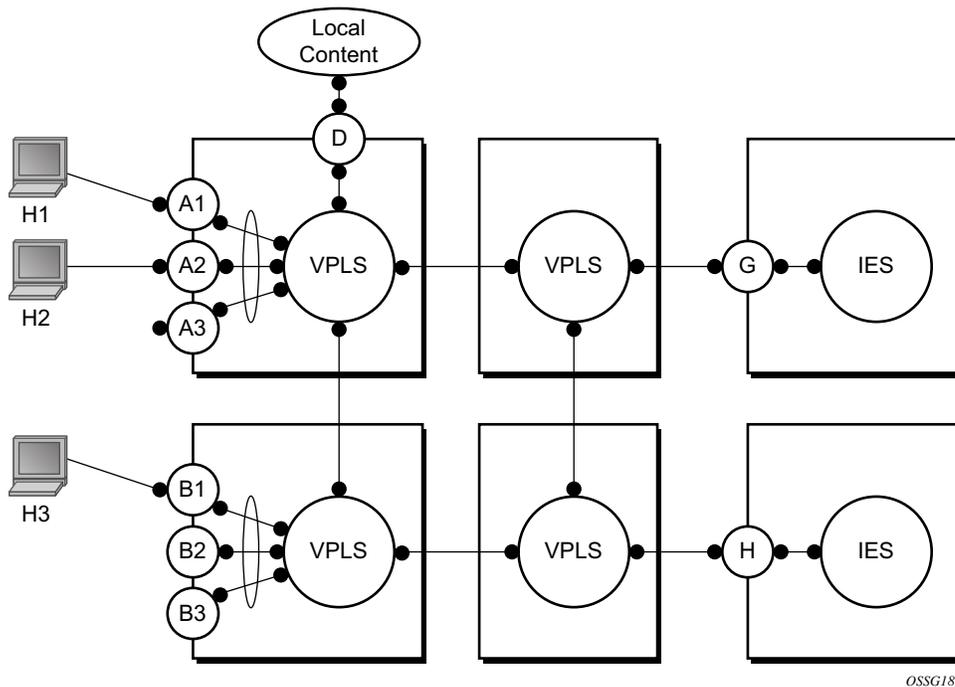


Figure 47: MAC Learning Protection

and 7950 XRS nodes 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 47) 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.

In order to eliminate the ability of a customer to cause a DOS attack, the node restricts the learning of protected MAC addresses based on a statically defined list. In addition 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.

DEI in IEEE 802.1ad

IEEE 802.1ad-2005 standard allows drop eligibility to be conveyed separately from priority in Service VLAN TAGs (STAGs) 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 three bit Priority Code Point (PCP) field and respectively in the DE Bit (Figure 48).

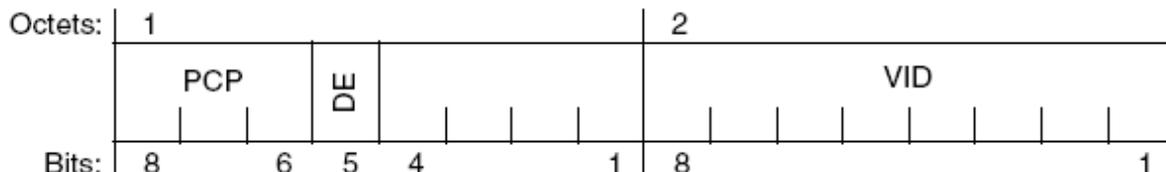


Figure 48: DE Bit in the 802.1ad S-TAG

The DE bit allows the S-TAG to convey eight forwarding classes/distinct emission priorities, each with a drop eligible indication.

When 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.

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 Alcatel-Lucent 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 ELAN capability. The LAG-emulating Ethernet tunnels and G.8031 protected Ethernet tunnels operate independently. (refer to LAG emulation using Ethernet Tunnels)

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 carries 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, 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 BVPLS using Ethernet tunnels the services are very similar.

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 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.

Table Management

The following sections describe VPLS features related to management of the Forwarding Information Base (FIB).

FIB Size

The following MAC table management features are required for each instance of a SAP or spoke SDP within a particular VPLS service instance:

- **MAC FIB size limits** — Allows users to specify the maximum number of MAC FIB entries that are learned locally for a SAP or remotely for a spoke SDP. If the configured limit is reached, then no new addresses will be learned from the SAP or spoke SDP until at least one FIB 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 FIB, 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.
-

FIB Size Alarms

The size of the VPLS FIB can be configured with a low watermark and a high watermark, expressed as a percentage of the total FIB size limit. If the actual FIB size grows above the configured high watermark percentage, an alarm is generated. If the FIB size falls below the configured low watermark percentage, the alarm is cleared by the system.

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 forwarding database (FIB). A local MAC address is a MAC address associated with a SAP because it ingresses 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 can happen in within tens of seconds beyond the age time. To minimize overhead, local MAC addresses on a LAG port and remote MAC addresses, in some circumstances, can take up to two times their respective age timer to be aged out.

Disable MAC Aging

The MAC aging timers can be disabled which will prevent any learned MAC entries from being aged out of the FIB. 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.

Disable MAC Learning

When MAC learning is disabled for a service, new source MAC addresses are not entered in the VPLS FIB, whether the MAC address is local or remote. MAC learning can be disabled for individual SAPs or spoke SDPs.

Unknown MAC Discard

Unknown MAC discard is a feature which discards all packets ingressing the service where the destination MAC address is not in the FIB. 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.

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.

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 re-learn rate for the MAC. When MAC move is enabled, the 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 simple level of control 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 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 given re-learn rate has been exceeded. The mechanism is based on the “cumulative” factor for every group of ports. Tertiary VPLS ports are blocked if the re-learn rate exceeds the configured threshold during one period while secondary ports are blocked only when re-learn 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).

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 then 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. **auto-learn-mac-protect** can be used to safe-guard against the possibility of MAC addresses being learned on the wrong SAP/SDP.

This feature provides the ability to automatically protect source MAC addresses which 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. It should be noted that 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 re-learned 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 contexts.
- 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 `restrict-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 given VPLS service.

Note, if `auto-learn-mac-protect` or `restrict-protected-src discard-frame` 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 given 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 given 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`

Operation

[Figure 49](#) shows a specific configuration using `auto-learn-mac-protect` and `restrict-protected-src discard-frame` in order to describe their operation.

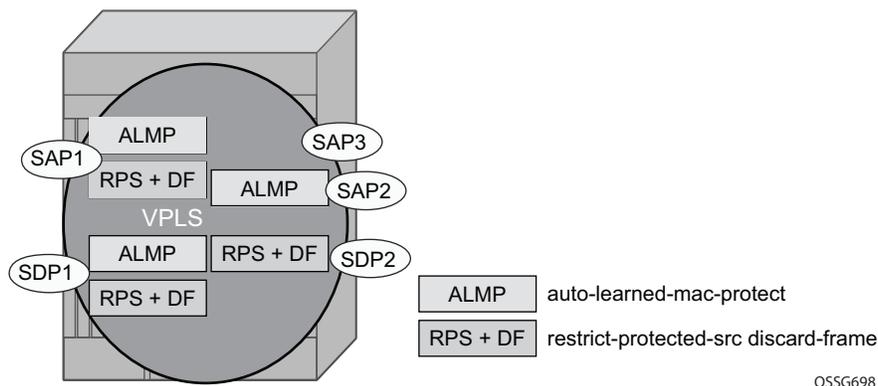


Figure 49: Auto-Learn-Mac-Protect Operation

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. auto-learn-mac-protect 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 describe 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 re-learned 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 on SAP2 and auto-learn-mac-protect is not enabled on SAP2, so the FDB would not be changed.
4. A frame with source MAC A enters SAP1, MAC A is re-learned 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, MAC A is re-learned on SDP1 as 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 re-learned 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 50 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.

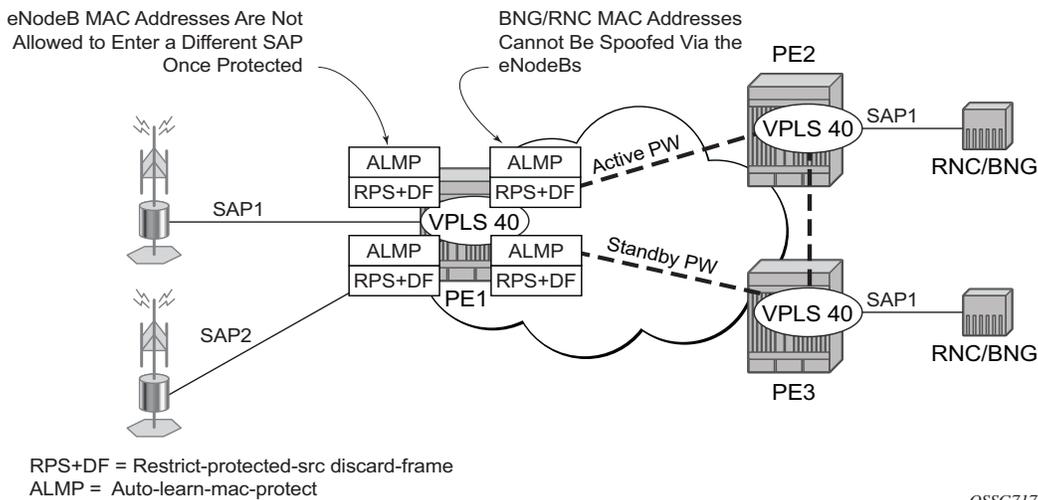


Figure 50: Auto-Learn-Mac-Protect Example

In order to protect the MAC addresses of the BNG/RNCs on PE1, **auto-learn-mac-protect** is enabled on the pseudo-wires connecting it to PE2 and PE3. Enabling **restrict-protected-src discard-frame** on the SAPs towards 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 **auto-learn-mac-protect** on the SAPs towards the eNodeBs will prevent the MAC addresses of the eNodeBs being learned on the wrong eNodeB SAP. Enabling **restrict-protected-src discard-frame** 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.

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).

VPLS and Spanning Tree Protocol

Alcatel-Lucent's VPLS service provides a bridged or switched Ethernet Layer 2 network. Equipment connected to SAPs forward Ethernet packets into the VPLS service. The 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. Alcatel-Lucent's implementation of the Spanning Tree Protocol (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.

Alcatel-Lucent's implementation of the Spanning Tree Protocol (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, refer to [Configuring a VPLS Service with CLI on page 483](#).

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 an mVPLS.

While the or 7950 XRS initially uses 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 and 7950 XRS support two BPDU encapsulation formats, and can dynamically switch between the following supported formats (on a per-SAP basis):

- IEEE 802.1D STP
- Cisco PVST

Multiple Spanning Tree

The Multiple Spanning Tree Protocol (MSTP) extends the concept of the IEEE 802.1w Rapid Spanning Tree Protocol (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 SR-Series implementation of Management VPLS (mVPLS) is used to group different VPLS instances under single RSTP instance. Introducing MSTP into the mVPLS allows interoperating with traditional Layer 2 switches in access network and provides an effective solution for dual homing of many business Layer 2 VPNs into a provider network.

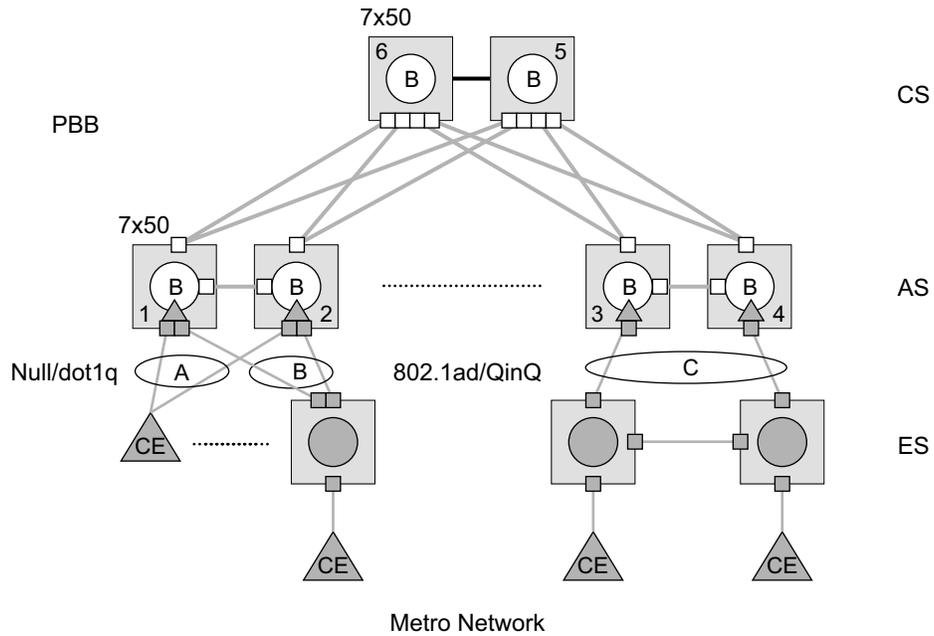
Redundancy Access to VPLS

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. In order 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 mVPLS on VPLS-PEs (only PEs directly connected to GigE MAN network) and then assign different managed-vlan ranges to different MSTP instances. Typically, the mVPLS 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 51](#) as example 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



OSSG205

Figure 51: 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 MVPLS.

MSTP for QinQ SAPs

MSTP runs in a MVPLS 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

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, 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.

MSTP General Principles

MSTP represents modification of RSTP which 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 single BPDU carries information for multiple MSTI which reduces overhead of the protocol.

Any given 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 given MSTI by configuring per VLAN assignment. This means that a 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 access network.

MSTP in the SR-Series Platform

The SR-Series platform uses a concept of mVPLS to group different SAPs under a single STP instance. The VLAN range covering SAPs to be managed by a given mVPLS is declared under a specific mVPLS SAP definition. MSTP mode-of-operation is only supported in an mVPLS.

When running MSTP, by default, all VLANs are mapped to the CIST. On 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.

Enhancements to the Spanning Tree Protocol

To interconnect routers (PE devices) across the backbone, service tunnels (SDPs) are used. These service tunnels are shared among multiple VPLS instances. Alcatel-Lucent's implementation of the Spanning Tree Protocol (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).

In order 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 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 towards the primary bridge. As a consequence, the path through the mesh are 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 which are running STP instances.

L2PT Termination

L2PT is used to transparently transport protocol data units (PDUs) of Layer 2 protocols such as STP, CDP, VTP and PAGP and UDLN. This allows running these protocols between customer CPEs without involving backbone infrastructure.

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 MAC address of the respective Layer 2 protocol.

and 7950 XRS routers support L2PT termination for STP BPDUs. More specifically:

- At ingress of every SAP/spoke SDP which is configured as L2PT termination, all PDUs with a MAC destination address, 01-00-0c-cd-cd-d0 will be intercepted and their MAC destination address will be overwritten to 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 redirection to CPM can be performed only at ingress. Therefore, to comply with the above requirement, as soon as at least 1 port of a given 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 given VPLS service.

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 and 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 and 7950 XRS devices. If enabled on a given SAP or spoke SDP, the system will intercept all BPDUs destined to 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 1 port within a given 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 given 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 given VPLS service.

L2PT and BPDU Translation

Cisco Discovery Protocol (CDP), Digital Trunking Protocol (DTP), Port Aggregation Protocol (PAGP), Uni-directional Link Detection (ULD) and Virtual Trunk Protocol (VTP) are supported. These protocols automatically pass the other protocols tunneled by L2PT towards 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.

Egress Multicast Groups

Efficient multicast replication is a method of increasing egress replication performance by combining multiple destinations into a single egress forwarding pass. In standard egress VPLS multicast forwarding, the complete egress forwarding plane is used per destination to provide ACL, mirroring, QoS and accounting for each path with associated receivers. In order to apply the complete set of available egress VPLS features, the egress forwarding plane must loop-back copies of the original packet so that each flooding destination may be processed. While each distributed egress forwarding plane only replicates to the destinations currently reached through its ports, this loop-back and replicate function can be resource intensive. When egress forwarding plane congestion conditions exist, unicast discards may be indiscriminate relative to forwarding priority. Another by-product of this approach is that the ability for the forwarding plane to fill the egress links is affected which could cause under-run conditions on each link while the forwarding plane is looping packets back to itself.

In an effort to provide highly scalable VPLS egress multicast performance for triple play type deployments, an alternative efficient multicast forwarding option is being offered. This method allows the egress forwarding plane to send a multicast packet to a set (called a chain) of destination SAPs with only a single pass through the egress forwarding plane. This minimizes the egress resources (processing and traffic management) used for the set of destinations and allows proper handling of congestion conditions and minimizes line under-run events. However, due to the batch nature of the egress processing, the chain of destinations must share many attributes. Also, egress port and ACL mirroring will be disallowed for packets handled in this manner.

Packets eligible for forwarding by SAP chaining are VPLS flooded packets (broadcast, multicast and unknown destination unicast) and IP multicast packets matching an VPLS Layer 2 (s,g) record (created through IGMP snooping).

Egress Multicast Group Provisioning

To identify SAPs in the chassis that are eligible for egress efficient multicast SAP chaining, an egress multicast group must be created. SAPs from multiple VPLS contexts may be placed in a single group to minimize the number of groups required on the system and to support multicast VPLS registration (MVR) functions.

Some of the parameters associated with the group member SAPs must be configured with identical values. The common parameters are checked as each SAP is provisioned into the group. If the SAP fails to be consistent in one or more parameters, the SAP is not allowed into the egress multicast group. Once a SAP is placed into the group, changing of a common parameter is not permitted.

Required Common SAP Parameters

Only SAPs created on Ethernet ports are allowed into an egress multicast group.

Required common parameters include:

- [SAP Port Encapsulation Type on page 393](#)
 - [SAP Port Dot1Q EtherType on page 393](#)
 - [Egress Multicast Groups on page 394](#)
 - [SAP Egress Filter on page 394](#)
-

SAP Port Encapsulation Type

The access port encapsulation type defines how the system will delineate SAPs from each other on the access port. SAPs placed in the egress multicast group must be of the same type. The supported access port encapsulation types are null and Dot1q. While all SAPs within the egress multicast group share the same encapsulation type, they are allowed to have different encapsulation values defined. The chained replication process will make the appropriate Dot1q value substitution per destination SAP.

The normal behavior of the system is to disallow changing the port encapsulation type once one or more SAPs have been created on the SAP. This being the case, no special effort is required to ensure that a SAP will be changed from null to Dot1q or Dot1q to null while the SAP is a member of a egress multicast group. Deleting the SAP will automatically remove the SAP from the group.

SAP Port Dot1Q EtherType

The access port dot1q-etype parameter defines which EtherType will be expected in ingress dot1q encapsulated frames and the EtherType that will be used to encapsulate egress dot1q frames on the port. SAPs placed in the same egress multicast group must use the same EtherType when dot1q is enabled as the SAPs encapsulation type.

The normal behavior of the system is to allow dynamic changing of the access port dot1q-etype value while SAPs are currently using the port. Once a dot1q SAP on an access port is allowed into an egress multicast group, the port on which the SAP is created will not accept a change of the configured dot1q-etype value. When the port encapsulation type is set to null, the port's dot1q-etype parameter may be changed at any time.

Egress Multicast Groups

Egress multicast groups to QinQ-encapsulated SAPs support includes:

- All SAP members of the given egress-multicast-group must have the same inner tag.
- A configuration flag, indicates, on a per egress-multicast-group basis, whether all member SAPs have the same inner or outer VLAN tag.

Membership rules for egress-multicast-groups in QinQ SAPs include:

- All SAPs that are members of the same egress-multicast-groups must have the same encapsulation type (as defined by `encap-type qinq` statement)
 - All SAP members of the given multicast group, port, or multicast-group must have the same inner Ethertype as well as outer Ethertype.
-

SAP Egress Filter

Due to the chaining nature of egress efficient multicast replication, only the IP or MAC filter defined for the first SAP on each chain is used to evaluate the packet. To ensure consistent behavior for all SAPs in the egress multicast group, when an IP or MAC filter is configured on one SAP it must be configured on all. To prevent inconsistencies, each SAP must have the same egress IP or MAC filter configured (or none at all) prior to allowing the SAP into the egress multicast group.

Attempting to change the egress filter configured on the SAP while the SAP is a member of an egress multicast group is not allowed.

If the configured common egress filter is changed on the egress multicast group, the egress filter on all member SAPs will be overwritten by the new defined filter. If the SAP is removed from the group, the previous filter definition is not restored.

SAP Egress QoS Policy

Each SAP placed in the egress multicast group may have a different QoS policy defined. When the egress forwarding plane performs the replication for each destination in a chain, the internal forwarding class associated with the packet is used to map the packet to an egress queue on the SAP.

In the case where customer SLA management is enabled on the SAP and the SAP queues are not available, the queues created by the `non-sub-addr-traffic` SLA-profile instance are used.

One caveat is that egress Dot1P markings for Dot1q SAPs in the replication chain are only evaluated for the first SAP in the chain. If the first SAP defines an egress Dot1P override for the

packet, all encapsulations in the chain will share the same value. If the first SAP in the chain does not override the egress Dot1P value, either the existing Dot1P value (relative to ingress) will be preserved or the value 0 (zero) will be used for all SAPs in the replication chain. The egress QoS policy Dot1P remark definitions on the other SAPs in the chain are ignored by the system.

Efficient Multicast Egress SAP Chaining

The egress XCM automatically creates the SAP chains on each egress forwarding plane. The size of each chain is based on the `dest-chain-limit` command defined on the egress multicast group to which the SAPs in the chain belong.

A set of chains is created by the XCM for each egress flooding list managed by the IOM. While SAPs from multiple VPLS contexts are allowed into a single egress multicast group, an egress flooding list is typically based on a subset of these SAPs. For instance, the broadcast/multicast/unknown flooding list for a VPLS context is limited to the SAPs in that VPLS context. With IGMP snooping on a single VPLS context, the flooding list is per Layer 2 IGMP (s,g) record and is basically limited to the destinations where IGMP joins for the multicast stream have been intercepted. When MVR (Multicast VPLS Registration) is enabled, the (s,g) flooding list may include SAPs from various VPLS contexts based on MVR configuration.

The system maintains a unique flooding list for each forwarding plane VPLS context (see section [VPLS Broadcast/Multicast/Unknown Flooding List on page 396](#)). This list will contain all SAPs (except for residential SAPs), spoke SDP and mesh SDP bindings on the forwarding plane that belong to that VPLS context. Each list may contain a maximum of 127 SAPs. In the case where the XCM is able to create an egress multicast chain, the SAPs within the chain are represented in the flooding list by a single SAP entry (the first SAP in the chain).

The system also maintains a unique flooding list for each Layer 2 IP multicast (s,g) record created through IGMP snooping (see sections [VPLS IGMP Snooping \(s,g\) Flooding List on page 397](#) and [MVR IGMP Snooping \(s,g\) Flooding List on page 397](#)). A flooding list created by IGMP snooping is limited to 127 SAPs, although it may contain other entries representing spoke and mesh SDP bindings. Unlike a VPLS flooding list, a residential SAP may be included in a Layer 2 IP multicast flooding list.

While the system may allow 30 SAPs in a chain, the uninterrupted replication to 30 destinations may have a negative effect on other packets waiting to be processed by the egress forwarding plane. Most notably, massive jitter may be seen on real time VoIP or other time-sensitive applications. The `dest-chain-limit` parameter should be tuned to allow the proper balance between multicast replication efficiency and the effect on time sensitive application performance. It is expected that the optimum performance for the egress forwarding plane will be found at around 16 SAPs per chain.

VPLS Broadcast/Multicast/Unknown Flooding List

The XCM includes all VPLS destinations in the egress VPLS Broadcast/Multicast/Unknown (BMU) flooding list that exist on a single VPLS context. Whenever a broadcast, multicast or unknown destination MAC is received in the VPLS, the BMU flooding list is used to flood the packet to all destinations. For normal flooding, care is taken at egress to ensure that the packet is not sent back to the source of the packet. Also, if the packet is associated with a split horizon group (mesh or spoke/SAP) the egress forwarding plane will prevent the packet from reaching destinations in the same split horizon context as the source SAP or SDP-binding.

The VPLS BMU flooding list may contain both egress multicast group SAPs and other SAPs or SDP bindings as destinations. The egress XCM will separate the egress multicast group SAPs from the other destinations to create one or more chains. Egress multicast group SAPs are placed into a chain completely at the discretion of the XCM and the order of SAPs in the list will be nondeterministic. When more SAPs exist on the VPLS context within the egress multicast group then are allowed in a single chain, multiple SAP chains will be created. The XCM VPLS egress BMU flooding list will then contain the first SAP in each chain plus all other VPLS destinations.

The SAPs in the same VPLS context must be in the same split horizon group to allow membership into the egress multicast group. The split horizon context is not required to be the same between VPLS contexts.

SAPs within the same VPLS context may be defined in different egress multicast groups, but SAPs in different multicast groups cannot share the same chain.

VPLS IGMP Snooping (s,g) Flooding List

When IGMP snooping is enabled on a VPLS context, a Layer 2 IP multicast record (s,g) is created for each multicast stream entering the VPLS context. Each stream should only be sent to each SAP or SDP binding where either a multicast router exists or a host exists that has requested to receive the stream (known as a receiver). To facilitate egress handling of each stream, the XCM creates a flooding list for each (s,g) record associated with the VPLS context. As with the BMU flooding list, source and split horizon squelching is enforced by the egress forwarding plane.

As with the BMU VPLS flooding list, the egress multicast group SAPs that have either static or dynamic multicast receivers for the (s,g) stream are chained into groups. The chaining is independent of other (s,g) flooding lists and the BMU flooding list on the VPLS instance. As the (s,g) flooding list membership is dynamic, the egress multicast group SAPs in chains in the list are also managed dynamically.

Since all SAPs placed into the egress multicast group for a particular VPLS context are in the same split horizon group, no special function is required for split horizon squelching.

MVR IGMP Snooping (s,g) Flooding List

When IGMP snooping on a SAP is tied to another VPLS context to facilitate cross VPLS context IP multicast forwarding, a Layer 2 IP multicast (s,g) record is maintained on the VPLS context receiving the multicast stream. This is essentially an extension to the VPLS IGMP snooped flooding described in [VPLS IGMP Snooping \(s,g\) Flooding List on page 397](#). The (s,g) list is considered to be owned by the VPLS context that the multicast stream will enter. Any SAP added to the list that is outside the target VPLS context (using the **from-vpls** command) is handled as an alien SAP. Split horizon squelching is ignored for alien SAPs.

When chaining the egress multicast group SAPs in an MVR (s,g) list, the XCM will keep the native chained SAPs in separate chains from the alien SAPs to prevent issues with split horizon squelching.

Mirroring and Efficient Multicast Replication

As previously stated, efficient multicast replication affects the ability to perform mirroring decisions in the egress forwarding plane. In the egress forwarding plane, mirroring decisions are performed prior to the egress chain replication function. Since mirroring decisions are only evaluated for the first SAP in each chain, applying a mirroring condition to packets that egress other SAPs in the chain has no effect. Also, the XCM manages the chain membership automatically and the user has no ability to provision which SAP is first in a chain. Thus, mirroring is not allowed for SAPs within a chain.

Port Mirroring

A SAP created on an access port that is currently defined as an egress mirror source may not be defined into an egress multicast group.

A port that has a SAP defined in an egress multicast group may not be defined as an egress mirror source. If egress port mirroring is desired, then all SAPs on the port must first be removed from all egress multicast groups.

Filter Mirroring

An IP or MAC filter that is currently defined on an egress multicast group as a common required parameter may not have an entry from the list defined as a mirror source.

An IP or MAC filter that has an entry defined as a mirror source may not be defined as a common required parameter for an egress multicast group.

If IP or MAC based filter mirroring is required for packets that egress an egress multicast group SAP, the SAP must first be removed from the egress multicast group and then an IP or MAC filter that is not associated with an egress multicast group must be assigned to the SAP.

SAP Mirroring

While SAP mirroring is not allowed within an IOM chain of SAPs, it is possible to define an egress multicast group member SAP as an egress mirror source. When the IOM encounters a chained SAP as an egress mirror source, it automatically removes the SAP from its chain, allowing packets that egress the SAP to hit the mirror decision. Once the SAP is removed as an egress mirror source, the SAP will be automatically placed back into a chain by the XCM.

It should be noted that all mirroring decisions affect forwarding plane performance due to the overhead of replicating the frame to the mirror destination. This is especially true for efficient multicast replication as removing the SAP from the chain also eliminates a portion of the replication efficiency along with adding the mirror replication overhead.

OAM Commands with EMG

There are certain limitations with using the OAM commands when egress multicast group (EMG) is enabled. This is because OAM commands work by looping the OAM packet back to ingress instead of sending them out of the SAP. Hence, if EMG is enabled, these OAM packets will be looped back once per chain and hence, will only be processed for the first SAP on each chain. Particularly, the **mac-ping**, **mac-trace** and **mfib-ping** commands will only list the first SAP in each chain.

XCM Chain Management

As previously stated, the XCM automatically creates the chain lists from the available egress multicast group SAPs. The XCM will create chains from the available SAPs based on the following rules:

1. SAPs from different egress multicast groups must be in different chains (a chain can only contain SAPs from the same group)
2. Alien and native SAPs must be in different chains
3. A specific chain cannot be longer than the defined dest-chain-limit parameter for the egress multicast group to which the SAPs belong

Given the following conditions for an XCM creating a multicast forwarding list (List 1) for a Layer 2 IP multicast (s,g) native to VPLS instance 100:

- Egress multicast group A
 - Destination chain length = 16
 - 30 member SAPs on VPLS 100 joined (s,g) (native to VPLS 100)
 - 41 member SAPs on other VPLS instances joined (s,g) (alien to VPLS 100)
- Egress multicast group B
 - Destination chain length = 8
 - 17 member SAPs on VPLS 100 joined (s,g) (native to VPLS 100)
- Egress multicast group C
 - Destination chain length = 12
 - 23 member SAPs on other VPLS instances joined (s,g) (alien to VPLS 100)

The system will build the SAP chains for List 1 according to [Table 10](#).

Table 10: SAP Chain Creation

Egress Forwarding List 1 SAP Chains					
Egress Multicast Group A Destination Chain Length 16		Egress Multicast Group B Destination Chain Length 8		Egress Multicast Group C Destination Chain Length 12	
Native Chains	Alien Chains	Native Chains	Alien Chains	Native Chains	Alien Chains
16	16	8			12
14	16	8			11
	9	1			

Adding a SAP to a Chain

A SAP must meet all the following conditions to be chained in a VPLS BMU flooding list:

1. The SAP is successfully defined as an egress multicast group member
2. The SAP is not currently an egress mirror source

Further, a SAP must meet the following conditions to be chained in an egress IP multicast (s,g) flooding list:

1. The SAP is participating in IGMP snooping
2. A static or dynamic join to the (s,g) record exists for the SAP or the SAP is defined as a multicast router port

Note: While an operationally down SAP is placed into replication chains, the system ignores that SAP while in the process of replication.

Based on the egress multicast group and the native or alien nature of the SAP in the list, a set of chains are selected for the SAP. The XCM will search the chains for the first empty position in an existing chain and place the SAP in that position. If an empty position is not found, the XCM will create a new chain with that SAP in the first position and add the SAP to the flooding list to represent the new chain.

Removing a SAP from a Chain

A SAP will be removed from a chain in a VPLS BMU flooding list or egress IP multicast (s,g) flooding list for any of the following conditions:

1. The SAP is deleted from the VPLS instance
2. The SAP is removed from the egress multicast group of which it was a member
3. The SAP is defined as an egress mirror source

Further, a SAP will be removed from an egress IP multicast (s,g) flooding list for the following conditions:

1. IGMP snooping removes the SAP as an (s,g) destination or the SAP is removed as a multicast router port

When the SAP is only being removed from the efficient multicast replication function, it may still need to be represented as a stand alone SAP in the flooding list. If the removed SAP is the first SAP in the list, the second SAP in the list is added to the flooding list when the first SAP is de-chained. If the removed SAP is not the first SAP, it is first de-chained and then added to the

flooding list. If the removed SAP is the only SAP in the chain, the chain is removed along with removing the SAP from the flooding list.

Moving a SAP from a chain to a stand alone condition or from a stand alone condition to a chain may cause a momentary glitch in the forwarding plane for the time that the SAP is being moved. Care is taken to prevent or minimize the possibility of duplicate packets being replicated to a destination while the chains and flooding lists are being manipulated.

Chain Optimization

Chains are only dynamically managed during SAP addition and removal events. The system does not attempt to automatically optimize existing chains. It is possible that excessive SAP removal may cause multiple chains to exist with lengths less than the maximum chain length. For example, if four chains exist with eight SAPs each, it is possible that seven of the SAPs from each chain are removed. The result would be four chains of one SAP each effectively removing any benefit of egress SAP replication chaining.

While it may appear that optimization would be beneficial each time a SAP is removed, this is not the case. Rearranging the chains each time a SAP is removed may cause either packet duplication or omitting replication to a destination SAP. Also, it could be argued that if the loop back replication load is acceptable before the SAP is removed, continuing with the same loop back replication load once the SAP is removed is also acceptable. It is important to note that the overall replication load is lessened with each SAP removal from a chain.

While dynamic optimization is not supported, a manual optimization command is supported in each egress multicast group context. When executed the system will remove and add each SAP, rebuilding the replication chains.

When the `dest-chain-limit` is modified for an egress multicast group, the system will reorganize the replication chains that contain SAPs from that group according to the new maximum chain size.

VPLS Redundancy

The VPLS standard (RFC 4762, *Virtual Private LAN Services Using LDP Signalling*) 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.

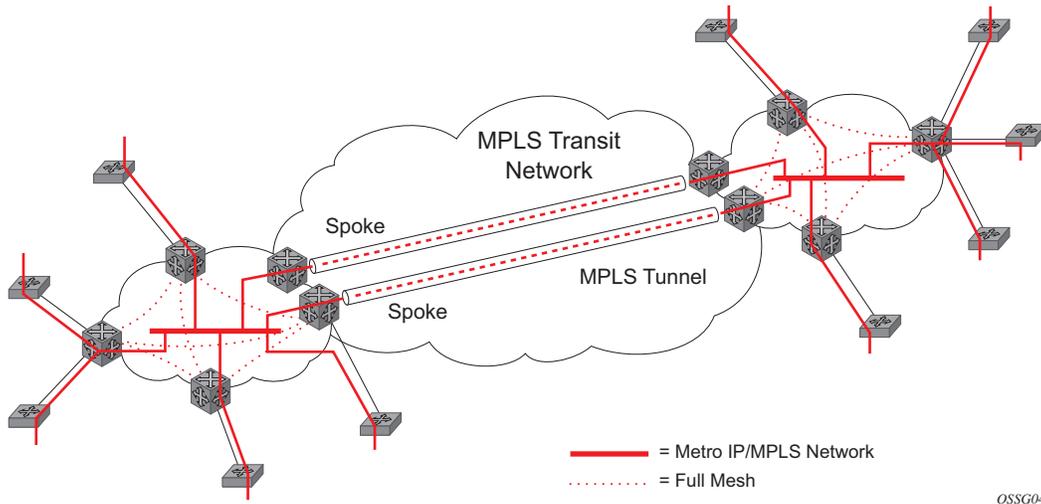
Alcatel-Lucent routers have implemented special features for improving the resilience of hierarchical VPLS instances, in both MTU and inter-metro applications.

Spoke SDP Redundancy for Metro Interconnection

When two or more meshed VPLS instances are interconnected by redundant spoke SDPs (as shown in [Figure 52](#)), a loop in the topology results. In order to remove such a loop from the topology, Spanning Tree Protocol (STP) can be run over the SDPs (links) which 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 which can associate a number of VPLSes to a single STP instance running over the redundant SDPs. Node redundancy is thus 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 mVPLS.

In the case of a failure of the active node, 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 re-learned 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 VPLSes with the two management VPLS services, load balancing across the spokes can be achieved.



OSSG045

Figure 52: HVPLS with Spoke Redundancy

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 mVPLS.

In the configuration example displayed in [Figure 52](#), the MTUs have spoke SDPs to two PEs 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, etc.), 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 generates 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.

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 53](#).

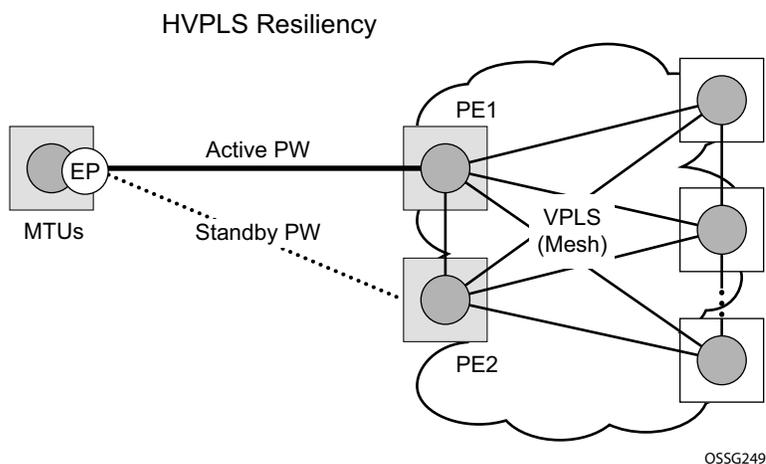


Figure 53: HVPLS Resiliency Based on AS Pseudowires

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 depicted in [Figure 54](#) can be used.

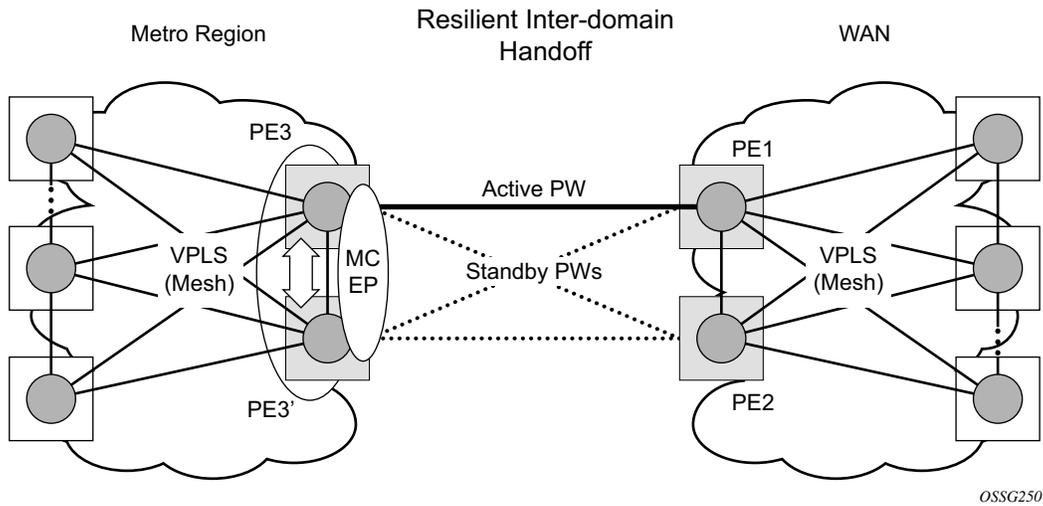


Figure 54: Multi-Chassis Pseudowire Endpoint for VPLS

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 point in 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 inter-connecting 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.

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 as 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 BFD session.

In order 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) then the following actions are not allowed under the interface: IP address change, “shutdown”, “no bfd” commands. If one of these action 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 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.

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.

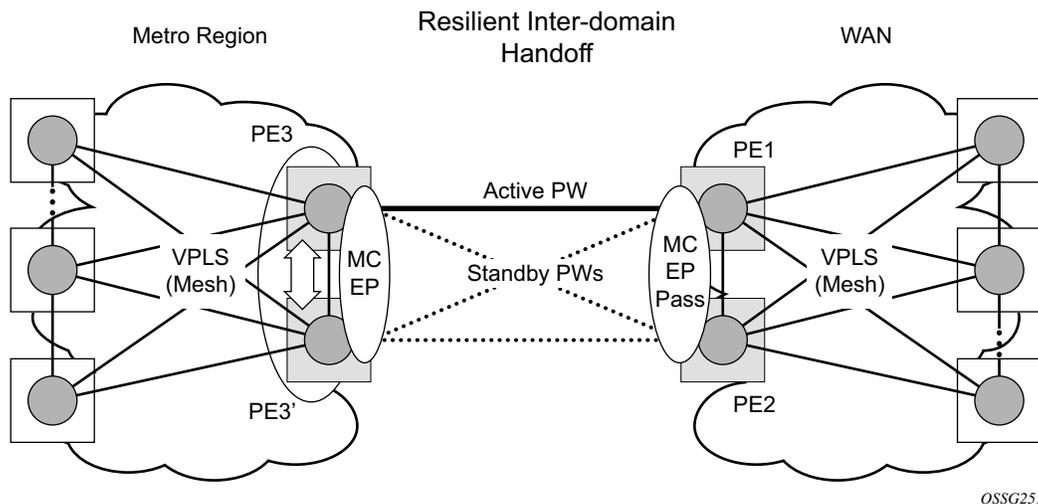


Figure 55: MC-EP in Passive Mode

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 mis-configuration 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
```

Support for Single Chassis Endpoint Mechanisms

In cases of SC-EP, there is 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 verify carefully 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 that 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 describes also how the main mechanisms used for single chassis endpoint are adapted for the MC-EP solution.

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 pseudowire in the MC-EP group. This section describes the MAC flush procedures that can be applied to ensure black-hole avoidance.

[Figure 56](#) depicts a pair of PE gateways (PE3 and PE3) running MC-EP towards 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 a MC-EP or a regular pseudowire endpoint.

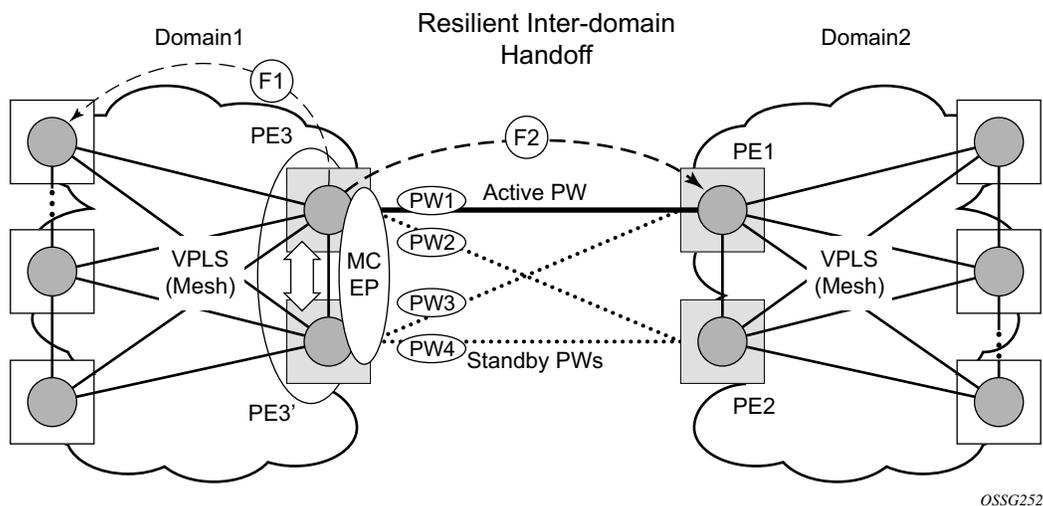


Figure 56: MAC Flush in the MC-EP Solution

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 towards VPLS mesh in the F1 direction.
 - PE3 sends a T-LDP **flush-all-but-mine** message towards PE2 in the F2 direction which is propagated by PE2 in the local VPLS mesh. Note that when MC-EP is in passive mode and the first spoke becomes active, a **no mac flush-all-but-mine** message will be generated.

Block-on-Mesh-Failure Support in MC-EP Scenario

The following rules describe how the block-mesh-on-failure must be ported to the MC-EP solution (see [Figure 56](#)):

- If PE3 does not have any forwarding path towards Domain1 mesh, it should block both PW1 and PW2 and inform PE3 so one of its pseudowires can be activated.
- In order to allow the use of block-on-mesh-failure for MC-EP, a new 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 from previous releases.
 - For a spoke SDP to be added under an endpoint, the setting for its **block-on-mesh-failure** parameter must be in sync 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 in any way the above behavior.

Prior to 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.

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-endpoint, the regular MC-EP selection algorithm (for example, the operational status -> precedence value) will be applied to determine the winner.

Revertive Behavior for Primary Pseudowire(s) in a MC-EP

For a single-chassis endpoint a revert-time command is provided under the VPLS endpoint. Refer to the [VPLS Services Command Reference on page 545](#) 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: i.e. if the active secondary pseudowire fails and is restored it will stay in standby until a configuration change or a force command occurs.

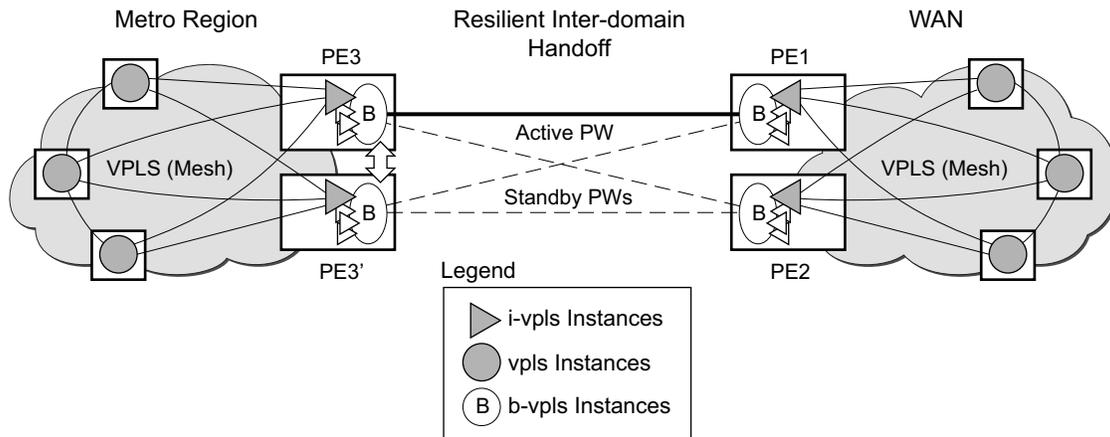
In the MC-EP case the revertive behavior is supported for pseudowire defined as primary (precedence 0). The following rules apply:

- 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
-

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 BVPLS instances. The MC-EP procedures described for regular VPLS apply.

PBB-VPLS can be also enabled just on the gateway MC-EP PEs as depicted in [Figure 57](#) below.



OSSG487

Figure 57: MC-EP with B-VPLS

Multiple I-VPLS instances may be used to represent in the gateway PEs the customer VPLS instances using 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 100s or even 1000s of customers 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 loadshare 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 PEs 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.

Note that for the PBB topology described in [Figure 57](#), 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.

MAC Flush Additions for PBB VPLS

The scenario depicted in [Figure 58](#) is used to define the blackholing problem in PBB-VPLS using MC-EP.

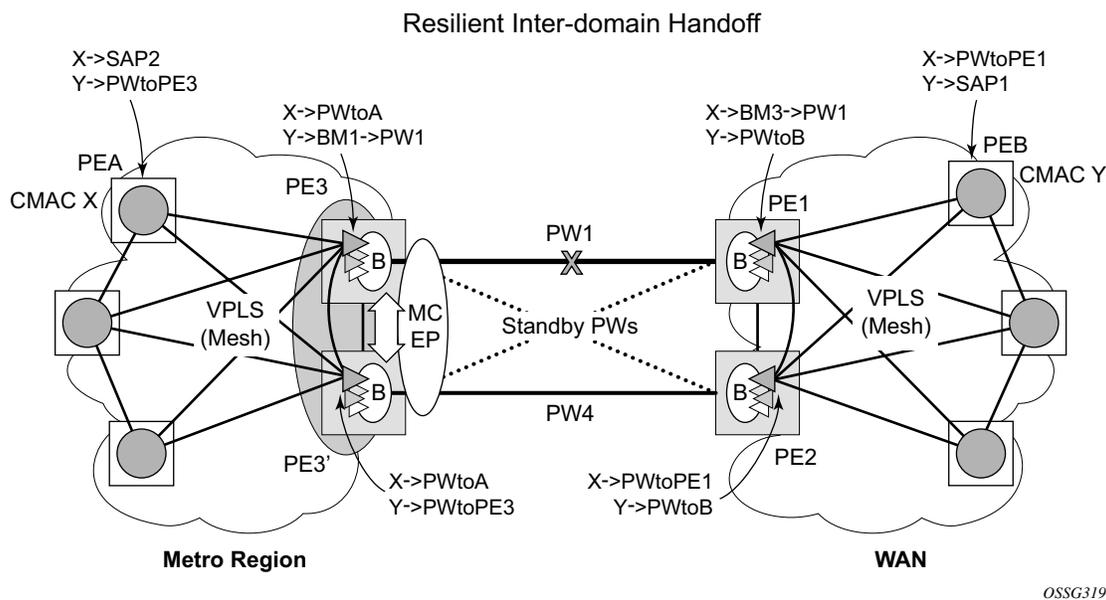


Figure 58: MC-EP with B-VPLS Failure Scenario

In topology displayed in [Figure 58](#), PE A and PE B are regular VPLS PEs participating in the VPLS mesh deployed in the metro and respectively WAN region. As the traffic flows between CEs with CMAC X and CMAC Y, the FIB 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 BVPLS FIBs. The traffic between CMAC X and CMAC Y will be blackholed as long as the entries from the VPLS and I-VPLS FIBs 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 FIB.

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 towards its local blue IVPLS mesh to PE3 and PEA when its MC-Endpoint 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 IVPLS mesh.

For consistency, a similar procedure is used for the BVPLS case as depicted in Figure 59.

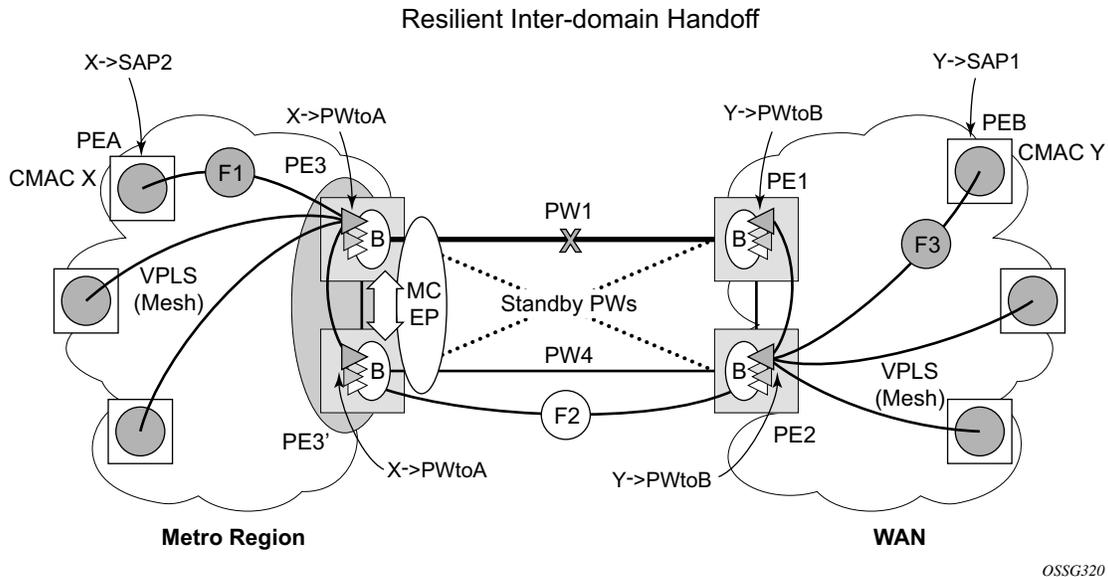


Figure 59: MC-EP with B-VPLS Mac Flush Solution

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 a LDP flush-all-from-me (marked with F1) to PE A and to the other regular VPLS PEs, including PE3. The following command enables this behavior on a per I-VPLS basis: **configure>service>vpls ivpls>send-flush-on-bvpls-failure.**
 - Result: PE A, PE3 and the other local VPLS PEs in the metro clear the VPLS FIB entries associated to PW to PE3.
- PE3 clears the entries associated to PW1 and sends in the B-VPLS context an LDP flush-all-but-mine (marked with F2) towards PE2 on the active PW4.
 - Result: PE2 clears the BVPLS FIB entries not associated with PW4.
- PE2 propagates the MAC flush-all-but-mine (marked with F3) from B-VPLS in the related I-VPLS context(s) towards all participating VPLS PEs – for example, in the blue IVPLS to PE B, PE1. It also clears all the CMAC entries associated with IVPLS pseudowires.

The following command enables this behavior on a per I-VPLS basis:

configure>service>vpls ivpls>propagate-mac-flush-from-bvpls

→ Result: PE B, PE1 and the other local VPLS PEs in the WAN clear the VPLS FIB entries associated to PW to PE2.

→ This command does not control though the propagation in the related IVPLS of the BVPLS 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.

Note that for an SC/MC endpoint configured in a BVPLS, failure/de-activation 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. In other words, connectivity between them exists only over the regular VPLS Mesh.

For the same reasons, it is recommended that static BMAC not be used on SC/MC endpoints.

VPLS Access Redundancy

A second application of hierarchical VPLS is using MTUs that are not MPLS-enabled which 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 mVPLS.
- Non-STP-based access using mechanisms such as MC_LAG, MC-APS, MC-RING.

STP-Based Redundant Access to VPLS

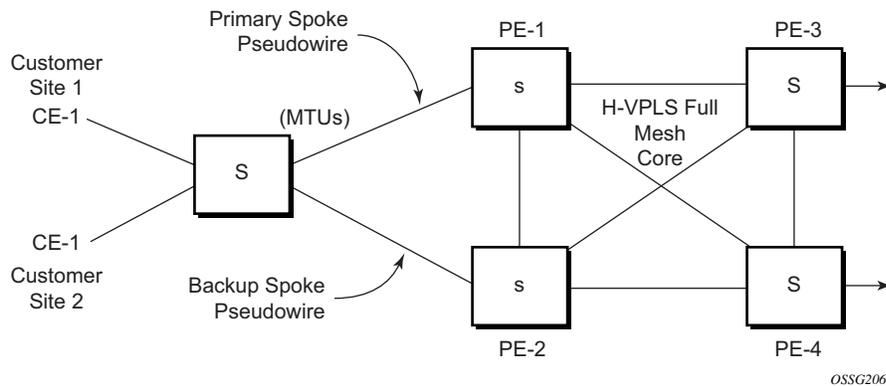


Figure 60: Dual Homed MTU-s in Two-Tier Hierarchy H-VPLS

In configuration shown in [Figure 60](#), STP is activated on the MTU and two PEs in order to resolve a potential loop. Note that STP only needs to run in a single VPLS instance, and the results of the STP calculations are applied to all VPLSes on the link.

In this configuration the scope of 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 case of STP as a 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, but the ones learned from originating PE. Such messages are sent to all PE peers connected through SDPs (mesh and spoke) in the context of VPLS service(s) which are managed by the given STP instance.

Redundant Access to VPLS Without STP

The Alcatel-Lucent implementation also 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 52](#) 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 that to a broadcasting of packets addressing affected hosts and re-learning process in case an alternative route exists.

Note that the message described here is different than the message described in 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 of the given CE device will open alternative link (L2-B switch in Figure 57) as well as on the speed 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 physical failure of the link.

Object Grouping and State Monitoring

This feature introduces a generic operational group object which 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.

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 in the past for LDP VPLS mesh. On the PE selected as the Designated Forwarder (DF), if the rest of the VPLS endpoints fail (pseudowire spoke(s)/pseudowire mesh and/or SAP(s)), 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 object(s) associated with the MH site.

Support for the feature is provided initially in VPLS and BVPLS 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 it 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 oper-group.
- Operational group feature may co-exist in parallel with the **block-on-mesh** feature as long as they are running in different VPLS instances

There are two steps involved in enabling the block on group failure 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** 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 then 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 member(s).
- For BGP-AD or BGP VPLS pseudowire(s) 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 location. 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
```

MAC Flush Message Processing

The previous sections described operation principle of several redundancy mechanisms available in context of VPLS service. All of them rely on MAC flush message as a tool to propagate topology change in a context of the given VPLS. This section aims to summarize basic rules for generation and processing of these messages.

As described on respective sections, the 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 requests clearing of all FDB entries which 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 message requests clearing all FDB entries learned from originating PE. This means that this message has exactly other effect then flush-all-but-mine message. This type is not included in 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 focuses on summarizing actions taken on reception and conditions individual messages are generated.

Upon reception of MAC flush messages (regardless the type) SR-Series PE will take following actions:

- Clears FDB entries of all indicated VPLS services conforming the definition.
- Propagates the message (preserving the type) to all LDP peers, if “propagate-mac-flush” flag is enabled at corresponding VPLS level.

The flush-all-but-mine message is generated under following conditions:

- The flush-all-but-mine message is received from LDP peer and propagate-mac-flush flag is enabled. The message is sent to all LDP peers in the context of VPLS service it was received in.
- 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 given STP instance (based on mVPLS definition). If all LDP peers are in the STP domain, i.e. the mVPLS 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 will not send a withdrawal if the mVPLS does not contain a mesh SDP. A mesh SDP must be configured in the mVPLS to send withdrawals.

- Flush-all-but-mine message is generated when switch over between spoke SDPs of the same endpoint occurs. The message is sent to LDP peer connected through newly active spoke SDP.

The flush-mine message is generated under following conditions:

- The flush-mine message is received from LDP peer and “propagate-mac-flush” flag is enabled. The message is sent to all LDP peers in the context of VPLS service it was received.
- The flush-mine message is generated when on a SAP or SDP transition from operationally up to an operationally down state and send-flush-on-failure flag is enabled in the context of the given VPLS service. The message is sent to all LDP peers connected in the context of the given VPLS service. The send-flush-on-failure flag is blocked in mVPLS and is only allowed to be configured in a VPLS service managed by mVPLS. This is to prevent that both messages are sent at the same time.
- The flush-mine message is generated when on a MC-LAG SAP or MC-APS SAP transition from an operationally up state to an operationally down state. The message is sent to all LDP peers connected in the context of the given VPLS service.
- The flush-mine message is generated when on a MC-RING SAP transition 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 given VPLS service.

Dual Homing to a VPLS Service

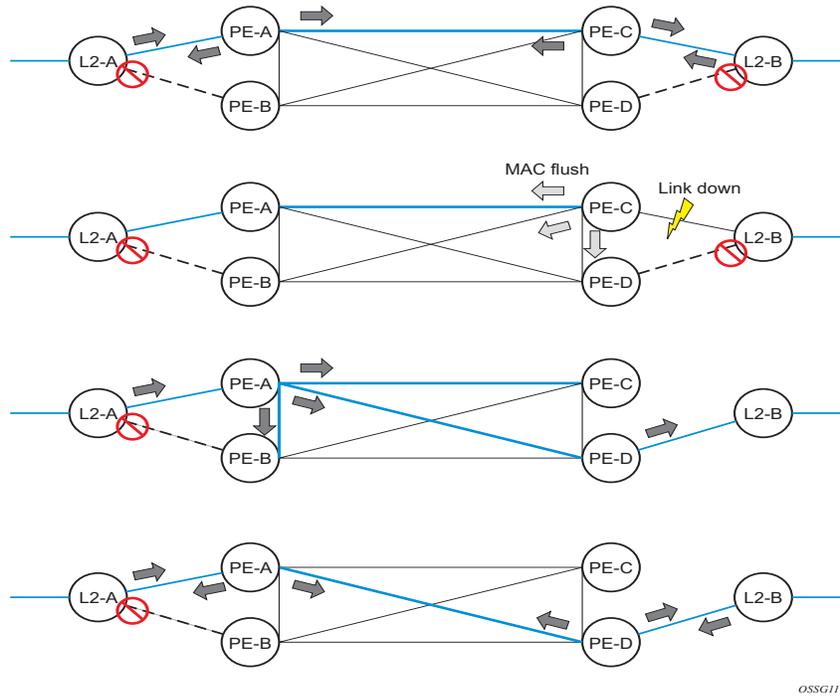


Figure 61: Dual Homed CE Connection to VPLS

Figure 61 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 that to a broadcasting of packets addressing affected hosts and re-learning process in case an alternative route exists.

Note that 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 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, 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 represent 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 of the given CE device will open alternative link (L2-B switch in [Figure 61](#)) as well as on the speed 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 physical failure of the link.

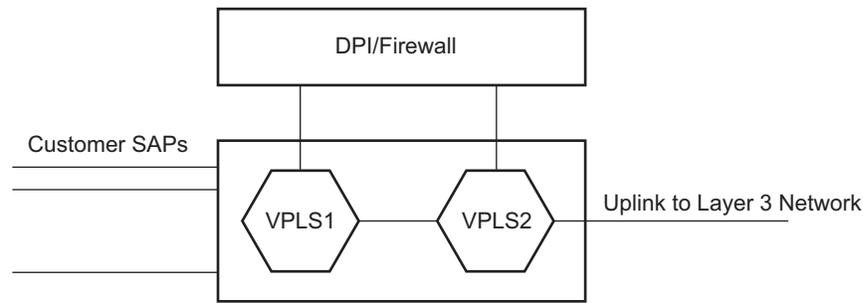
MC-Ring and VPLS

The use of multi-chassis ring control in a combination with the plain VPLS SAP is supported FDB in individual ring nodes in case of the link (or ring node) failure cannot be cleared.

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 mVPLS can be applied.

ACL Next-Hop for VPLS



OSSG207

Figure 62: Application 1 Diagram

The ACL next-hop for VPLS feature enables an ACL that has a forward next-hop 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 to the same gateway to be split and forwarded differently based on the ACL.

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. 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 or lack thereof. The SAP can either terminate a Layer 2 firewall, 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 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.

SDP Statistics for VPLS and VLL Services

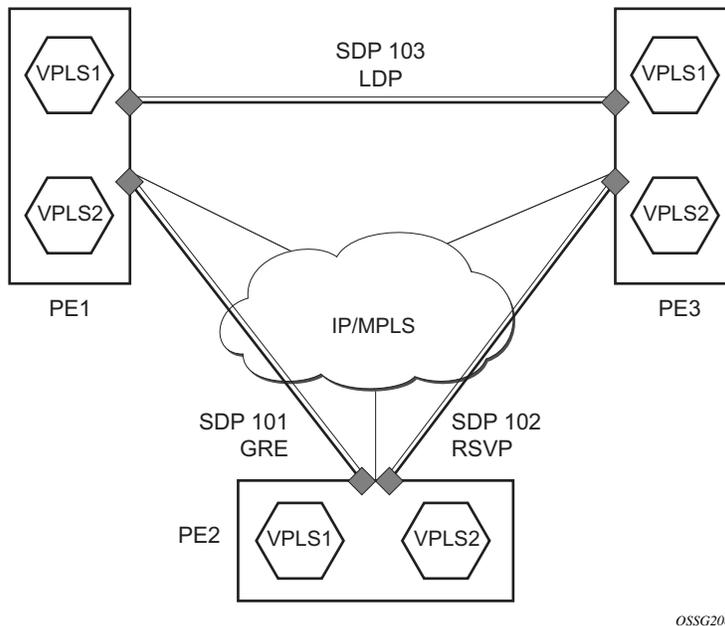


Figure 63: SDP Statistics for VPLS and VLL Services

The simple three-node network described in [Figure 63](#) 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.

SDP statistics allow providers to bill customers on a per-SDP per-byte basis. This destination-based billing model is 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 a bulk manner.

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 is an excellent complement to 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 and to change the signaling protocol.

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 may be found at: <http://www.iana.org/assignments/address-family-numbers> and SAFI <http://www.iana.org/assignments/safi-namespace>.

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, unique network wide identifier; same value assigned for all VPLS switch instances (VSIs) belonging to the same VPLS; 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 bytes identifier (usually the system IP of the VPLS PE); encoded and carried in the corresponding BGP NLRI.

In order 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 take care of signaling the service labels later on.

The format of the BGP AD NLRI is very similar with the one used for IP VPN as depicted in [Figure 64](#). The system IP may be used for the last 4 bytes of the VSI ID further simplifying the addressing and the provisioning process.

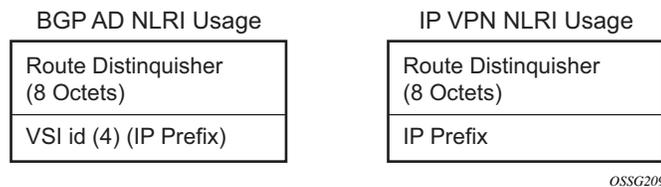


Figure 64: BGP AD NLRI versus IP VPN NLRI

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 a 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).

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 it was used extensively in the initial VPWS and VPLS deployments. The simple format is easy to understand but it does not provide the required information model for BGP auto-discovery function. In order 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 respectively, Target Attachment Individual Identifier (TAII). 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 depicted in [Figure 65](#).

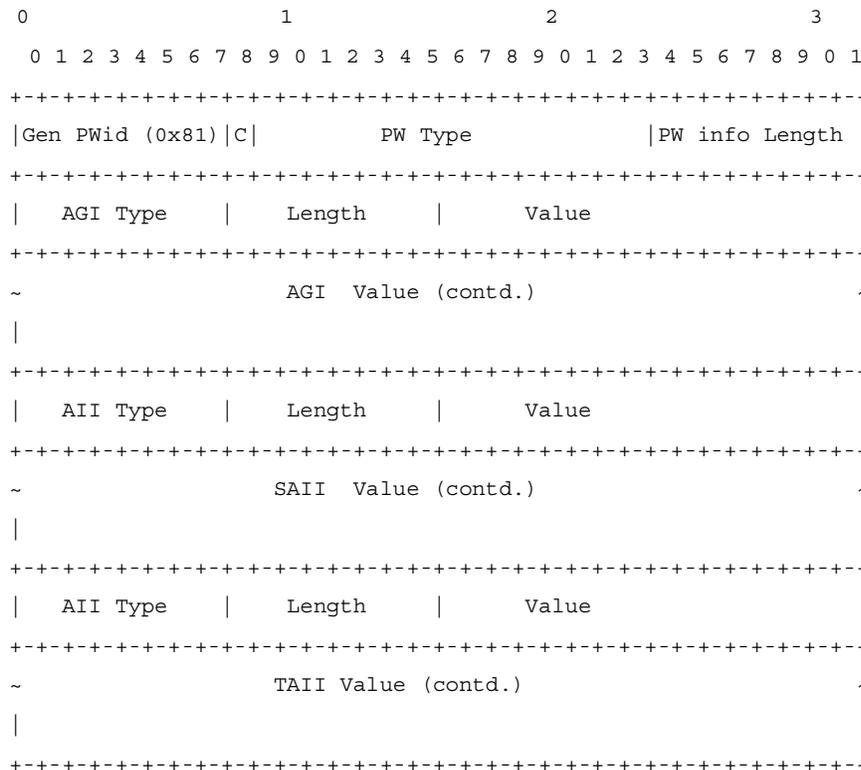


Figure 65: Generalized Pseudowire-ID FEC Element

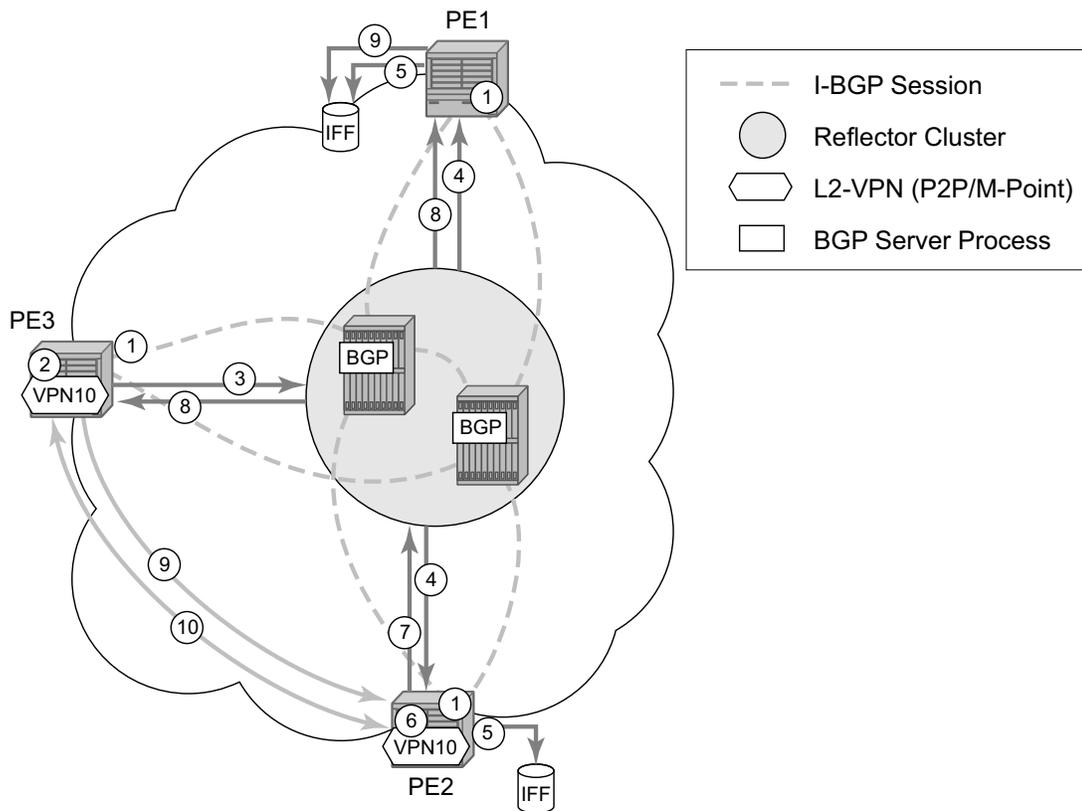
Each of the FEC fields are designed as a sub-TLV equipped with its own type and length providing support for new applications. To accommodate the BGP AD information model the following FEC formats are used:

- AGI (type 1) is identical in format and content with the BGP extended community attribute used to carry the VPLS-ID value.
 - Source AII (type 1) is a 4 bytes value destined to carry the local VSI-id (outgoing NLRI minus the RD).
 - Target AII (type 1) is a 4 bytes value destined to carry the remote VSI-ID (incoming NLRI minus the RD).
-

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 66](#).



OSSG210

Figure 66: 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.

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 OS MPLS Guide for more information about MPLS, LDP, and RSVP.

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 gleaned 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 as long as 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 in order to 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.

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 gleaned 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.

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 VSIs 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 great care as they may impact many customers at once.

The Alcatel-Lucent 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, then 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.

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 point in 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.

Resiliency Schemes

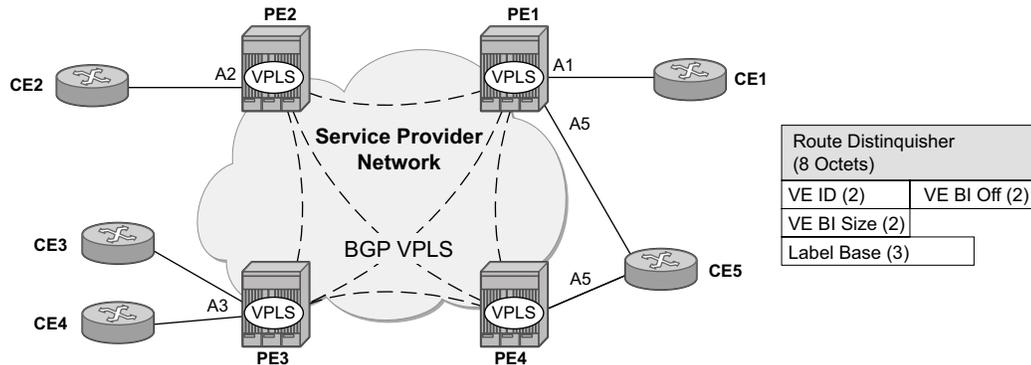
The use of BGP AD on the network side, or in the backbone, does not affect the different resiliency schemes Alcatel-Lucent has developed in the access network. This means that both Multi-Chassis Link Aggregation (MC-LAG) and Management-VPLS (M-VPLS) can still be used.

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.

BGP VPLS

The Alcatel-Lucent BGP VPLS solution, compliant with RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*, is described in this section.



OSSG488

Figure 67: BGP VPLS Solution

Figure 67 depicts 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 two (2) bytes VE ID assigned to the VPLS instance.
 - BGP-AD for LDP VPLS: 4 bytes 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. Same AFI, SAFI 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. As a result the BGP NLRI content includes also 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, sequencing.
-

Pseudowire Signaling Details

The pseudowire is setup 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.
 - Alcatel-Lucent implementation uses always a value of eight (8).
- Label Base (LB): local allocated label base.
 - The next eight (8) labels allocated for remote PEs.

This BGP update is telling the other PE(s) that accept the RT: “in order to reach me (VE-ID = x) use a pseudowire label of $LB + VE-ID - VBO$ using the BGP NLRI for which $VBO \leq \text{local VE-ID} < VBO + VBS$.”

Here is an example of how this algorithm works assuming PE1 has VE-ID 7 configured:

- PE1 finds a Label Block of eight (8) consecutive labels available, starting with LB = 1000
- PE1 then starts sending 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 one (1) to eight (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 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 now 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-ID 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 towards the originator VPLS PE.

The format is described below:

```

+-----+
| 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:

- 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, see below for details.
- Layer-2 MTU is 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.

The detailed format for the Control Flags bit vector is described below:

```

0 1 2 3 4 5 6 7
+-----+
|D| MBZ      |C|S| (MBZ = MUST Be Zero)
+-----+
    
```

The following bits in the Control Flags are defined:

- 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, Alcatel-Lucent 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 & 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 the “config service vpls shutdown”
- 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 idea 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 deleted — a BGP Update with unreachable-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 pseudowire(s) mapped to that SDP goes down. There is no BGP-update sent.

Supported VPLS Features

BGP VPLS just added support for a new type of pseudowire signaling based on MP-BGP. It is based on the existing VPLS instance hence it inherited all the existing Ethernet switching functions. Here are some of the most important existing VPLS features ported also to BGP VPLS:

- VPLS data plane features: for example FIB management, SAPs, LAG access, BUM rate limiting.
- MPLS tunneling: LDP, LDP over RSVP-TE, RSVP-TE, MP-BGP based on RFC3107 (Option C solution)
- HVPLS topologies, Hub and Spoke traffic distribution
- Coexists with LDP VPLS (with or without BGP-AD) in the same VPLS instance.
→ LDP, BGP-signaling should operate in disjoint domains to simplify loop avoidance

BGP VPLS

- Coexist with BGP-based multi-homing.
- BGP VPLS is supported as the control plane for BVPLS.
- Supports IGMP/PIM snooping
- 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

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 iVPLS, bVPLS, and rVPLS)
- 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 command **config>service>vpls>spoke-sdp>bfd-template** *name* 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 command **config>service>vpls>bgp>pw-template-binding>bfd-template** *name* and enable it using the command **config>service>vpls>bgp>pw-template-binding>bfd-enable**.

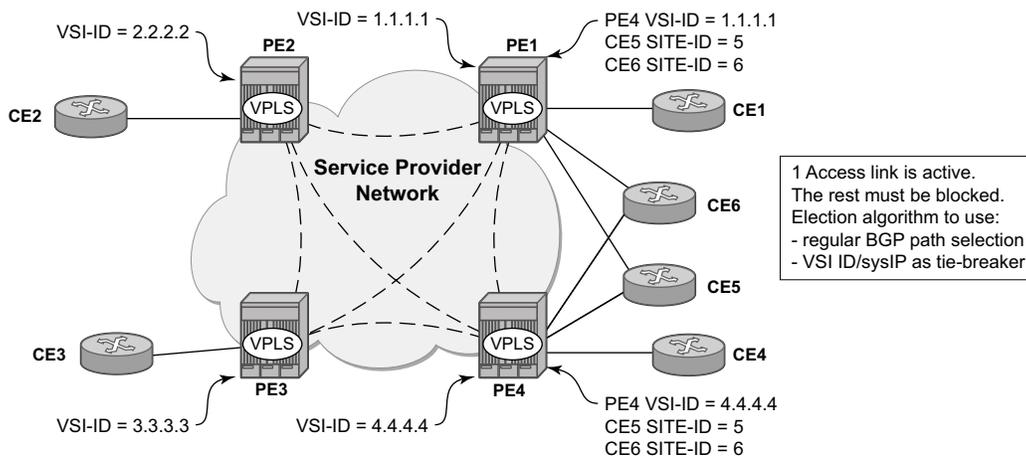
For BGP-AD VPLS, a BFD template is referenced from the pseudowire template context. To configure VCCV BFD for BGP-AD, use the command **config>service>vpls>bgp-ad>pw-template-binding>bfd-template** *name* and enable it using the command **config>service>vpls>bgp-ad>pw-template-binding>bfd-enable**.

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 PE(s) 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 towards the core PEs and access PEs or CEs
- The interaction with pseudowire signaling (BGP/LDP)



OSSG489

Figure 68: BGP Multi-Homing for VPLS

Figure 68 depicts the VPLS using BGP Multi-homing for the case of multi-homed CEs. Although the picture depicts the case of a pseudowire infrastructure signaled with LDP for a LDP VPLS using BGP-AD for discovery, the procedures are identical for BGP VPLS or for a mix of BGP and LDP signaled pseudowires.

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 69](#), CE5 will be assigned the same site-id on both PE1 and PE4. For the same VPLS instance though, 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 NRIs in [Figure 69](#).

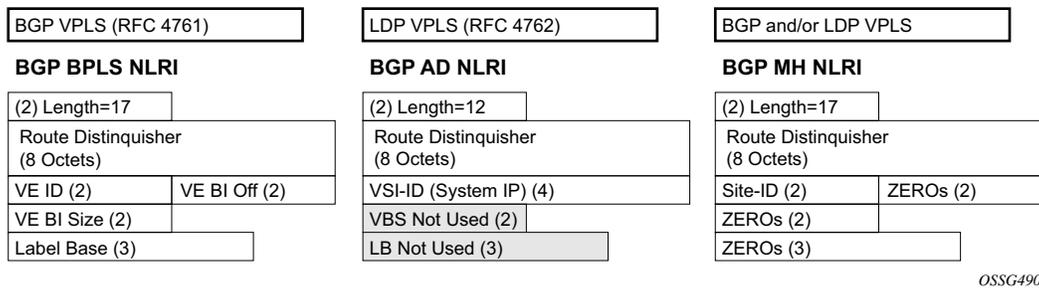


Figure 69: BGP MH-NLRI for VPLS Multi-Homing

The BGP VPLS NLRI described in RFC 4761 is used to carry a two (2) byte site-ID that identifies the MH Site. The last seven (7) 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 setup 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. Alcatel-Lucent BGP-based multi-homing solution uses the DF election procedure described in the IETF working group document *draft-ietf-l2vpn-vpls-multihoming*. 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.

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.

Supported Services and Multi-Homing Objects

This feature is supported for the following services:

- LDP VPLS with or without BGP-AD
- BGP VPLS
- mix of the above
- PBB BVPLS on BCB
- PBB I-VPLS (see the *IEEE 802.1ah PBB Guide* 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:
SAP(s), pseudowires (BGP VPLS, BGP-AD, provisioned and LDP signaled spoke SDP and mesh SDP)

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.

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) towards 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 object(s) going down or more generically, the entire site going from Designated Forwarder (DF) to non-DF.
 - Note that 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.
-

Indicating non-DF status towards 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 de-activation is linked to a CCM failure on a SAP that has a down MEP configured.

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 depicted in [Figure 70](#).

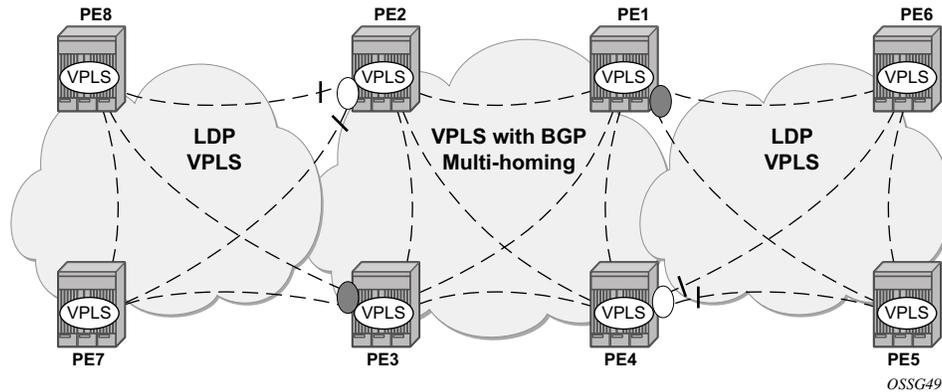


Figure 70: BGP MH Used in an HVPLS Topology

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 indicate the number of pseudowire failures that are required for the MH site to be declared down.

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.

In order to improve this condition, VPLS is IP multicast aware so it forwards IP multicast traffic based on multicast states.

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.

Plain PIM Snooping

In plain PIM snooping configuration, VPLS PE routers only snoop, PIM messages generated on their own. Join/prune suppression must be disabled on CE routers.

When plain PIM snooping is configured, a VPLS PE router detects a condition where join/prune suppression is not disabled on one or multiple CE routers, the PE router should put PIM snooping into PIM proxy state. A trap is generated which reports the condition to the operator and is logged to syslog. If the condition changes, for example, join/prune suppression was disabled on CE routers, the PE reverts to plain PIM snooping state. A trap is generated and is logged to syslog.

PIM Proxy

For PIM proxy configurations, VPLS PE routers perform the following:

- Snoop hellos and flood hellos in 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 PE(s) 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, but 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 syslog. If the condition changes, for example, join/prune suppression enabled on a CE router, PIM proxy is placed back into operational state. Again, a trap is generated to report the condition to the operator and is logged to syslog.

Multicast Listener Discovery (MLD) Snooping and MAC-Based Multicast Forwarding

VPLS-based transport is a popular architecture as it better handles IPv6 multicast on the transport configurations for those backbones who use IPv6 instead of IPv4.

The VPLS based transport architecture combines MLD snooping and MAC based multicast forwarding.

MLD Snooping

MLD snooping is basically a IPv6 version of IGMP snooping. The guidelines and procedures are similar to IGMP snooping as well.

MAC-Based IPv6 Multicast Forwarding

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.

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 . Thus, when both MLD snooping and native IPv6 multicast are enabled on the same device, both formats are supported on the same XMA, although they are used for different services.

PIM and IGMP Snooping Interaction

This section describes how to handle the scenario where IGMP snooping and PIM snooping are both enabled for the same VPLS.

When both PIM snooping and IGMP snooping are enabled for a VPLS, multicast traffic is forwarded based on the combined multicast forwarding table.

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.

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 given 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/B-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. Note that 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

The allow-ip-int-bind VPLS Flag

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.

Note that 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.

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.

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.

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.

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.

Routed VPLS Feature Restrictions

When the **allow-ip-int-bind** flag is set on a VPLS service, the following features cannot be enabled (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.
- 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.
- Network domain on SDP bindings.
- Per Service Hashing not supported
- No BGP-VPLS
- IOM3+ cards only

Note: IES/VPRN Saps can be on non IOM3+ cards but traffic on them will not be forwarding on Routed VPLS/Routed I-VPLS

- No Time of Day accounting on Routed VPLS SAPs.
 - No Ingress Queuing for Split-Horizon Groups
 - No Multiple Virtual Router support
-

Routed I-VPLS Feature Restrictions

- No Multicast support
- No VC-VLAN on SDPs
- force-qtag-forwarding is Not supported
- No Control word on B-VPLS SDPs with Routed I-VPLS
- No Hash Label on B-VPLS SDPs with Routed I-VPLS

IES IP Interface VPLS Binding and Chassis Mode Interaction

It is possible to bind both IES and VPRN IP interfaces to a VPLS in chassis mode A. Chassis - mode D is not required.

VPRN IP Interface VPLS Binding and Forwarding Plane Constraints

When an IP interface within a VPRN service context is bound to a VPLS or an I-VPLS service name, all of the SAPs within the VPRN service context must be created on ports that are attached to FP2 forwarding planes or better. If a VPRN SAP is on a non-supported forwarding plane, the service name cannot be bound to the VPRN's IP interface. Once an IP interface on the VPRN service is bound to a service name, a SAP on the VPRN service cannot be created on a port (or LAG) on an FP1 forwarding plane.

This restriction prevents a packet from entering the VPRN service on a port that cannot reach a routed VPLS next-hop.

Route Leaking Between Routing Contexts

While the system prevents a routing context from existing on FP1 based forwarding planes while a VPLS service is bound to the routing context, it is possible to create conditions using route leaking (importing or exporting routes using routing policies) where an FP1 based IP interface is asked to route to a routed VPLS next-hop. The system reacts to this condition by populating the next-hop in the FP1 forwarding plane with a null egress IP interface index. This causes any packets that are associated with that next-hop on an FP1 forwarding plane to be discarded. If ICMP destination unreachable messaging is enabled, unreachable messages will be sent.

Ingress LAG and FP1 to Routed VPLS Discards

If the chassis is connected by LAG to an upstream router and the LAG is split between FP1 and FP2 forwarding plane ports while routes have been shared between routing contexts, flows that are sent to the FP2 ports by the upstream router are capable of reaching a next-hop in a routed VPLS while flows going to the FP1 ports cannot.

IPv4 Multicast Routing Support

IPv4 multicast routing is supported in a routed VPLS service when the source of the multicast stream is on the IP side and when the source is on the VPLS side. 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.

The IP interface of a routed VPLS supports the configuration of both PIM and IGMP for IPv4 multicast.

To enable the ability to forward IPv4 multicast traffic from the VPLS side of the routed VPLS service to the IP side, the `forward-ipv4-multicast-to-ip-int` parameter must be configured as shown below:

```
configure
  service
    vpls <service-id>
      allow-ip-int-bind
        forward-ipv4-multicast-to-ip-int
      exit
    exit
  exit
```

Enabling IGMP snooping in the VPLS service is optional. If IGMP snooping is enabled, it is mandatory to enable IGMP on the routed VPLS IP interface in order for multicast traffic to be sent into, or received from, the VPLS service.

If both IGMP and PIM are configured on the routed VPLS IP interface, it is necessary to configure the associated IP interface to be both the PIM designated router and the IGMP querier in order that the multicast traffic is sent into the VPLS service, as IGMP joins are only propagated to the IP interface if it is the IGMP 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
        exit
      exit
    exit
  exit
```

This is useful to achieve 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 snooping is enabled in the VPLS service. On failure of the active router, this configuration avoids the remaining router having to wait until it sends an IGMP query into the VPLS service before it starts

receiving IGMP joins, and consequently sending the multicast traffic into the VPLS service. When the mrouter port is configured as above, all IGMP joins (and multicast traffic) are sent to the VPLS service IP interface.

IGMP snooping should only be enabled when systems, as opposed to PIM routers, are connected to the VPLS service. If IGMP snooping is enabled when the VPLS service is used for transit traffic for connected PIM routers, the IGMP 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` configured, the IPv4 multicast traffic received in the VPLS service which is forwarded through the IP interface will be IMPM managed even without IGMP snooping being enabled. This does not apply to traffic which is only flooded within the VPLS service.

When IPv4 multicast traffic is forwarded from a VPLS SAP or spoke SDP to both VPLS and IP outbound interfaces, for example to another VPLS SAP or spoke SDP and to a regular 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

An example scenario is shown in [Figure 71](#). 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.

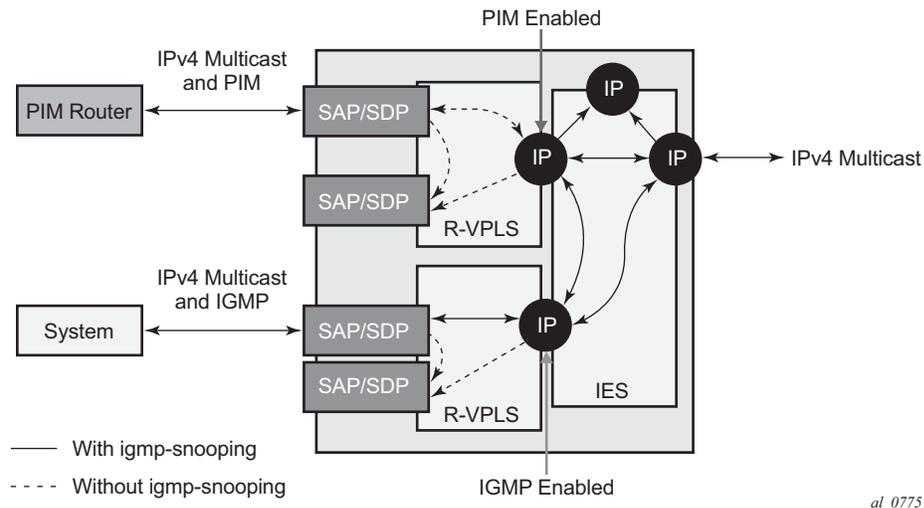


Figure 71: IPv4 Multicast with a Router VPLS service

The IPv4 multicast traffic entering the IES service through the regular IP interface is replicated to both the other regular IP interface and the two routed VPLS interfaces if PIM/IGMP joins have been received on the respective IP interfaces. This traffic will be flooded into both VPLS services unless IGMP snooping is enabled in the lower VPLS service, in which case it is only sent to the system originating the IGMP join.

The IPv4 multicast traffic entering the upper VPLS service from the PIM router will be flooded in that VPLS service and, if related joins have been received, forwarded to the regular IP interfaces. It will also be forwarded to the lower VPLS service if an IGMP join is received on its IP interface, again being flooded or not in that VPLS service depending on whether IGMP snooping is enabled.

The IPv4 multicast traffic entering the lower VPLS service from the system will be flooded in that VPLS service, unless IGMP 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 service if related joins have been received on those interfaces and it will also be forwarded to the upper VPLS service if a PIM join is received on its IP interface, this being flooded in that VPLS service.

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.

Routed VPLS Caveats

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 given packet type (IPv4 or IPv6) is not overridden, the SAP specified filter is applied to the packet (if defined).

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.

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.

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.
 - 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.
-

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

Spanning Tree and Split Horizon

A routed VPLS context supports all spanning tree and split horizon capabilities that a non-routed VPLS service supports.

VPLS Service Considerations

This section describes the and 7950 XRS service features and any special capabilities or considerations as they relate to VPLS services.

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 and 7950 XRS VPLS service:

- Ethernet null
 - Ethernet Dot1q
 - Ethernet QinQ
-

VLAN Processing

The SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs to:

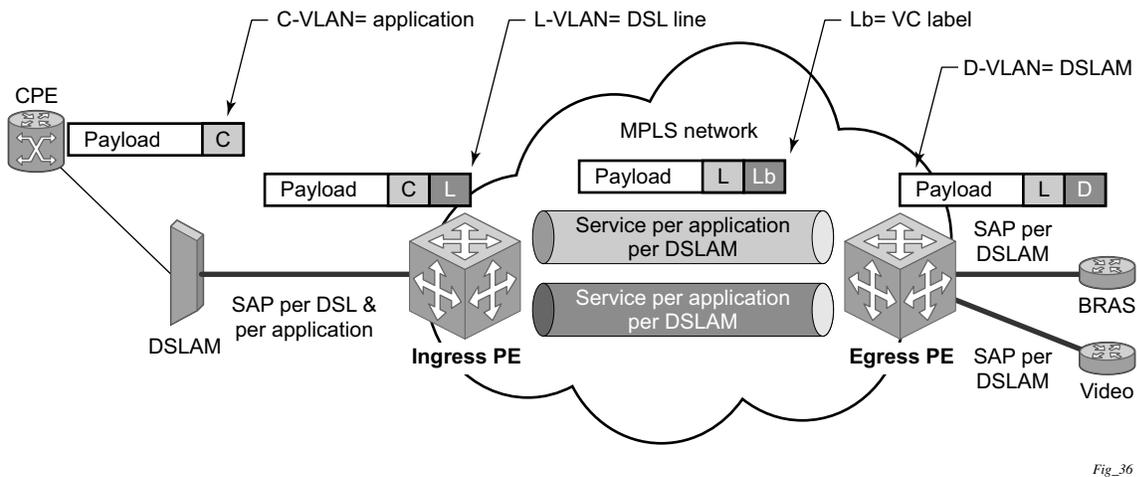
1. Null encapsulation defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP.
2. Dot1q encapsulation defined on ingress — Only first label is considered.
3. QinQ encapsulation defined on ingress— Both labels are considered.
Note that the SAP can be defined 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.

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.



Fig_36

Figure 72: Ingress VLAN Swapping

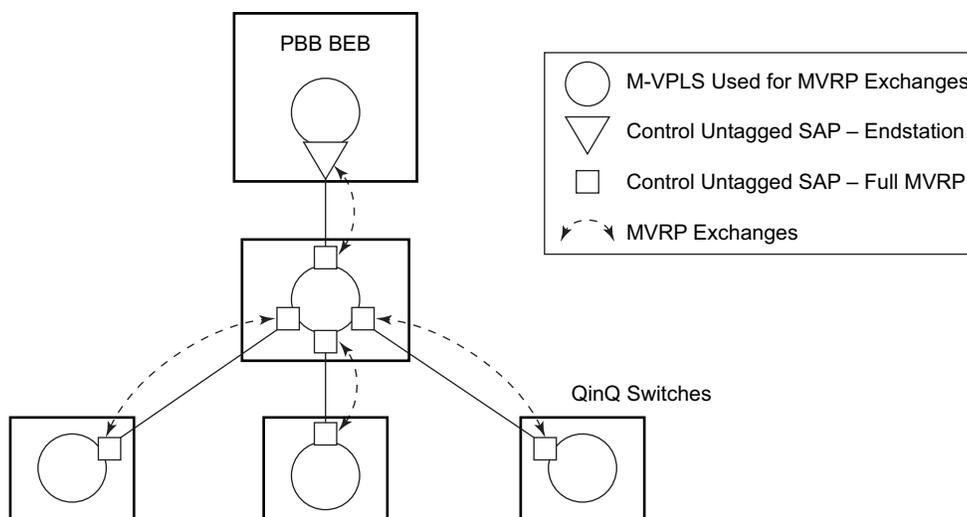
The network diagram describes the network where at user access side (DSLAM facing SAPs) every subscriber is represented by several QinQ SAPs with inner-tag encoding service and outer-tag encoding subscriber (DSL line). The aggregation side (BRAS or PE facing SAPs) the is represented by DSL line number (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.

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 (see IEEE 802.1ad) or Backbone VLAN IDs (BVID) in a Provider Backbone Bridging (PBB) domain (see IEEE 802.1ah).

The initial focus of Alcatel-Lucent 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 depicted in [Figure 73](#).



OSSG492

Figure 73: 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 FIB, trunks in the MVRP, M-VPLS scope
 - The VLAN FIBs (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. IVPLS 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

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 SAP(s) 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

Instantiate Related VLAN FIBs and Trunks in MVRP Scope

This involves the configuration in the M-VPLS, under vpls-group of the following attributes: VLAN range(s), 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 FIB and related SAP(s) 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 SAP(s) 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-control** command under vpls-group:

- If mvrp-control is disabled the auto-created VPLS instance(s) and related SAP(s) 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 enable 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 FIBs.

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 attribute(s) 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 FIBs and related trunks in non-PBB switches:

```

mrp
  no shutdown
  mvrp
    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
  vpls-sap-template-binding Autosapl
  mvrp-control
  no shutdown

```

A similar M-VPLS configuration may be used to auto-instantiate the VLAN FIBs 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:

```

config>service>vpls control-mvrp m-vpls create customer 1
[.]
sap 1/1/1:0
  mrp mvrp
    endstation-vid-group 1 vlan-id 100-2000
  no shutdown

```

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 IVPLS 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 depicted in [Figure 74](#).

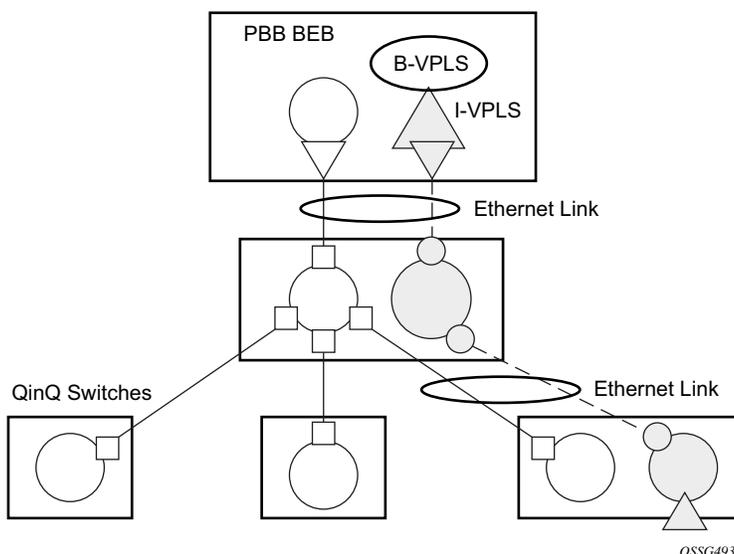


Figure 74: Service Instantiation with MVRP - QinQ to PBB Example

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 IVPLS SAP must be configured towards 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 depicted 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.

Note that 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.

Service Auto-Discovery using Multiple VLAN Registration Protocol (MVRP)

In the data VPLS instance (VLAN 100) controlled by MVRP, on the QinQ switch:

```
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:

```
config>service>vpls 1000 i-vpls
    sap 8/1/2:100 //sap (using MVRP VLAN 100 on endstation port in
    VPLS.)
    no shutdown
```

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 Alcatel-lucent implementation. Alcatel-Lucent MVRP may be used also in a non- MSTP, loop free topology.

STP-MVRP Interaction

The following table captures the expected interaction between STP (MSTP or P-MSTP) and MVRP:

Table 11: 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.

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

Note that 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.

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.

In order 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 FIB 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. Note that 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 FIB 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.

VPLS E-Tree Services

This section describes the following topics:

- [VPLS E-Tree Services Overview on page 476](#)
 - [Leaf-ac and Root-ac SAPs on page 477](#)
 - [Leaf-ac and Root-ac SDP Binds on page 478](#)
 - [Root-leaf-tag SAPs on page 478](#)
 - [Root-leaf-tag SDP Binds on page 479](#)
 - [Interaction between VPLS E-Tree Services and Other Features on page 480](#)
-

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-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 75](#) 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.

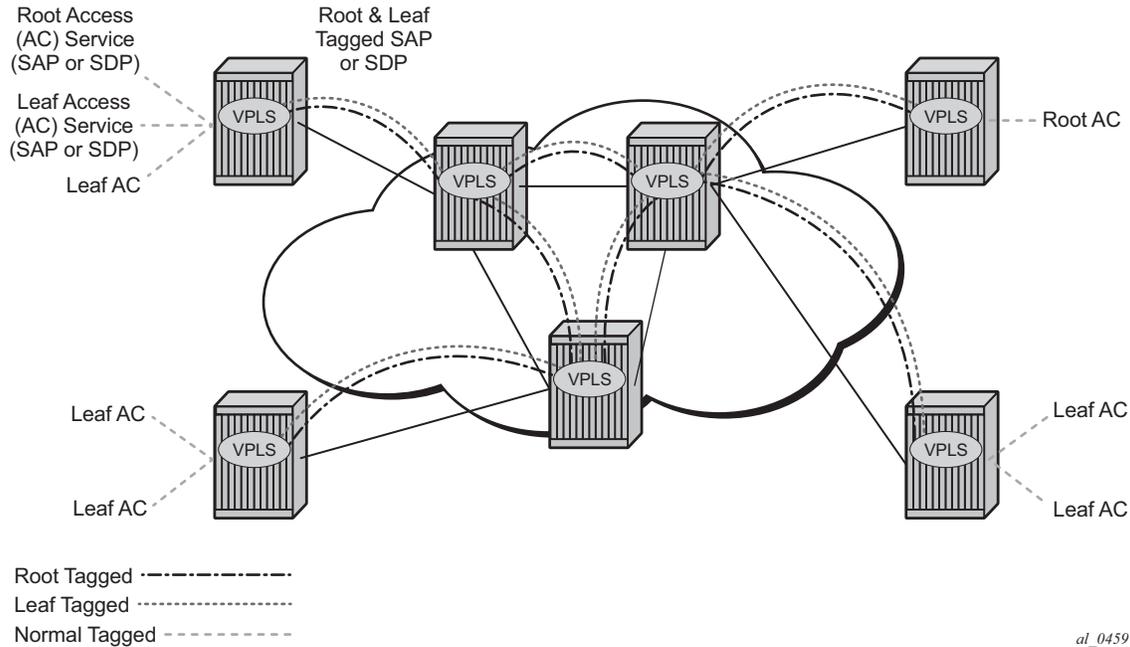


Figure 75: E-Tree Service

Leaf-ac and Root-ac SAPs

Figure 76 illustrates the terminology used for E-Tree in draft-ietf-l2vpn-vpls-pe-etree and a mapping to SROS 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 SROS is concerned, these are normal SAPs with either no tag (Null)/ priority tag or dot1Q or QinQ encapsulation on the frame. Note that, 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.

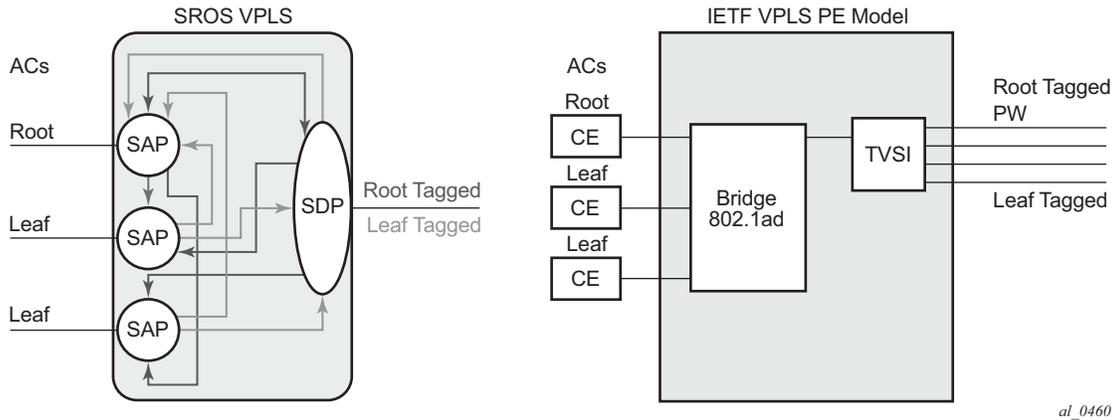


Figure 76: Mapping PE Model to 7x50 VPLS Service

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.

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 77 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
```

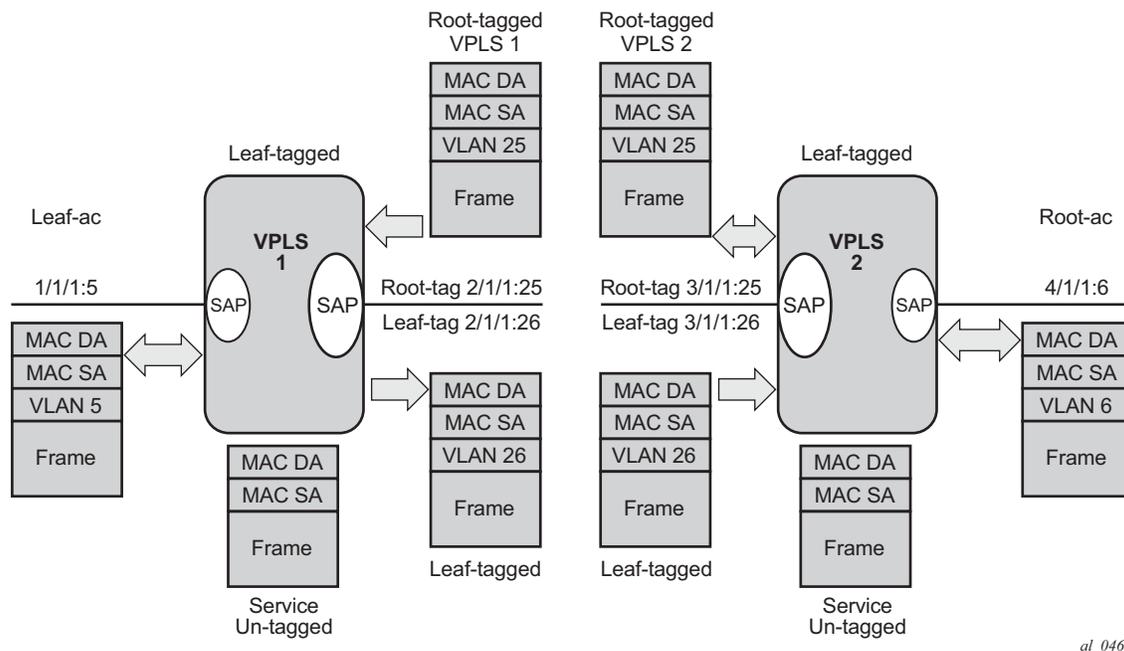


Figure 77: Leaf and Root Tagging Dot1q

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.

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 SROS implementation uses a fixed VLAN ID 1 for root and fixed VLAN ID 2 for leaf. The root and leaf tags are considered a global value and signaling is not supported. Note that 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 78 illustrates the behavior when leaf-ac or root-ac interfaces exchange traffic over a root-leaf-tag SDP-binding.

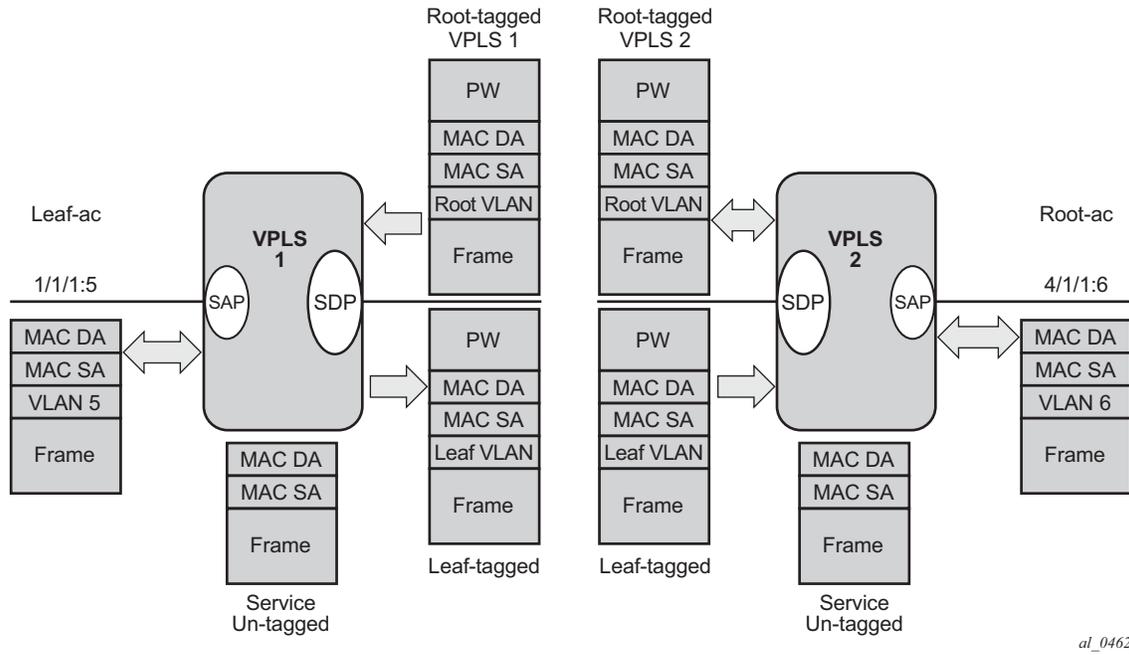


Figure 78: Leaf and Root Tagging PW

al_0462

Interaction between VPLS E-Tree Services and Other Features

As a general rule, any CPM-generated traffic is always root traffic (STP, OAM, etc.) 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 work on VPLS E-Tree services. Note that multicast 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, it is recommended 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.

Configuring a VPLS Service with CLI

This section provides information to configure VPLS services using the command line interface.

Topics in this section include:

- [Basic Configuration on page 484](#)
- [Common Configuration Tasks on page 486](#)
 - [Configuring VPLS Components on page 487](#)
 - [Creating a VPLS Service on page 489](#)
 - [Configuring a VPLS SAP on page 500](#)
- [Configuring VPLS Redundancy on page 523](#)
 - [Creating a Management VPLS for SAP Protection on page 523](#)
- [Configuring VPLS E-Tree Services on page 538](#)
- [Service Management Tasks on page 539](#)
 - [Modifying VPLS Service Parameters on page 539](#)
 - [Modifying Management VPLS Parameters on page 540](#)
 - [Deleting a VPLS Service on page 542](#)
 - [Disabling a VPLS Service on page 542](#)
 - [Re-Enabling a VPLS Service on page 543](#)

Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPLS service:

- Customer ID (refer to the *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 displays 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 displays 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#
```

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 egress multicast groups (optional):

1. Define egress multicast group name(s)
2. Specify the destinations per pass
3. Define SAP common requirements

For VPLS services:

1. Associate VPLS service with a customer ID
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)
 - Optional — Specify SAP egress multicast-group name
3. Associate SDPs for (distributed services)
4. Modify STP default parameters (optional) (see [VPLS and Spanning Tree Protocol on page 383](#))
5. Enable service

Configuring VPLS Components

Use the CLI syntax displayed below to configure the following entities:

- [Configuring Egress Multicast Groups on page 488](#)
- [Creating a VPLS Service on page 489](#)
 - [Enabling MAC Move on page 492](#)
- [Configuring a VPLS SAP on page 500](#)
 - [Local VPLS SAPs on page 500](#)
 - [Distributed VPLS SAPs on page 501](#)
 - [Configuring SAP-Specific STP Parameters on page 503](#)
 - [STP SAP Operational States on page 507](#)
 - [Configuring VPLS SAPs with Split Horizon on page 509](#)
 - [Configuring Overrides on Service SAPs on page 511](#)
- [Configuring SDP Bindings on page 510](#)
 - [Mesh SDP on page 512](#)
 - [Spoke SDP on page 513](#)
- [Configuring VPLS Redundancy on page 523](#)

Configuring Egress Multicast Groups

Use the following CLI syntax to configure egress multicast groups:

```
CLI Syntax: config>service# egress-multicast-group group-name
                description description-string
                dest-chain-limit [destinations per pass]
                sap-common-requirements
                dot1q-etype 0x0600..0xffff
                egress-filter [ip ip-filter-id]
                egress-filter [ipv6 ipv6-filter-id]
                egress-filter [mac mac-filter-id]
                qinq-etype [0x0600..0xffff]
                qinq-fixed-tag-value tag-value
```

The following example displays an egress multicast group configuration:

```
A:ALA-48>config>service>egress-multicast-group# info
-----
                dest-chain-limit 10
                sap-common-requirements
                dot1q-etype 0x060e
                egress-filter ip 10
                exit
-----
A:ALA-48>config>service>egress-multicast-group#
```

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 displays 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#
```

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 displays 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 mmp 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 mmp

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 mmp

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#

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 displays 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#

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:

$$2 \times (\text{Bridge_Forward_Delay} - 1.0 \text{ seconds}) \geq \text{Bridge_Max_Age}$$
$$\text{Bridge_Max_Age} \geq 2 \times (\text{Bridge_Hello0_Time} + 1.0 \text{ seconds})$$

The following STP parameters can be modified at VPLS level:

- [Bridge STP Admin State on page 494](#)
- [Mode on page 495](#)
- [Bridge Priority on page 495](#)
- [Max Age on page 496](#)
- [Forward Delay on page 496](#)
- [Hello Time on page 497](#)
- [MST Instances on page 498](#)
- [MST Max Hops on page 498](#)
- [MST Name on page 498](#)
- [MST Revision on page 498](#)

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 or 7950 XRS. When STP on the VPLS is administratively enabled, but the administrative state of a SAP is down, BPDUs received on such a SAP 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 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 on page 383](#) 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 on page 506](#)).

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 on page 496](#).

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 its own tree within the region, thus 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.

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 displays 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#
```

Configuring a VPLS SAP

A default QoS policy is applied to each ingress 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 on page 500](#)
 - [Distributed VPLS SAPs on page 501](#)
-

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 displays 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#
```

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, see the *Services Overview Guide*. For SDP binding information, see [Configuring SDP Bindings on page 510](#).

The following example displays 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
```

Configuring VPLS Components

```
        stp
          shutdown
        exit
      sap 1/1/3:33 create
        description "VPLS SAP"
        multi-service-site "West"
      exit
    exit
  ...
-----
*A:ALA-3>config>service#
```

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 on page 503](#)
 - [SAP Virtual Port Number on page 504](#)
 - [SAP Priority on page 504](#)
 - [SAP Path Cost on page 505](#)
 - [SAP Edge Port on page 505](#)
 - [SAP Auto Edge on page 506](#)
 - [SAP Link Type on page 506](#)
-

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 towards 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 its 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 — 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.

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 on page 504](#) 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 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 thus 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 on page 496](#)). 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 on page 505](#)).

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

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 on page 507](#)
 - [Operationally Discarding on page 507](#)
 - [Operationally Learning on page 507](#)
 - [Operationally Forwarding on page 508](#)
-

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 proper STP state for each SAP while not forwarding actual user traffic. The duration of the discarding state is explained in section [Forward Delay on page 496](#).

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 FIB. Layer 2 frames received on other forwarding interfaces and destined for the SAP are also forwarded.

SAP BPDUs Encapsulation State

Each SAP has a Read-Only operational state that shows which BPDUs 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 BPDUs encapsulation state until a PVST encapsulated BPDUs is received in which case, the SAP will convert to the PVST encapsulation state. Each received BPDUs 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 BPDUs encapsulation state until a dot1d encapsulated BPDUs is received, in which case, the SAP reverts to the dot1d encapsulation state. Each received BPDUs 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 BPDUs encapsulation state for SAPs defined on Ethernet interface with encapsulation type set to null.

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 displays 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#
```

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 on page 522](#)).

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.

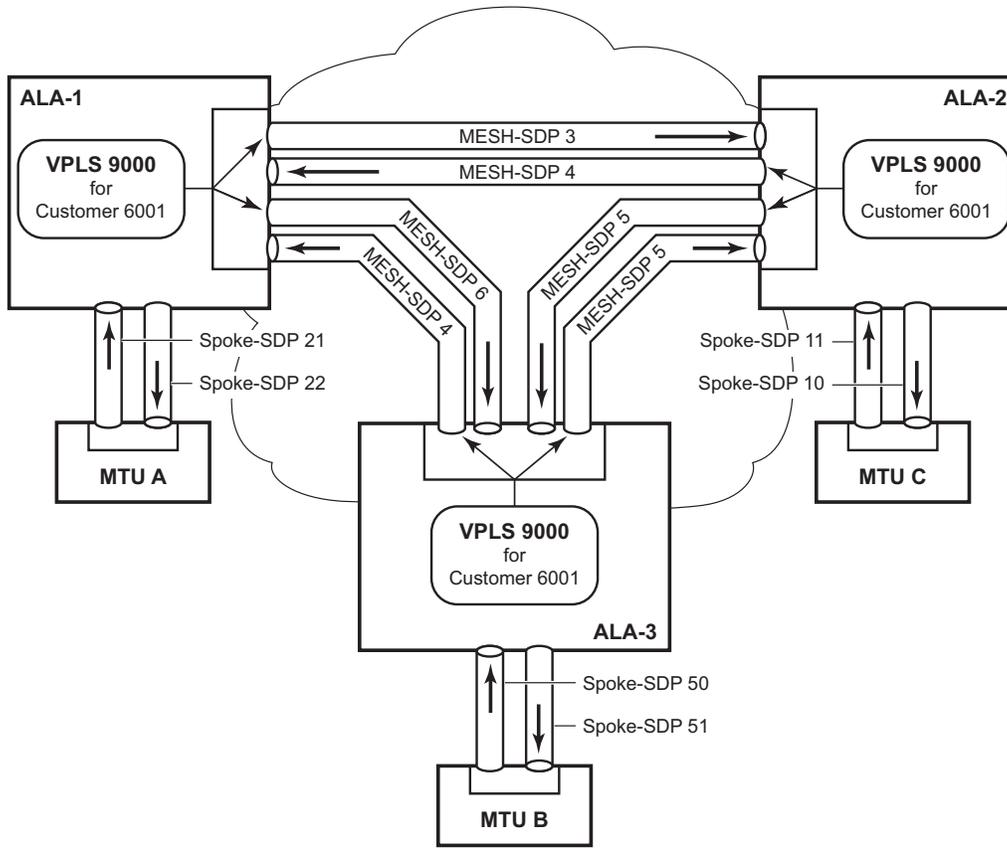
A VC-ID can be specified with the SDP-ID. The VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer SRs on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

[Figure 79](#) displays an example of a distributed VPLS service configuration of spoke and mesh SDPs (uni-directional tunnels) between routers and MTUs.

Configuring Overrides on Service SAPs

The following output displays 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#
```



OSSG032

Figure 79: 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 *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-hori-
                zon-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 displays 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."
```

Configuring VPLS Components

```
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#
```

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 on page 515](#)
 - [Spoke SDP Virtual Port Number on page 516](#)
 - [Spoke SDP Priority on page 516](#)
 - [Spoke SDP Path Cost on page 517](#)
 - [Spoke SDP Edge Port on page 517](#)
 - [Spoke SDP Auto Edge on page 518](#)
 - [Spoke SDP Link Type on page 518](#)
-

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 towards 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 its 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 — 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 on page 516](#) 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 thus 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 on page 496](#)). 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 on page 517](#)).

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

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 on page 519](#)
 - [Operationally Discarding on page 519](#)
 - [Operationally Learning on page 519](#)
 - [Operationally Forwarding on page 520](#)
-

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 proper STP state for each spoke SDP while not forwarding actual user traffic. The duration of the discarding state is explained in section [Forward Delay on page 496](#).

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 FIB. Layer 2 frames received on other forwarding interfaces and destined for the spoke SDP are also forwarded.

Spoke SDP BPDUs Encapsulation States

IEEE 802.1D (referred as dot1d) and Cisco's per VLAN Spanning Tree (PVST) BPDUs 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 BPDUs.

Table 12 shows differences between dot1D and PVST Ethernet BPDUs encapsulations based on the interface encap-type field:

Table 12: Spoke SDP BPDUs 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)
SNAP PID	N/A	N/A	N/A	01 0b
CONFIG or TCN BPDUs	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d
TLV: Type & 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.

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 displays 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
  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#
```

Configuring VPLS Redundancy

This section discusses the following service management tasks:

- [Creating a Management VPLS for SAP Protection on page 523](#)
 - [Creating a Management VPLS for Spoke SDP Protection on page 524](#)
-

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](#) . The tasks below should be performed on both nodes providing the protected VPLS service.

Before configuring a management VPLS, first read [VPLS Redundancy on page 402](#) 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 towards the MTU. Note that 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.

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 80](#). 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 on page 500](#) 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 on page 527](#)).
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 7 are 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.

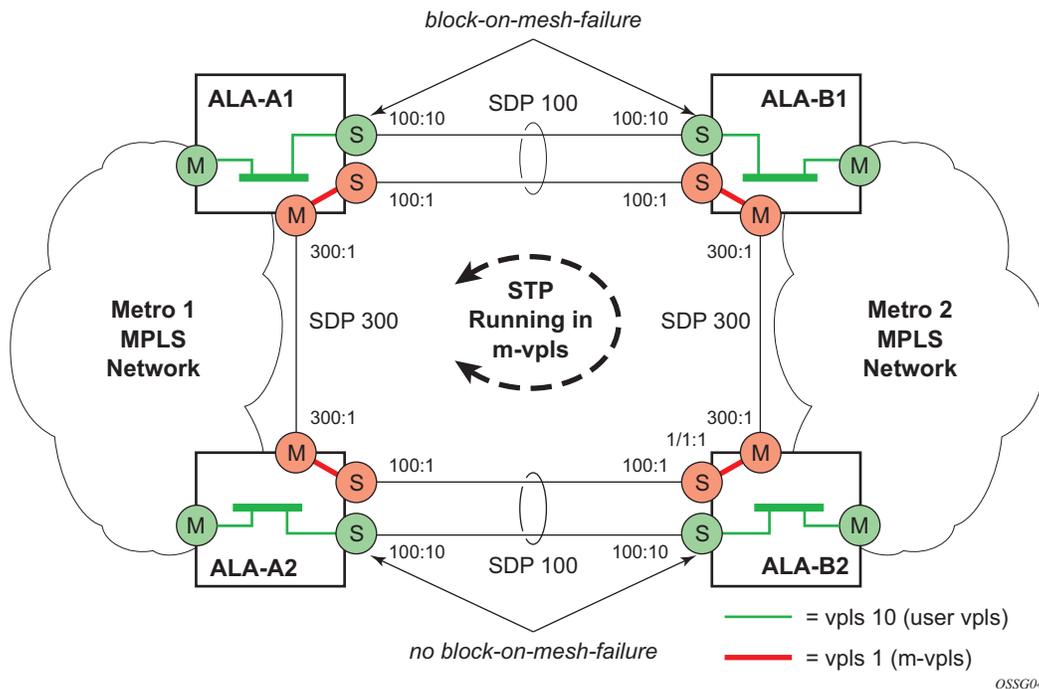


Figure 80: Example Configuration for Protected VPLS Spoke SDP

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`

Configuring VPLS Redundancy

The following example displays 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#
```

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.

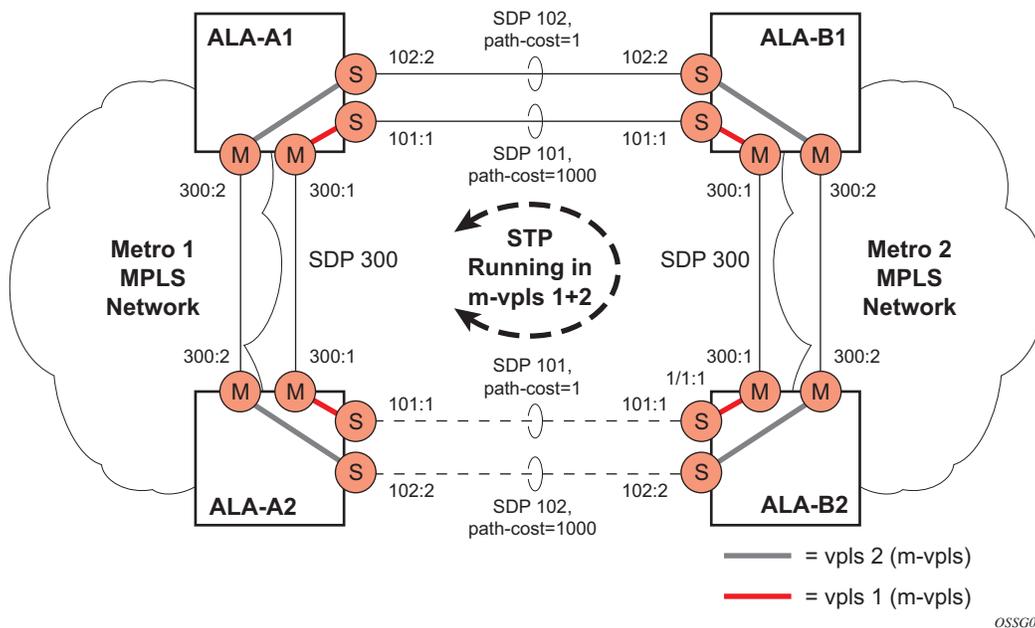


Figure 81: Example Configuration for Load Balancing Across Two Protected VPLS Spoke SDPs

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`

Configuring VPLS Redundancy

Note: the STP path costs in each peer node should be reversed.

The following example displays 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 displays 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#
```

Configuring VPLS Redundancy

The following example displays 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
    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
    mesh-sdp 300:2 create
    exit
    stp
    exit
    no shutdown
exit
...
-----
*A:ALA-A1>config>service#
```

The following example displays 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#
```

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`

Configuring Multi-Chassis Endpoints

The following output displays 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 on page 405](#) are expressed in this output.

Node Mapping to figures the document:

- 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
-----
```

Configuring VPLS Redundancy

```
        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

```

Configuring VPLS Redundancy

```
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#
```


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. Note that 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.

Note that 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
-----
```

Service Management Tasks

This section discusses the following service management tasks:

- [Modifying VPLS Service Parameters on page 539](#)
- [Modifying Management VPLS Parameters on page 540](#)
- [Deleting a Management VPLS on page 540](#)
- [Disabling a Management VPLS on page 541](#)
- [Deleting a VPLS Service on page 542](#)

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, and then enter the new information.

The following displays 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
stp
shutdown
exit
sap 1/1/5:22 create
description "VPLS SAP"
exit
exit
no shutdown
-----
*A:ALA-1>config>service>vpls#
```

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`

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`

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 desired, 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
```

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
```

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
```

Re-Enabling a VPLS Service

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`

VPLS Services Command Reference

Command Hierarchies

- [Global Commands on page 546](#)
- [Oper Group Commands on page 552](#)
- [SAP Commands on page 553](#)
- [Mesh SDP Commands on page 564](#)
- [Spoke SDP Commands on page 567](#)
- [Provider Tunnel Commands on page 571](#)
- [Show Commands on page 574](#)
- [Clear Commands on page 577](#)
- [Debug Commands on page 578](#)

VPLS Service Configuration Commands

Global Commands

```

config
  — service
    — vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls] [etree] [create]
    — no vpls service-id
      — allow-ip-int-bind
      — [no] allow-ip-int-bind
        — forward-ipv4-multicast-to-ip-int
        — [no] forward-ipv4-multicast-to-ip-int
      — backbone-smac ieee-address
      — no backbone-smac
      — 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...(up to 5 max)}]
        — no pw-template-binding policy-id
          — [no] bfd-enable
          — bfd-template [256 chars max]
          — 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]}}
        — 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*
- **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**]
 - **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*]
- **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**
 - **delayed-enable** *seconds*
 - **no delayed-enable**
 - **description** *description-string*
 - **no description**
 - **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**
 - **[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-protect**
 - **[no] mac** *ieee-address*
- **mac-subnet-length** *subnet-length*
- **no mac-subnet-length**
- **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** *ipv-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
  — oper-group seconds
  — no oper-group
  — mode mode
  — [no] shutdown
— [no] propagate-mac-flush
— [no] propagate-mac-flush-from-bvpls
— remote-age aging-timer
— no remote-age
— send-bvpls-flush {[all-but-mine] [all-from-me]}
— no send-bvpls-flush
— [no] send-flush-on-bvpls-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..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-pacing-interval** *milli-seconds*
- **no lsp-pacing-interval**
- **retransmit-interval** *seconds*
- **no retransmit-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** **alarm-only**
 - **restrict-protected-src** **[discard-frame]**
 - **no restrict-protected-src**
- **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*
 - **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**
- **vpls-group** *id*
 - **service-range** *startid-endid [vlan-id startvid]*
 - **vpls-template-binding** *name/id*
 - **vpls-sap-template-binding** *name/id*
 - **[no] mvrp-control**
- **vxlan vni** *vni-id create*
- **no vxlan vni**

Oper Group Commands

```
config
  — service
    — vpls service-id (See the Layer 2 Services Guide)
      — [no] interface ip-int-name
        — monitor-oper-group name
        — no monitor-oper-group
```

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] [root-leaf-tag leaf-tag-vid | leaf-ac]
— no sap sap-id
— accounting-policy acct-policy-id
— no accounting-policy
— [no] auto-learn-mac-protect
— 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
— bpdu-translation {auto | pvst | stp}
— no bpdu-translation
— [no] cflowd
— [no] collect-stats
— description description-string
— no description
— [no] disable-aging
— [no] disable-learning
— [no] discard-unknown
— egress
— [no] agg-rate
— [no] limit-unused-bandwidth
— [no] queue-frame-based-accounting
— rate {max | rate}
— no rate
— 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**
- **multicast-group** *group-name*
- **no multicast-group**
- **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**]
 - **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**
 - **cbs** *size-in-kbytes*
 - **no cbs**
 - **high-prio-only** *percent*
 - **no high-prio-only**
 - **mbs** *size-in-kbytes*
 - **no mbs**
 - **rate** *pir-rate* [**cir** *cir-rate*]
 - **no rate**
 - **wred-queue-policy** *slope-policy-name*
 - **no wred-queue-policy**
- [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**
 - **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**
- **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]
- **vmep-extensions**
- **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] feature**
- **[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..32000]
 - **no max-num-grp-sources**
 - **mcac**
 - **mc-constraints**
 - **level** *level-id* **bw** *bandwidth*
 - **no level** *level-id*
 - **number-down** *number-lag-port-down* **level** *level-id*
 - **no number-down**
 - **policy** *policy-name*
 - **no policy**
 - **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*
 - **no unconstrained-bw**
 - **no use-lag-port-weight**
 - **[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*]
- **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**
 - **high-prio-only** *percent*
 - **no high-prio-only**
 - **mbs** *size-in-kbytes*
 - **no mbs**
 - **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**
- **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..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**
- **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] shutdown**
- **monitor-oper-group** *group-name*
- **no monitor-oper-group**
- **oper-group** *group-name*
- **no oper-group**
- **multi-service-site** *customer-site-name*
- **no multi-service-site**
- **pim-snooping**
 - **max-num-groups** *num-groups*
 - **[no] monitor-oper-group** *name*
 - **[no] oper-group** *name*
- **[no] process-cpm-traffic-on-sap-down**
- **restrict-protected-src** **alarm-only**
- **restrict-protected-src** **[discard-frame]**
- **no restrict-protected-src**
- **restrict-unprotected-dst**
- **no restrict-unprotected-dst**
- **[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** [pt-pt | shared]
- **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**
- **tod-suite** *tod-suite-name*
- **no tod-suite**

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 | pvst | stp}
- **no bpdu-translation**
- [no] **collect-stats**
- **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**
- **egress**
 - **agg-rate-limit** *agg-rate* [queue-frame-based-accounting]
 - **no agg-rate-limit**
 - [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*
 - **no qos**
 - **scheduler-policy** *scheduler-policy-name*
 - **no scheduler-policy**
- **ingress**
 - **agg-rate-limit** *agg-rate*
 - **no agg-rate-limit**
 - **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**

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
        — 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]
        — eth-cfm
          — [no] collect-lmm-stats
          — mep mep-id domain md-index association ma-index [direction {up | down}]
          — 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**
 - **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 *vlan-id*]
- [no] **snelch-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..32000]
 - **no max-num-grp-sources**
 - **mcac**
 - **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** *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*]
 - **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
— 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 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
— 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..4094
— no vlan-vc-tag [0..4094]

```

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 | pvst | stp}
        — no bpdu-translation
        — [no] collect-stats
        —
        — description description-string
        — no description
        — [no] disable-aging
        — [no] disable-learning
        — [no] discard-unknown-source
        — eth-cfm
          — [no] collect-lmm-stats
          — mep mep-id domain md-index association ma-index [direction {up | down}]
          — 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
          — noccm-padding-size ccm-padding
          — [no] csf-enable
            — multiplier multiplier-value
            — no multiplier
          — [no] description
          — [no] eth-test-enable
            — test-pattern {all-zeros | a ll-ones} [crc-enable]
            — no test-pattern
          — fault-propagation-enable {use-if-tlv | suspend-ccm}

```

- **no fault-propagation-enable**
- **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 *vlan-id*]
- [no] **snelch-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..32000]
 - **no max-num-grp-sources**
 - **mcac**
 - **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*]
 - **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**
 - **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] monitor-oper-group** *name*
 - **[no] oper-group** *name*
- **precedence** *precedence-value* | **primary**
- **no precedence**
- **[no] pw-status-signaling**
- **propagate-mac-flush** [*precedence-value* | **primary**]
- **no propagate-mac-flush**
- **[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** [**pt-pt** | **shared**]
 - **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..4094*
- **no vlan-vc-tag** [*0..4094*]

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

```

Egress Multicast Group Commands

```

config
  — service
    — egress-multicast-group group-name [create]
    — no egress-multicast-group group-name
      — description description-string
      — no description
      — dest-chain-limit destinations per pass
      — no dest-chain-limit
      — sap-common-requirements
        — dot1q-etype 0x0600..0xffff
        — no dot1q-etype
        — egress-filter [ip ip-filter-id]
        — egress-filter [ipv6 ipv6-filter-id]
        — egress-filter [mac mac-filter-id]
        — no egress-filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
        — encap-type {dot1q | null}
        — no encap-type
        — qinq-etype [0x0600..0xffff]
        — no qinq-etype
        — qinq-fixed-tag-value tag-value
        — no qinq-fixed-tag-value

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]
      — no sap sap-id
        — egress
          — multicast-group group-name
          — no multicast-group

```

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
```

Show Commands

```

show
  — service
    — egress-label egress-label1 [egress-label2]
    — fdb-info
    — fdb-mac ieee-address [expiry]
    — id service-id
      — all
      — base
      — bgp-evpn
      — etree
      — fdb [sap sap-id] [expiry] | [sdp sdp-id] [expiry] | [mac ieee-address] [expiry] |
        endpoint endpoint | [detail] [expiry] [pbb]
      — igmp-snooping
        — all
        — base
        — mroute mroute [detail]
        —
        — proxy-db [detail]
        — proxy-db [group grp-ip-address]
        — querier
        — static [sap sap-id | sdp sdp-id:vc-id]
      — isid-policy
      — labels
      — l2pt disabled
      — l2pt [detail]
      — mac-move
      — mac-protect
      — mfib [brief | statistics] [ip | mac] brief
      — mfib [group grp-address | *] [statistics]
      — proxy-arp [ip-address ip-address] [detail]
      — proxy-nd [ip-address ip-address] [detail]
      — sap [sap-id] [detail]
      — sdp [sdp-id | far-end ip-addr] [detail]
      — sdp sdp-id:vc-id {mrp | mrrp}
      — site [detail]
      — site name
      — split-horizon-group [group-name]
      — stp mst-instance mst-inst-number
      — stp [detail]
      — 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 *filter-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
- **vxlan**

Clear Commands

```

clear
  — service
    — id service-id
      — 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]]
        — querier
      — mesh-sdp sdp-id[:vc-id] ingress-vc-label
      — 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
        — counters
        — l2pt
        — mesh-sdp sdp-id[:vc-id] {all | counters | stp | mrp}
        — spoke-sdp sdp-id[:vc-id] {all | counters | stp | l2pt | mrp}
        — stp
      — sap sap-id {all | cem | counters | l2pt | stp | mrp}
      — sdp sap-id {keep-alive}

```

Debug Commands

```

debug
  — service
    — id service-id
      — [no] arp-host
      — igmp-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] 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}
      — 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
  
```

Tools Commands

```

tools
  — dump
    — service
      — proxy-arp usage
      — proxy-nd usage
      — vpls id
        — provider-tunnels type
  
```

- **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 7950 OS OAM and Diagnostic Guide for information about CLI commands and syntax for OAM and diagnostics commands.

VPLS Service Configuration Commands

Generic Commands

shutdown

Syntax	<code>[no] shutdown</code>
Context	<pre> config>service>vpls config>service>vpls>spb>level config>service>vpls>snooping config>service>vpls>igmp-snooping config>service>vpls>mac-move config>service>vpls>gsmp config>service>vpls>gsmp>group config>service>vpls>gsmp>group>neighbor config>service>vpls>interface config>service>vpls>split-horizon-group config>service>vpls>sap config>service>vpls>sap>stp config>service>vpls>sap>arp-host config>service>vpls>sap>sub-sla-mgmt config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp config>service>vpls>spoke-sdp>stp config>service>vpls>stp config>service>vpls>spoke-sdp>stp config>service>vpls>mrp 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>bgp-ad config>service>vpls>eth-cfm>mep config>service>vpls>mesh-sdp>eth-cfm>mep config>service>vpls>>spoke-sdp>eth-cfm>mep </pre>
Description	<p>This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.</p> <p>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.</p> <p>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.</p> <p>The no form of this command places the entity into an administratively enabled state.g</p>

- 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 two SAPs or if one SDP 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.
 - 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 once bound to the VPLS.
 - Routed VPLS with forward-ipv4-multicast-to-ip-int and IGMP Snooping** — In order 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.

description

- Syntax** `description description-string`
no description
- Context** `config>service>vpls`
`config>service>vpls>gsmp>group`
`config>service>vpls>gsmp>group>neighbor`
`config>service>vpls>igmp-snooping>mvr`
`config>service>vpls>interface`
`config>service>vpls>split-horizon-group`
`config>service>vpls>sap`
`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.
- Default** No description associated with the configuration context.

Parameters *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.

VPLS Service Commands

vpls

Syntax	vpls <i>service-id</i> customer <i>customer-id</i> vpn <i>vpn-id</i> [m-vpls] [bvpls i-vpls] [etree] [create] no vpls <i>service-id</i>
Context	config>service
Description	<p>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 <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> exists, the context for editing the service is entered.</p> <p>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.</p> <p>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 <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> 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.</p> <p>Once a service is created, the use of the customer <i>customer-id</i> is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect <i>customer-id</i> specified will result in an error.</p> <p>More than one VPLS service may be created for a single customer ID.</p> <p>By default, no VPLS instances exist until they are explicitly created.</p> <p>The no form of this command deletes the VPLS service instance with the specified <i>service-id</i>. 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.</p>
Parameters	<p><i>service-id</i> — 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 <i>service-id</i> must be the same number used for every router on which this service is defined.</p> <p>Values <i>service-id:</i> 1 — 2147483648 <i>svc-name:</i> 64 characters maximum</p> <p>customer <i>customer-id</i> — 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.</p> <p>Values 1 — 2147483647</p> <p>vpn <i>vpn-id</i> — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.</p> <p>Values 1 — 2147483647</p> <p>Default null (0)</p>

m-vpls — Specifies a management VPLS.

e-tree — Specifies a VPLS service as an E-TreeVPLS. 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.

Note that a 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.

backbone-smac

Syntax	backbone-smac <i>ieee-address</i>
Context	config>service>vpls
Description	This command configures the backbone source MAC address used for PBB. This command allows a per B-VPLS control of the B-SMAC and the B-Mcast MAC. All I-VPLS provisioned under this B-VPLS will share the provisioned value.
Default	backbone-smac address is chassis MAC address
Parameters	<i>ieee-address</i> — Specifies the backbone source MAC address.

backbone-vpls

Syntax	backbone-vpls <i>vpls-id[:isid]</i> no backbone-vpls
Context	config>service>vpls
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.

VPLS Service Commands

- 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 — 16777215

stp

- Syntax** [no] stp
- Context** config>service>vpls>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 signalled 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 | pvst | stp}
no bpdu-translation
- Context** config>service>vpls>spoke-sdp
config>service>vpls>sap
- Description** This command enables the translation of BPDUs to a given format, meaning that all BPDUs transmitted on a given SAP or spoke SDP will have a specified format.
The **no** form of this command reverts to the default setting.
- Default** no bpdu-translation
- Parameters** **auto** — Specifies that appropriate format will be detected automatically, based on type of bpdus received on such port.
pvst — Specifies the BPDU-format as PVST. Note that the correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).

stp — Specifies the BPDU-format as STP.

cflowd

Syntax	[no] cflowd
Context	config>service>vpls>sap
Description	<p>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 I2-ip template enabled.</p> <p>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.</p> <p>For L2 services, only ingress sampling is supported.</p>
Default	no cflowd

lag-link-map-profile

Syntax	lag-link-map-profile <i>link-map-profile-id</i> no lag-link-map-profile
Context	config>service>vpls>sap
Description	<p>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/network interface egress traffic will be re-hashed over LAG as required by the new configuration.</p> <p>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.</p>
Default	no lag-link-map-profile
Parameters	<i>link-map-profile-id</i> — An integer from 1 to 64 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

I2pt-termination

Syntax	I2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp] no I2pt-termination
Context	config>service>vpls>spoke-sdp config>service>vpls>sap

VPLS Service Commands

Description	This command enables Layer 2 Protocol Tunneling (L2PT) termination on a given SAP or spoke SDP. L2PT termination will be supported only for STP BPDUs. PDUs of other protocols will be discarded. This feature can be enabled only if STP is disabled in the context of the given VPLS service.
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). uddl — Specifies unidirectional link detection. vtp — Specifies the virtual trunk protocol.

def-mesh-vc-id

Syntax	[no] def-mesh-vc-id <i>vc-id</i>
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	<i>vc-id</i> — Specifies the default mesh vc-id. Values 1 — 4294967295

default-gtw

Syntax	default-gtw
Context	config>service>vpls
Description	This command configures a service default gateway.

ip

Syntax	ip <i>ip-address</i> no ip
Context	config>service>vpls>defgw
Description	This command configures the default gateway IP address.

mac

Syntax	mac <i>ieee-address</i>
Context	config>service>vpls>defgw
Description	This command configures the default gateway MAC address.

disable-aging

Syntax	[no] disable-aging
Context	config>service>vpls config>service>vpls>spoke-sdp config>service>vpls>sap config>template>vpls-template
Description	<p>This command disables MAC address aging across a VPLS service or on a VPLS service SAP or spoke SDP.</p> <p>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.</p> <p>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.</p> <p>When the disable-aging command is entered at the VPLS level, the disable-aging state of individual SAPs or SDPs will be ignored.</p> <p>The no form of this command enables aging on the VPLS service.</p>
Default	no disable-aging

disable-learning

Syntax	[no] disable-learning
Context	config>service>vpls config>service>vpls>sap config>service>vpls>spoke-sdp config>template>vpls-template

VPLS Service Commands

- 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>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 FIB 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.

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
- 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 *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.

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, refer to [Auto-Learn MAC Protect on page 377](#).

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 will ignore standby-bit received from TLDP peers for the given spoke SDP and performs internal tasks without taking it into account.

This command is present at endpoint level as well as spoke SDP level. If the spoke SDP is part of the explicit-endpoint, it is not possible to change this setting 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 given explicit-endpoint will inherit this setting from the endpoint configuration.

Default enabled

restrict-protected-src

VPLS Service Commands

Syntax	restrict-protected-src restrict-protected-src [discard-frame] no restrict-protected-src
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 This command indicates the action to take whenever a relearn request for a protected MAC is received on a restricted SAP belonging to this SHG When enabled, the agent will protect the MAC from being learned or re-learned on a SAP that has restricted learning enabled.
Default	restrict-protected-src
Parameters	discard-frame — Specifies that the SAP will start discarding the frame in addition to generating sapReceivedProtSrcMac notification.

revert-time

Syntax	revert-time <i>revert-time</i> 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	<i>revert-time</i> — 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 — 600 <i>infinite</i> — Specifying this keyword makes endpoint non-revertive.

static-mac

Syntax	static-mac <i>ieee-address</i> [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	<i>ieee-address</i> — 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 given 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 given 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 <i>high-water-mark</i>
Context	config>service>vpls config>template>vpls-template
Description	This command specifies the value to send logs and traps when the threshold is reached.
Parameters	<i>high-water-mark</i> — Specify the value to send logs and traps when the threshold is reached.
Values	0— 100
Default	95%

fdb-table-low-wmark

Syntax	[no] fdb-table-low-wmark <i>low-water-mark</i>
Context	config>service>vpls config>template>vpls-template
Description	This command specifies the value to send logs and traps when the threshold is reached.
Parameters	<i>low-water-mark</i> — Specify the value to send logs and traps when the threshold is reached.
Values	0— 100
Default	90%

fdb-table-size

Syntax	fdb-table-size <i>table-size</i> no fdb-table-size [<i>table-size</i>]
Context	config>service>vpls config>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	250 — Forwarding table of 250 MAC entries.
Parameters	<i>table-size</i> — Specifies the maximum number of MAC entries in the FDB.
Values	1 — 511999 Chassis-mode A or B limit: 131071 Chassis-mode C limit: 196607 Chassis-mode D limit: 511999

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>service>vpls
Description	This command creates an IP interface.

address

Syntax	address <i>ip-address</i> [/ <i>mask</i>]> [<i>netmask</i>]
---------------	---

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.

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 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

arp-timeout

Syntax	arp-timeout <i>seconds</i> 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. 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	<i>seconds</i> — 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 — 65535

delayed-enable

Syntax	delayed-enable <i>seconds</i> no delayed-enable
Context	config>service>vpls>interface
Description	This command will cause a delay in the activation of an IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up. 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 expires.
Parameters	<i>seconds</i> — Specifies a delay, in seconds, to make the interface operational. Values 1 — 1200

mac

Syntax	mac <i>ieee-address</i> 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

VPLS Service Commands

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 <i>ieee-address</i> [create] sap <i>sap-id</i> monitor <i>fwd-status</i> mac <i>ieee-address</i> [create] spoke-sdp <i>sdp-id:vc-id</i> monitor <i>fwd-status</i> no mac <i>ieee-address</i>
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. 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). It cannot be all zeros. create — This keyword is mandatory while creating a static MAC. monitor fwd-status — Specifies that this static mac is based on the forwarding status of the SAP or spoke SDP for multi-homed operation.

unnumbered

Syntax	unnumbered [<i>ip-int-name</i> <i>ip-address</i>] no unnumbered
Context	config>service>ies>if config>service>vpls>if config>service>vprn>if
Description	This command configures the interface as an unnumbered interface. Unnumbered IP interface is supported on a SONET/SDH access port with the PPP, ATM, Frame Relay, cisco-HDLC encapsulation. It is not supported on access ports that do not carry IP traffic, but are used for native TDM circuit emulation.
Parameters	<i>ip-int-name</i> — Specifies the name of the IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. <i>ip-address</i> — 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	<p>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.</p> <p>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.</p>
Default	no default

entry

	entry <i>id</i> create no entry
Context	config>service>vpls>isid-policy
Description	<p>This command creates or edits an isid-policy entry. Multiple entries can be created using unique entry-id numbers within the isid-policy.</p> <p>Default: No entry</p> <p>entry-id — 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.</p> <p>The following rules govern the usage of multiple entry statements:</p> <ul style="list-style-type: none"> • overlapping values are allowed: <ul style="list-style-type: none"> – 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 to execute the associated action for that entry. <p>no isid - removes all the previous statements under one entry.</p> <p>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.</p> <p>Values 1-65535</p>

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	<i>isid</i> — Specifies the ISID value in 24 bits. When singular, ISID identifies a particular ISID to be used for matching. Values 0..16777215 <i>to isid</i> — 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>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	<i>aging-timer</i> — The aging time for local MACs expressed in seconds. Values 60 — 86400

mac-move

Syntax	[no] mac-move
Context	config>service>vpls config>template>vpls-template
Description	<p>This command enables 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.</p> <p>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 given 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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of this command disables MAC move.</p>

mac-protect

Syntax	mac-protect
Context	config>service>vpls
Description	<p>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.</p>
Default	disabled

mac

Syntax	[no] mac <i>ieee-address</i>
Context	config>service>vpls>mac-protect
Description	This command adds a protected MAC address entry.
Parameters	<i>ieee-address</i> — 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 <i>subnet-length</i> 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 FIB entries. Then, it will install the first 28 bits in the FIB 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 FIB lookup to determine the next hop. The no form of this command switches back to full MAC lookup.
Parameters	<i>subnet-length</i> — Specifies the number of bits to be considered when performing MAC learning or MAC switching. Values 24 — 48

mac-notification

Syntax	mac-notification
Context	config>service>vpls bvpls
Description	This command controls the settings for the MAC notification message. The mac-notification message must be generated under the following events: <ol style="list-style-type: none"> 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/Epipes) 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 <i>value</i>
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	<i>value</i> — Specifies the frequency of subsequent MAC notification messages.
	Values 100 ms – 10 sec, in increments of 100 ms up to 1 sec and then in increments of 1 second up to 10 sec.

renotify

Syntax	renotify <i>value</i> 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	<i>value</i> — Specifies the time interval between re-notification in seconds.
	Values 240—840 seconds

count

Syntax	[no] count <i>value</i>
Context	config>service>vpls>mac-notification
Description	This command configures how often MAC notification messages are sent.
Parameters	<i>value</i> — Specifies, in seconds, how often MAC notification messages are sent.

VPLS Service Commands

Values 1—10
Default Inherits the chassis level configuration from config>service>mac-notification

move-frequency

Syntax **move-frequency** *frequency*
no move-frequency

Context config>service>vpls>mac-move
config>template>vpls-template>mac-move

Description This command indicates the maximum rate at which MAC's 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 MAC's.
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 — 100

number-retries

Syntax **number-retries** *number-retries*
no number-retries

Context config>service>vpls>mac-move
config>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 — 255

primary-ports

Syntax **primary-ports**

Context config>service>vpls>mac-move
config>template>vpls-template>mac-move

Description This command enables 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 <i>cumulative-factor</i> no cumulative-factor
Context	configure->service->vpls->mac-move->primary-ports configure->service->vpls->mac-move->secondary-ports config>template>vpls-template>mac-move>primary-ports config>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	<i>cumulative-factor</i> — Specifies a MAC relearn period to be used for MAC relearn rate.
Values	3 — 10

sap

Syntax	sap [split-horizon-group <i>group-name</i>] [create] no sap <i>sap-id</i>
Context	config>service>vpls>mac-move>primary-ports config>service>vpls>mac-move>secondary-ports
Description	This command declares a given SAP as a primary (or secondary) VPLS port.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1271 for command syntax.

spoke-sdp

Syntax	[no] spoke-sdp <i>spoke-id</i>
Context	config>service>vpls>mac-move>primary-ports config>service>vpls>mac-move>secondary-ports
Description	This command declares a given spoke SDP as a primary (or secondary) VPLS port.
Parameters	<i>spoke-id</i> — Specifies the SDP ID to configure as the primary VPLS port.
Values	1 — 17407
	<i>vc-id</i> — The virtual circuit identifier.
Values	1 — 4294967295

cumulative-factor

Syntax	[no] cumulative-factor <i>factor</i>
Context	config>service>vpls>mac-move>primary-ports config>service>vpls>mac-move>secondary-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	<i>factor</i> — Specifies the factor defining the number of mac-relearn measurement periods can be used to measure mac-relearn rate.
	Values 2 — 10

secondary-ports

Syntax	secondary-ports
Context	config>service>vpls>mac-move config>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 <i>timeout</i> no retry-timeout
Context	config>service>vpls>mac-move config>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	10 (when mac-move is enabled)
Parameters	<i>timeout</i> — Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.
Values	0 — 120

mfib-table-high-wmark

Syntax	[no] mfib-table-high-wmark <i>high-water-mark</i>
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	<i>high-water-mark</i> — Specifies the multicast FIB high watermark as a percentage.
Values	1 — 100
Default	95%

mfib-table-low-wmark

Syntax	[no] mfib-table-low-wmark <i>low-water-mark</i>
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	<i>low-water-mark</i> — Specifies the multicast FIB low watermark as a percentage.
Values	1 — 100
Default	90%

mfib-table-size

Syntax	mfib-table-size <i>size</i> 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 <i>mfib-table-size</i> 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

VPLS Service Commands

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* — The maximum number of (s,g) entries allowed in the Multicast FIB.

Values 1 — 16383

mld-snooping

Syntax **mld-snooping**

Context config>service>vpls
config>service>vpls>sap

Description This command configures MLD snooping parameters.

remote-age

Syntax **remote-age seconds**
no remote-age

Context config>service>vpls
config>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* — The aging time for remote MACs expressed in seconds.

Values 60 — 86400

send-bvpls-flush

Syntax	send-bvpls-flush {[all-but-mine] [all-from-me]} no send-bvpls-flush
Context	config>service>vpls
Description	<p>This command enables generation of LDP MAC withdrawl “flush-all-from-me” in the B-VPLS domain when the following triggers occur in the related IVPLS:</p> <ul style="list-style-type: none"> • MC-LAG failure • Failure of a local SAP • Failure of a local pseudowire/SDP binding <p>Note that failure means transition of link SAP/pseudowire to either down or standby status.</p> <p>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.</p>
Default	no send-bvpls-flush
Parameters	<p>all-but-mine — Specifies to send an LDP flush all-but-mine and also sent into the B-VPLS. Note that both parameters can be set together.</p> <p>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. Note that both parameters can be set together.</p>

send-flush-on-bvpls-failure

Syntax	[no] send-flush-on-bvpls-failure
Context	config>service>vpls ivpls
Description	<p>This command enables the generation in the local I-VPLS of a LDP MAC flush-all-from-me following a failure of SAP/the whole endpoint/spoke-SDP in the related B-VPLS. Note that the failure of mesh-SDP in B-VPLS does not generate the I-VPLS MAC flush.</p> <p>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.</p>
Default	no send-flush-on-bvpls-failure

propagate-mac-flush-from-bvpls

Syntax	[no] propagate-mac-flush-from-bvpls
Context	config>service>vpls ivpls
Description	<p>This command enables the propagation in the local I-VPLS 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 I-VPLS FIB. 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.</p>

VPLS Service Commands

The **no** form of this command disables the propagation of LDP MAC Flush in I-VPLS from the related B-VPLS.

Default no propagate-mac-flush-from-bvpls

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 “oper-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>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

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 — 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 given 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 and 7950 XRS 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 given service once it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

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. In Release 8.0 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, and it cannot be an I-VPLS service bound to a B-VPLS context
- 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	forward-ipv4-multicast-to-ip-int no forward-ipv4-multicast-to-ip-int
Context	config>service>vpls>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

site

Syntax	site <i>name</i> [create] no site <i>name</i>
Context	config>service>vpls
Description	This command configures a VPLS site. The no form of the command removes the name from the configuration.
Parameters	<i>name</i> — 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 <i>seconds</i> 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	<i>seconds</i> — Specifies the site boot-timer in seconds. Values 0 — 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 <i>group-name</i> 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 <i>name</i> 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 <i>sap-id</i> 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. See [Common CLI Command Descriptions on page 1271](#) for command syntax.

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 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 — 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 **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.

VPLS Service Commands

minated 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 **configure>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 — 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 — 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.

spb

Syntax	spb [create] no spb
Context	config>service>vpls
Description	This command configures Shortest Path Bridging.

level

Syntax	level [1..1]
Context	config>service>vpls>spb
Description	This command enables the context to configure SPB level information.

bridge-priority

Syntax	bridge-priority <i>bridge-priority</i> 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	<i>bridge-priority</i> — Specifies the bridge priority. Values 0 — 15

ect-algorithm

Syntax	ect-algorithm <i>fid-range</i> <i>fid-range</i> { low-path-id high-path-id } no ect-algorithm <i>fid-range</i> <i>fid-range</i>
Context	config>service>vpls>spb>level
Description	This command configures the ECT algorithm of forwarding range.
Parameters	fid-range <i>fid-range</i> — Specifies the FID range. Values 1..4095 — 1..4095 low-path-id — Keyword to specify the low path ID. high-path-id — Keyword to specify 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

hello-interval

Syntax	hello-interval <i>seconds</i> no hello-interval
Context	config>service>vpls>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 intersystem. 9 — Hello interval default for non-designated intersystems.
Parameters	<i>seconds</i> — The hello interval in seconds expressed as a decimal integer.
	Values 1 — 20000

hello-multiplier

Syntax	hello-multiplier <i>multiplier</i> no hello-multiplier
Context	config>service>vpls>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	<i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer.
	Values 2 — 100

Isp-lifetime

Syntax	Isp-lifetime <i>seconds</i> no Isp-lifetime
Context	config>service>vpls>spb
Description	<p>This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.</p> <p>Each LSP received is maintained in an LSP database until the Isp-lifetime expires unless the originating router refreshes the LSP. By default, each router refreshes its LSP's every 20 minutes (1200 seconds) so other routers will not age out the LSP.</p> <p>The LSP refresh timer is derived from this formula: $\text{Isp-lifetime}/2$</p> <p>The no form of the command reverts to the default value.</p>
Default	1200 — LSPs originated by the router should be valid for 1200 seconds (20 minutes).
Parameters	<p><i>seconds</i> — The time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.</p> <p>Values 350 — 65535</p>

Isp-refresh-interval

Syntax	Isp-refresh-interval <i>seconds</i> no Isp-refresh-interval
Context	config>service>vpls>spb
Description	<p>This command configures the LSP refresh timer interval. When configuring the LSP refresh interval, the value that is specified for Isp-lifetime must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.</p> <p>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 Isp-refresh-interval command will be rejected.</p>
Default	600 seconds
Parameters	<p><i>seconds</i> — Specifies the refresh interval.</p> <p>Values 150— 65535</p>

Isp-wait

Syntax	Isp-wait <i>Isp-wait</i> [<i>Isp-initial-wait</i> [<i>Isp-second-wait</i>]]
Context	config>service>vpls>spb
Description	This command is used to customize the throttling of LSP-generation. Timers that determine when to generate the first, second and subsequent LSPs can be controlled with this command. Subsequent

VPLS Service Commands

LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

Parameters *lsp-max-wait* — Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated.

Values 1 — 120

Default 5

lsp-initial-wait — Specifies the initial LSP generation delay in seconds.

Values 0 — 100

Default 0

lsp-second-wait — Specifies the hold time in seconds between the first and second LSP generation.

Values 1 — 100

Default 1

overload-on-boot

Syntax **overload-on-boot** [*timeoutseconds*]
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 and/or show router isis status commands to display the administrative and operational state as well as all timers.
Parameters	timeout <i>seconds</i> — Configure the timeout timer for overload-on-boot in seconds. Values 60 — 1800

overload

Syntax	overload [<i>timeout seconds</i>] 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 be 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	<i>seconds</i> — The time, in seconds, that this router must operate in overload state. Default infinity (overload state maintained indefinitely) Values 60 — 1800

metric

Syntax	metric <i>ipv4-metric</i> no metric
Context	config>service>vpls>spb>level
Description	This command configures the IS-IS interface metric for IPv4 unicast.
Parameters	<i>ipv4-metric</i> — Specifies the IS-IS interface metric for IPv4 unicast. Values 1 — 16777215

lsp-pacing-interval

Syntax	lsp-pacing-interval <i>milliseconds</i> no lsp-pacing-interval
Context	config>service>vpls>spb
Description	This command configures the interval between LSP packets are sent from the interface. To avoid bombarding adjacent neighbors with excessive data, pace the Link State Protocol Data Units (LSP's). If a value of zero is configured, no LSP's are sent from the interface. The no form of the command reverts to the default value.
Default	100 — LSPs are sent in 100 millisecond intervals.
Parameters	<i>milliseconds</i> — The interval in milliseconds that IS-IS LSP's can be sent from the interface expressed as a decimal integer. Values 0 — 65535

retransmit-interval

Syntax	retransmit-interval <i>seconds</i> no retransmit-interval
Context	config>service>vpls>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	100
Parameters	<i>seconds</i> — The interval in seconds that IS-IS LSPs can be sent on the interface. Values 1 — 65535

split-horizon-group

Syntax	[no] split-horizon-group [<i>group-name</i>] [<i>residential-group</i>]
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.

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)

Default A split horizon group is by default not created as a residential-group.

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** version 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

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	<p>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.</p> <p>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). In order 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.</p>
Default	no auto-learn-mac-protect

restrict-protected-src

Syntax	restrict-protected-src [<i>alarm-only</i> <i>discard-frame</i>] 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	<p>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:</p> <ul style="list-style-type: none"> • 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/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 FP2 per 10 minutes in a given 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). In order 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 both the “**restrict-protected-src [alarm-only]**” command and with the configuration of manually protected MAC addresses within a given VPLS.

Note that the **alarm-only** parameter is not supported on the or 7950 XRS

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/mesh-SDP.

Default no alarm-only

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 FP2 per MAC address per 10 minutes within a given VPLS service.

Default no discard-frame

Default no restrict-protected-src

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 to 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.

VPLS Service Commands

Default no restrict-unprotected-dst

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 — 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.
 - Note that 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—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-VPLSes to control same VLAN range on different ports.

Values 1—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—1024

vpls-sap-template-binding

Syntax	vpls-sap-template-binding <i>name/id</i> no vpls-sap-template-binding
Context	config>service>vpls>vpls-group
Description	<p>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.</p> <p>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 vpls-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.</p>
Default	no vpls-sap-template-binding
Parameters	<p><i>name</i> — Specifies the name of the VPLS template.</p> <p>Values ASCII character string</p> <p><i>id</i> — Specifies the ID of the VPLS template</p> <p>Values 1—8196</p>

mvrp-control

Syntax	[no] mvrp-control
Context	config>service>vpls>vpls-group
Description	<p>This command enables MVRP control in the VPLS instances instantiated using the templates for the specified vpls-group. That means the flooding FIB 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.</p> <p>The no form of this command disallows MVRP control over this VPLS. The VPLS will be created with a regular FIB and will become as a result active upon creation time. Command change is allowed only when the related vpls-group is in shutdown state.</p>
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 <i>value</i> no hold-time
Context	config>service>vpls>mrp>mvrp
Description	<p>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.</p> <p>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.</p> <p>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.</p>
Default	no hold-time
Parameters	<i>value</i> — Specifies the hold time in minutes
Values	1—30 minutes

endstation-vid-group

Syntax	endstation-vid-group <i>id</i> vlan-id <i>startvid-endvid</i> no endstation-vid-group <i>id</i>
Context	config>service>vpls>mrp>mvrp

VPLS Service Commands

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.

Note that VPLS-groups are not allowed under the associated Management VPLS (MVPLS) once the endstation is configured under one SAP. VPLS-groups can be supported in the chassis using a different MVPLS.

The **no** form of the command removes the specified group id.

Default no endstation-vid-group

Parameters *id* — Specifies the range index.

Values 1—4094

starvid-endvid — Specifies the range of VLANs to be controlled by MVRP.

Values 1—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

tod-suite

Syntax **tod-suite** *tod-suite-name*
no tod-suite

Context config>service>vpls>sap

Description This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the **config>system>cron** context.

Default no tod-suite

Parameters *tod-suite-name* — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

vxlan

Syntax	vxlan vni <i>vni-id</i> create no vxlan vni <i>vni-id</i>
Context	config>service>vpls
Description	This command enables the use of vxlan in the vpls service.
Default	none
Parameters	vni — Specifies the vxlan network identifier configured in the vpls service. All the 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 — 16777215

VPLS Interface Commands

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>service>vpls
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>By default, no default IP interface names are defined within the system. All VPLS IP interfaces must be explicitly defined in an enabled state.</p> <p>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.</p> <p>For VPLS services, the IP interface must be shutdown before the SAP on that interface is removed.</p> <p>For VPLS service, ping and traceroute are the only applications supported.</p>
Parameters	<p><i>ip-int-name</i> — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP.</p> <p>An interface name:</p> <ul style="list-style-type: none"> • 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. <p>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.</p>

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 ipaddress, 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. Note that 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

VPLS Interface Commands

128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

General Switch Management Protocol Commands

gsmp

Syntax	gsmp
Context	config>service>vpls
Description	This command enables the context to configure General Switch Management Protocol (GSMP) connections maintained in this service.
Default	not enabled

group

Syntax	[no] group <i>name</i>
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

General Switch Management Protocol Commands

config>service>vprn>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 — 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 — 25

neighbor

Syntax **[no] neighbor** *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 towards 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, ef, 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 — 7

persistency-database

Syntax	persistency-database no persistency-database
Context	config>service>vpls <service id>gsmp config>service>vprn<service id>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

VPLS STP Commands

stp

Syntax	stp
Context	config>service>vpls config>service>vpls>sap config>service>vpls>spoke-sdp config>template>vpls-template
Description	This command enables the context to configure the Spanning Tree Protocol (STP) parameters. Alcatel-Lucent'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 Alcatel-Lucent'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 on page 641). 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)

General Switch Management Protocol Commands

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>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 — 30

hello-time

Syntax **hello-time** *hello-time*
no hello-time

Context config>service>vpls>stp

```
config>template>vpls-template>stp
```

Description	<p>This command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.</p> <p>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.</p> <p>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).</p> <p>The configured hello-time can also be used to calculate the forward delay. See auto-edge on page 641.</p> <p>The no form of this command returns the hello time to the default value.</p>
Default	2 seconds
Parameters	<p><i>hello-time</i> — The hello time for the STP instance in seconds.</p> <p>Values 1 — 10</p>

hold-count

Syntax	<p>hold-count <i>BDPU tx hold count</i></p> <p>no hold-count</p>
Context	<pre>config>service>vpls>stp config>template>vpls-template>stp</pre>
Description	<p>This command configures the peak number of BPDUs that can be transmitted in a period of one second.</p> <p>The no form of this command returns the hold count to the default value</p>
Default	6
Parameters	<p><i>BDPU tx hold count</i> — The hold count for the STP instance in seconds.</p> <p>Values 1 — 10</p>

link-type

Syntax	<p>link-type {<i>pt-pt</i> <i>shared</i>}</p> <p>no link-type</p>
Context	<pre>config>service>vpls>sap>stp config>service>vpls>spoke-sdp>stp</pre>
Description	<p>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.</p> <p>The no form of this command returns the link type to the default value.</p>

General Switch Management Protocol Commands

Default pt-pt

mst-instance

Syntax **mst-instance** *mst-inst-number*

Context config>service>vpls>sap>stp

Description This command enables 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 — 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 command specifies path-cost within a given instance, expressing probability that a given 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 towards the spanning tree regional root which include this port.

Values 1 — 200000000

mst-priority

Syntax **mst-priority** *stp-priority*
no mst-priority

Context config>service>vpls>sap>stp>mst-instance

Description This command specifies the port priority within a given instance, expressing probability that a given 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 <i>seconds</i> no max-age
Context	config>service>vpls>stp config>template>vpls-template>stp
Description	<p>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 thus reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.</p> <p>STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.</p> <p>The no form of this command returns the max age to the default value.</p>
Default	20 seconds
Parameters	<i>seconds</i> — 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>template>vpls-template>stp
Description	<p>This command specifies the version of Spanning Tree Protocol the bridge is currently running. See section Spanning Tree Operating Modes on page 383 for details on these modes.</p> <p>The no form of this command returns the STP variant to the default.</p>
Default	rstp
Parameters	<p>rstp — Corresponds to the Rapid Spanning Tree Protocol specified in IEEE 802.1D/D4-2003.</p> <p>dot1w — Corresponds to the mode where the Rapid Spanning Tree is backward compatible with IEEE 802.1w.</p> <p>compdot1w — Corresponds to the Rapid Spanning Tree Protocol fully conformant to IEEE 802.1w.</p> <p>mstp — Sets MSTP as the STP mode of operation. Corresponds to the Multiple Spanning Tree Protocol specified in 802.1Q REV/D5.0-09/2005</p> <p>pmstp — The PMSTP mode is only supported in VPLS services where the mVPLS flag is configured.</p>

mst-instance

General Switch Management Protocol Commands

Syntax	[no] mst-instance <i>mst-inst-number</i>
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	<i>mst-inst-number</i> — Specifies the Multiple Spanning Tree instance. Values 1 — 4094

mst-priority

Syntax	mst-priority <i>bridge-priority</i> 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 <i>bridge-priority</i> 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 BPDU's 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	<i>bridge-priority</i> — Specifies the priority of this specific Multiple Spanning Tree Instance for this service. Values 0 — 65535

vlan-range

Syntax	[no] vlan-range [<i>vlan-range</i>]
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 given mst-instance is shutdown.

The **no** form of this command removes the **vlan-range** from given **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 — 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 — 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

General Switch Management Protocol Commands

Syntax	mst-revision <i>revision-number</i>
Context	config>service>vpls>stp
Description	<p>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.</p> <p>The no form of this command returns MST configuration revision to its default value.</p>
Default	0
Parameters	<i>revision-number</i> — Specifies the MSTP region revision number to define the MSTP region.
	Values 0 — 65535

path-cost

Syntax	path-cost <i>sap-path-cost</i> no path-cost
Context	config>service>vpls>sap>stp config>service>vpls>spoke-sdp>stp
Description	<p>This command configures the Spanning Tree Protocol (STP) path cost for the SAP or spoke SDP.</p> <p>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.</p> <p>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 and 7950 XRS the STP path cost is a purely static configuration.</p> <p>The no form of this command returns the path cost to the default value.</p> <p><i>path-cost</i> — The path cost for the SAP or spoke SDP.</p> <p>Values 1 — 200000000 (1 is the lowest cost)</p> <p>Default 10</p>

port-num

Syntax	[no] port-num <i>virtual-port-number</i>
Context	config>service>vpls>sap>stp config>service>vpls>spoke-sdp>stp
Description	<p>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</p>

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 <i>bridge-priority</i> no priority
Context	config>service>vpls>stp config>template>vpls-template>stp
Description	The <i>bridge-priority</i> 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	<i>bridge-priority</i> — The bridge priority for the STP instance. Values Allowed values are integers in the range of 4096 — 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 <i>stp-priority</i> no priority
Context	config>service>vpls>spoke-sdp config>service>vpls>sap>stp
Description	This command configures the Alcatel-Lucent 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.

General Switch Management Protocol Commands

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, thus the actual value range is 0 to 240 in increments of 16.

Default 128

VPLS SAP Commands

sap

sap *sap-id* [**split-horizon-group** *group-name*] [**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 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 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. See [Common CLI Command Descriptions on page 1271](#) for command syntax.

create — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

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

General Switch Management Protocol Commands

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.

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.

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 only available when the VPLS is designated as an E-Tree VPLS.

cflowd

Syntax	[no] cflowd
Context	config>service>vpls>sap
Description	<p>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.</p> <p>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.</p> <p>For Layer 2 services, only ingress sampling is supported.</p>
Default	no cflowd

discard-unknown-source

Syntax	[no] discard-unknown-source
Context	config>service>vpls>sap config>service>vpls>spoke-sdp
Description	<p>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 on page 663) 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.</p> <p>When disabled, the packets are forwarded based on the destination MAC addresses.</p> <p>The no form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.</p>
Default	no discard-unknown

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 enables 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 <i>ring-id</i> 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 Eth-ring path encapsulation. The no form of this command removes eth-ring from this SAP
Default	no eth-ring
Parameters	<i>ring-id</i> — Specifies the ring ID. Values 1-128

path

Syntax	path <i>path-index tag qtag[.qtag]</i> no path <i>path-index</i>
Context	config>service>vpls>sap>eth-tunnel

ETH-CFM Service Commands

Description	This command configures Ethernet tunnel SAP path parameters. The no form of the command removes the values from the configuration.
Default	none
Parameters	<i>path-index</i> — Specifies the path index value. Values 1 — 16 <i>tag qtag[.qtag]</i> — Specifies the qtag value. Values 0 — 4094, *

mep

Syntax	mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [direction { up down }] primary-vlan-enable [vlan <i>vlan-id</i>] no mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i>
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	<i>mep-id</i> — Specifies the maintenance association end point identifier. Values 1 — 8191 <i>md-index</i> — Specifies the maintenance domain (MD) index value. Values 1 — 4294967295 <i>ma-index</i> — Specifies the MA index value. Values 1 — 4294967295 direction up down — Indicates the direction in which the maintenance association (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 towards 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 can not 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 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 MEP.

vlan-id — Must match the *vlan-id* under the bridge-identifier for the MA that is appropriate for this service

Values 0 — 4094

mip

Syntax **mip** [**mac** *mac-address*] **primary-vlan-enable** [**vlan** *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 dependant on the mhf-creation configuration for the MA. This MIP option is only available for default and static mhf-creation methods.

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. 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 — 4094

Default **no mip**

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 dependant on the mhf-creation configuration for the MA. This MIP option is only available for default and static mhf-creation methods.

ETH-CFM Service Commands

Parameters	<p>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 can not 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.</p> <p>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.</p> <p><i>vlan-id</i> — Must match the vlan-id under the bridge-identifier for the MA that is appropriate for this service</p> <p>Values 0 — 4094</p>
-------------------	--

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 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 DOWN MEP is configured.

client-meg-level

Syntax	client-meg-level <i>[[/level/ [/level/ ...]]</i> 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	<i>level</i> — Specifies the client MEG level.
Values	1 — 7
Default	1

ccm-padding-size

Syntax	ccm-padding-size <i>ccm-padding</i> no ccm-padding-size <i>ccm-padding</i>
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	<i>ccm-padding</i> — specifies the byte size of the Optional Data TLV
Values	3 — 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 <i>multiplier-value</i> 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

ETH-CFM Service Commands

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	<i>multiplier-value</i> — Specifies the multiplier used for timing out CSF.
Values	0.0, 2.0 .. 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 <i>priority-value</i> 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	<i>priority-value</i> — Specify the priority value of the AIS messages originated by the node.
Values	0 — 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 <i>priority</i> 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	<i>priority</i> — Specifies the priority of CCM and LTM messages. Values 0 — 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 <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [<i>priority priority</i>] [<i>data-length data-length</i>] 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.

ETH-CFM Service Commands

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

Description config>service>vpls>mesh-sdp 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 — 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.

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

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 <i>mac-address</i> 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	<i>mac-address</i> — 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 <i>seconds</i>
Context	config>service>vpls>sap>eth-cfm>mep
Description	This command enables/disables eth-test functionality on MEP.
Parameters	<i>seconds</i> — Specifies the one way delay threshold, in seconds.
Values	0..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

ETH-CFM Service Commands

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.

Parameters	accept — Share fate with the facility tunnel MEP ignore — Do not share fate with the facility tunnel MEP
Default	ignore (Service Level) accept (SAP Level for Epipe and VPLS)

vmep-extensions

Syntax	[no] vmep-extensions
Context	<code>config>service>vpls>eth-cfm</code>
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 BVPLS 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	<code>no vmep-extensions</code> (for B-VPLS) <code>vmep-extensions</code> (for VPLS contexts not B-VPLS based)

vmep-filter

Syntax	[no] vmep-filter
Context	<code>config>service>vpls>eth-cfm>sap</code> <code>config>service>vpls>eth-cfm>spoke-sdp</code> <code>config>service>vpls>eth-cfm>mesh-sdp</code>
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	<code>no vmep-filter</code>

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 — 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 — When specified, 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 given SAP for duration of its age-timer. The age of the MAC address entry in the FIB is set by the age timer. If mac-aging is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with mac-pinning enabled will remain in the FIB on this SAP/SDP forever. Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).
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 <i>table-size</i> no max-nbr-mac-addr
Context	config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>endpoint
Description	This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP, spoke SDP, or endpoint.

ETH-CFM Service Commands

When the configured limit has been reached, and `discard-unknown-source` has been enabled for this SAP or spoke SDP (see [discard-unknown-source on page 652](#)), 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 — 511999

Chassis-mode C limit: 196607

Chassis-mode D limit: 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 — 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 end-point peer.

ip-address — Specifies the IP address of multi-chassis end-point peer.

multi-service-site

Syntax	multi-service-site <i>customer-site-name</i> no multi-service-site
Context	config>service>vpls>sap
Description	<p>This command associates the SAP with a <i>customer-site-name</i>. If the specified <i>customer-site-name</i> does not exist in the context of the service customer ID an error occurs and the command will not execute. If <i>customer-site-name</i> exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within <i>customer-site-name</i> as parent schedulers.</p> <p>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.</p> <p>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.</p>
Default	None
	<p><i>customer-site-name</i> — 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.</p> <p>Values Any valid customer-site-name created within the context of the customer-id.</p>

precedence

Syntax	precedence [<i>precedence-value</i> 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	<p><i>precedence-value</i> — Specifies the precedence of this SDP bind.</p> <p>Values 1 — 4</p> <p><i>primary</i> — Assigns this as the primary spoke-sdp.</p>

static-isid

Syntax	[no] static-isid range <i>entry-id isid</i> [to <i>isid</i>] [create]
Context	config>service>vpls><instance> b-vpls>sap

config>service>vpls><instance> b-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 to execute the associated action for that entry.

no isid - 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..16777215

to isid — Identifies upper value in a range of ISIDs to be used as matching criteria.

Values 0..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.
- Note that 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 enables 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 <i>vlan-range</i>
Context	config>service>vpls>sap>managed-vlan-list
Description	<p>This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.</p> <p>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.</p> <p>To modify the range of VLANs, first the new range should be entered and afterwards the old range removed. See Modifying VPLS Service Parameters on page 539.</p>
Default	None
Parameters	<i>vlan-range</i> — 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 — 4094 end-vlan: 0 — 4094

VPLS SAP ATM Commands

atm

Syntax	atm
Context	config>service>vpls>sap
Description	<p>This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:</p> <ul style="list-style-type: none"> • 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. <p>If ATM functionality is not supported for a given context, the command returns an error.</p>

egress

Syntax	egress
Context	config>service>vpls>sap>atm
Description	This command enables the context to configure egress ATM attributes for the SAP.

encapsulation

Syntax	encapsulation <i>atm-encap-type</i>
Context	config>service>vpls>sap>atm
Description	<p>This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684, <i>Multiprotocol Encapsulation over ATM AAL5</i>, and to the ATM Forum LAN Emulation specification.</p> <p>Ingress traffic that does not match the configured encapsulation will be dropped.</p>
Default	The encapsulation is driven by the services for which the SAP is configured. For IES and VPRN service SAPs, the default is aal5snap-routed .
Parameters	<p><i>atm-encap-type</i> — Specify the encapsulation type.</p> <p>Values</p> <ul style="list-style-type: none"> 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.

ingress

Syntax	ingress
Context	config>service>vpls>sap>atm
Description	This command enables the context to configure ingress ATM attributes for the SAP.

traffic-desc

Syntax	traffic-desc <i>traffic-desc-profile-id</i> no traffic-desc
Context	config>service>vpls>sap>atm>ingress config>service>vpls>sap>atm>egress
Description	<p>This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).</p> <p>When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.</p> <p>When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.</p> <p>The no form of the command reverts the traffic descriptor to the default traffic descriptor profile.</p>
Default	The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.
Parameters	<i>traffic-desc-profile-id</i> — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax	oam
Context	config>service>vpls>sap>atm
Description	<p>This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.</p> <p>The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):</p> <ul style="list-style-type: none">• 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	<p>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.</p> <p>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 DOWN, or enters a fault state and comes back 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).</p> <p>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. Note that when alarm-cells is disabled, a PVCC will change operational status to UP from 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.</p>
Default	Enabled for PVCCs delimiting VPLS SAPs.

VPLS Filter and QoS Policy Commands

egress

Syntax	egress
Context	config>service>vpls>sap
Description	This command enables 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 enables 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 to be active on the context's object (SAP, subscriber, VPORT etc.).

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 (or disable) aggregate rate overrun protection on the agg-rate context.

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 enable (or disable) frame based accounting on all queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports.

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 <i>group-name</i> [type <i>group-type</i>] [qos-per-member] [create] no encap-group <i>group-name</i>
Context	config>service>vpls>sap>egress>encap-defined-qos
Description	<p>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.</p> <p>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.</p> <p>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. Note that 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.</p> <p>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:</p> <pre>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</pre> <p>Note that 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 no qos command until all members are deleted from the encap-group.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>Note that 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</p>

LAG. The set of scheduler/**agg-rate** instances will be replicated per link or perXMA 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** — This 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 <i>ip-filter-id</i> filter ipv6 <i>ipv6-filter-id</i> filter mac <i>mac-filter-id</i> no filter [ip <i>ip-filter-id</i>] [mac <i>mac-filter-id</i>] [ipv6 <i>ipv6-filter-id</i>]
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	<p>This command associates an IP filter policy or MAC filter policy with an ingress or egress Service Access Point (SAP) or IP interface.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Special Cases	VPLS — Both MAC and IP filters are supported on a VPLS service SAP.
Parameters	<p>ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.</p> <p>Values 1 — 65535</p> <p>ipv6 <i>ipv6-filter-id</i> — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.</p> <p>Values 1 — 65535</p> <p>mac <i>mac-filter-id</i> — 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.</p> <p>Values 1 — 65535</p>

hsmda-queue-override

Syntax	[no] hsmda-queue-override
Context	config>service>vpls>sap>egress

Description This command enables 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 given SAP QoS policy.

Parameters *queue-id* — Specifies the queue ID to override.

Values 1 — 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 — 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 — 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

Description This command specifies the administrative PIR by the user.

Parameters *pir-rate* — Configures the administrative PIR specified by the user.

Values 1 — 40000000, max

wrr-weight

Syntax	wrr-weight <i>value</i> 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	<i>percentage</i> — Specifies the weight for the HSMDA queue. Values 1— 32

wrr-policy

Syntax	wrr-policy <i>hsmda-wrr-policy-name</i> 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	<i>hsmda-wrr-policy-name</i> — Specifies the existing HSMDA WRR policy name to associate to the queue.

secondary-shaper

	secondary-shaper <i>secondary-shaper-name</i> no secondary-shaper
Context	config>service>vpls>sap>egress>hsmda-queue-over
Description	This command configures an HSMDA secondary shaper. Note that an shaper override can only be configured on an HSMDA SAP.
Parameters	<i>secondary-shaper-name</i> — Specifies a secondary shaper name up to 32 characters in length.

multicast-group

Syntax	multicast-group <i>group-name</i> no multicast-group
Context	config>service>vpls>sap>egress
Description	This command places a VPLS Ethernet SAP into an egress multicast group. The SAP must comply with the egress multicast group's common requirements for member SAPs. If the SAP does not

comply, the command will fail and the SAP will not be a member of the group. Common requirements for an egress multicast group are listed below:

- If an egress-filter is specified on the egress multicast group, the SAP must have the same egress filter applied.
- If an egress-filter is not defined on the egress multicast group, the SAP cannot have an egress filter applied.
- If the egress multicast group has an encap-type set to null, the SAP must be defined on a port with the port encapsulation type set to null.
- If the egress multicast group has an encap-type set to dot1q, the SAP must be defined on a port with the port encapsulation type set to dot1q and the port's dot1q-etype must match the dot1q-etype defined on the egress multicast group.
- The access port the SAP is created on cannot currently be an egress mirror source.

Once a SAP is a member of an egress multicast group, the following rules apply:

- The egress filter defined on the SAP cannot be removed or modified. Egress filtering is managed at the egress multicast group for member SAPs.
- If the encapsulation type for the access port the SAP is created on is set to dot1q, the port's dot1q-etype value cannot be changed.
- Attempting to define an access port with a SAP that is currently defined in an egress multicast group as an egress mirror source will fail.

Once a SAP is included in an egress multicast group, it is then eligible for efficient multicast replication if the egress forwarding plane performing replication for the SAP is capable. If the SAP is defined as a Link Aggregation Group (LAG) SAP, it is possible that some links in the LAG are on forwarding planes that support efficient multicast replication while others are not. The fact that some or all the forwarding planes associated with the SAP cannot perform efficient multicast replication does not affect the ability to place the SAP into an Egress multicast group.

A SAP may be a member of one and only one egress multicast group. If the multicast-group command is executed with another egress multicast group name, the system will attempt to move the SAP to the specified group. If the SAP is not placed into the new group, the SAP will remain a member of the previous egress multicast group. Moving a SAP into an egress multicast group may cause a momentary gap in replications to the SAP destination while the move is being processed.

The **no** form of the command removes the SAP from any egress multicast group in which it may currently have membership. The SAP will be removed from all efficient multicast replication chains and normal replication will apply to the SAP. A momentary gap in replications to the SAP destination while it is being moved is possible. If the SAP is not currently a member in an egress multicast group, the command has no effect.

Default no multicast-group

Parameters *group-name* — The *group-name* is required when specifying egress multicast group membership on a SAP. An egress multicast group with the specified egress-multicast-group-name must exist and the SAP must pass all common requirements or the command will fail.

Values Any valid egress multicast group name.

Default None, an egress multicast group name must be explicitly specified.

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.

ETH-CFM Service Commands

Parameters *rate* | **max** — Specifies the max rate override in kilobits-per-second or use the maximum.
Values 1 — 20000000 Kbps, max

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 — 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* — 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 — 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** — This keyword signifies that size is expressed in bytes.
- kilobytes** — The optional kilobytes keyword signifies that size is expressed in kilobytes.
- Values** 0 – 16777216 or default

policer-control-policy

- Syntax** **policer-control-policy** *policy-name* [**create**]
no policer-control-policy
- Context** config>service>vpls>sap>egress
config>service>vpls>sap>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 thus 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 (in-profile / out-of-profile) 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.

In order 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	<p><i>policy-name</i> — 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.</p> <p>create — The keyword is required when a new policy is being created and the system is configured for explicit object creation mode.</p>

policer-override

Syntax	[no] policer-override
Context	config>service>vpls>sap>egress config>service>vpls>sap>ingress
Description	<p>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.</p> <p>The no form of the command is used to remove any existing policer overrides.</p>
Default	no policer-overrides

policer

Syntax	policer <i>policer-id</i> [create] no policer <i>policer-id</i>
Context	config>service>vpls>sap>egress>policer-override config>service>vpls>sap>ingress>policer-override
Description	<p>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.</p> <p>The no form of the command is used to remove any existing overrides for the specified policer-id.</p>

ETH-CFM Service Commands

- 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 — 16777216 or default

mbs

- Syntax** **mbs** *size* [**bytes** | *kilobytes*]
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 is used to restore the *policer?* *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 – 16777216

byte — When byte is defined, the value given for size is interpreted as the queue? MBS value given in bytes. When kilobytes is defined, the value is interpreted as the queue? MBS value given in kilobytes.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> }
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. The no packet-byte-offset command is used to restore the policer? packet-byte-offset setting to the policy defined value.
Default	no packet-byte-offset
Parameters	add <i>add-bytes</i> — 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 — 31 subtract <i>sub-bytes</i> — 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 — 64

rate

Syntax	rate { <i>rate</i> max } [cir { max <i>rate</i> }]
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 rate command is used to restore the policy defined metering and profiling rate to a policer.
Parameters	{ <i>rate</i> 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? 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 — 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 — 2000000000, max

stat-mode

Syntax **stat-mode** *stat-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. An egress policer has multiple types of offered packets (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 at 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. Once a policer has been made a child to a parent policer, the stat-mode cannot be changed to no-stats unless the policer parenting is first removed.

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. 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 no-stats.

The stat-mode setting defined for the policer in the QoS policy may be overridden on an sla-profile or 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 stat-mode command attempts to return the policer's stat-mode setting to no-stats. The command will fail if the policer is currently configured as a child policer using the policer's parent command. The no parent command must first be executed for the no stat-mode command to succeed.

Parameters

stat-mode — Specifies the mode of statistics collected by this policer.

Values no-stats, minimal, offered-profile-no-cir, offered-profile-cir, offered-total-cir

no-stats — Counter resource allocation: 0

The no-stats mode is the default stat-mode for the policer. The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policers parent command will fail.

When collect-stats is enabled, the lack of counters causes the system to generate the following statistics:

- | | |
|----------------|-----|
| a. offered-in | = 0 |
| b. offered-out | = 0 |
| c. discard-in | = 0 |
| d. discard-out | = 0 |
| e. forward-in | = 0 |
| f. forward-out | = 0 |

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation: 1 The minimal mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (soft or hard profile) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

- | | |
|--------------|---|
| 1. offered | <= soft-in-profile-out-of-profile, profile in/out |
| 2. discarded | <= Same as 1 |
| 3. forwarded | <= Derived from 1 – 2 |

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- | | |
|----------------|-----|
| a. offered-in | = 1 |
| b. offered-out | = 0 |
| c. discard-in | = 2 |
| d. discard-out | = 0 |
| e. forward-in | = 3 |
| f. forward-out | = 0 |

Counter 0 indicates that the accounting statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation: 2

The offered-profile-no-cir mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The offered-profile-no-cir mode is most useful when profile based offered, discard and forwarding stats are required from the ingress policer, but a CIR is not being used to recolor the soft in-profile and out-of-profile packets. This mode does not prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-in <= soft-in-profile, profile in
2. offered-out <= soft-out-of-profile, profile out
3. dropped-in <= Same as 1
4. dropped-out <= Same as 2
5. forwarded-in <= Derived from 1 – 3
6. forwarded-out <= Derived from 2 – 4

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

offered-profile-cir — Counter resource allocation: 3

The offered-profile-cir mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The offered-profile-cir mode is most useful when profile based offered, discard and forwarding stats are required from the ingress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

1. offered-in-that-stayed-green-or-turned-red <= profile in
2. offered-soft-that-turned-green <= soft-in-profile-out-of-profile
3. offered-soft-or-out-that-turned-yellow-or-red <= soft-in-profile-out-of-profile, profile out
4. dropped-in-that-stayed-green-or-turned-red <= Same as 1
5. dropped-soft-that-turned-green <= Same as 2
6. dropped-soft-or-out-that-turned-yellow-or-red <= Same as 3
7. forwarded-in-that-stayed-green <= Derived from 1 – 4
8. forwarded-soft-that-turned-green <= Derived from 2 – 5
9. forwarded-soft-or-out-that-turned-yellow <= Derived from 3 – 6

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2 + 3
- c. discard-in = 4
- d. discard-out = 5 + 6
- e. forward-in = 7 + 8
- f. forward-out = 9

offered-total-cir — Counter resource allocation: 2

The offered-total-cir mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The offered-total-cir mode is most useful when profile based offered stats are not required from the ingress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

1. offered-that-turned-green <= soft-in-profile-out-of-profile, profile in/out
2. offered- that-turned-yellow-or-red<= soft-in-profile-out-of-profile, profile in/out
3. dropped-offered-that-turned-green<= Same as 1
4. dropped-offered-that-turned-yellow-or-red<= Same as 2
5. forwarded-offered-that-turned-green<= Derived from 1 – 3
6. forwarded-offered-that-turned-yellow<= Derived from 2 – 4

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1 + 2 (Or 1 and 2 could be summed on b)
- b. offered-out = 0
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

Counter 0 indicates that the accounting statistic returns a value of zero.

qos

Syntax	qos <i>policy-id</i> [shared-queuing multipoint-shared] [fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i>] 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 given 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	<p><i>policy-id</i> — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.</p> <p>Values 1 — 65535</p> <p>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.</p> <p>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, in addition to the unicast packets, multipoint packets also used shared queues.</p> <p>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.</p> <p>When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.</p> <p>When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.</p> <p>Values Multipoint or not present.</p> <p>Default Present (the queue is created as non-multipoint).</p> <p>fp-redirect-group — This keyword creates an instance of a named queue group template on the ingress forwarding plane of a given 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.</p>

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/IMM/XMA ingress forwarding plane.

qos

Syntax	qos policy-id [port-redirect-group <i>queue-group-name</i> instance <i>instance-id</i>] no qos
Context	config>service>vpls>sap>egress
Description	<p>This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).</p> <p>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.</p> <p>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.</p> <p>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 given type will return an error.</p> <p>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.</p> <p>By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.</p> <p>The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.</p>
Default	none
Parameters	<p>port-redirect-group — This keyword associates a SAP egress with an instance of a named queue group template on the egress port of a given IOM/IMM/XMA. The queue-group-name and instance-id are mandatory parameters when executing the command.</p> <p><i>queue-group-name</i> — 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 <code>config>port>ethernet>access>egress</code>.</p> <p>instance <i>instance-id</i> — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.</p>

Values 1 — 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 enables 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* — The queue ID whose parameters are to be overridden.

Values 1 — 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 50×20 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 500×1.1 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 7500×1.1 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 — 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 given 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 drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

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 desired, 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 — 131072 or default

high-prio-only

ETH-CFM Service Commands

Syntax	high-prio-only <i>percent</i> no high-prio-only
Context	config>service>vpls>sap>egress>queue-override>queue config>service>vpls>sap>ingress>queue-override>queue
Description	<p>This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The high-prio-only command configures the percentage of buffer space for the queue, used exclusively by high priority packets.</p> <p>The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The high-prio-only parameter is used to override the default value derived from the network-queue command.</p> <p>The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.</p> <p>The no form of this command restores the default high priority reserved size.</p>
Parameters	<p><i>percent</i> — The <i>percent</i> parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.</p> <p>Values 0 — 100, default</p>

mbs

Syntax	mbs size [bytes kilobytes] no mbs
Context	config>service>vpls>sap>egress>queue-override>queue config>service>vpls>sap>ingress>queue-override>queue config>service>vpls>sap>egress>hsmda-queue-override>queue
Description	<p>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.</p> <p>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 proper 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.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The no form of this command returns the MBS size assigned to the queue.</p>
Default	default

Parameters *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

For `sap>egress>queue-override>queue` and `sap>ingress>queue-override>queue`

Values [0 — 1073741824, **default** in **bytes** or **kilobytes**

For `sap>egress>hsmda-queue-override>queue`

Values [0 — 2625][**kilobytes**] | [0 — 2688000]**bytes** | **default**

mbs

Syntax **mbs** {*size-in-kbytes* | **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 enqueueing 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 proper 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.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns the MBS size assigned to the queue to the default value.

Default default

Parameters *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

Values 0 — 131072 or default

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] 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	<p>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.</p> <p>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.</p> <p>The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-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 as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.</p> <p>The CIR can be used by the queue’s parent commands <i>cir-level</i> and <i>cir-weight</i> parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.</p> <p>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 <i>queue-id</i>.</p> <p>The no form of the command returns all queues created with the <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters (max, 0).</p>
Default	rate max cir 0 — The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.
Parameters	<p><i>pir-rate</i> — 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.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queue’s adaptation-rule parameters and the actual hardware where the queue is provisioned.</p> <p>For egress>queue-override>queue and ingress>queue-override>queue:</p> <p>Values 1 — 2000000000, max in Kbps</p> <p>Default max</p> <p>For egress>hsmda-queue-over>queue:</p> <p>Values 1 — 100000000, max in Kbps</p> <p>Default max</p> <p>cir <i>cir-rate</i> — 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.</p> <p>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 or queues.</p>

For **egress>queue-override>queue** and **ingress>queue-override>queue**:

Values 0 — 2000000000, **max** in Kbps

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 enables 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 <i>queue-id</i>
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	<i>queue-id</i> — The queue ID whose parameters are to be overridden.
Values	1 — 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 50×20 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 500×1.1 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 7500×1.1 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.

ETH-CFM Service Commands

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 — 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 given 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 drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

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 desired, 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 — 131072 or default

high-prio-only

Syntax	high-prio-only <i>percent</i> no high-prio-only
Context	config>service>vpls>sap>egress>queue-override>queue config>service>vpls>sap>ingress>queue-override>queue
Description	<p>This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The high-prio-only command configures the percentage of buffer space for the queue, used exclusively by high priority packets.</p> <p>The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The high-prio-only parameter is used to override the default value derived from the network-queue command.</p> <p>The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.</p> <p>The no form of this command restores the default high priority reserved size.</p>
Parameters	<p><i>percent</i> — The <i>percent</i> parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.</p> <p>Values 0 — 100, default</p>

mbs

Syntax	mbs { <i>size-in-kbytes</i> default } no mbs
Context	config>service>vpls>sap>egress>queue-override>queue
Description	<p>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.</p> <p>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 proper 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.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The no form of this command returns the MBS size assigned to the queue.</p>
Default	default

ETH-CFM Service Commands

Parameters *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

Values 0 — 131072 or default

mbs

Syntax **mbs** {*size-in-kbytes* | **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 enqueueing 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 proper 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.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns the MBS size assigned to the queue to the default value.

Default default

Parameters *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

Values 0 — 131072 or default

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 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 packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile 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** — The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

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 — 100000000

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 or queues.

Values 0 — 100000000, **max**, **sum**

Default 0

wred-queue-policy

Syntax **wred-queue-policy** *slope-policy-name*
no wred-queue-policy

Context config>service>vpls>sap>egress>queue-override>queue

ETH-CFM Service Commands

- Description** The `wred-queue-policy` command is used on an egress SAP to override the slope policy associated with a WRED queue. When specified, the SAP egress QoS policy derived slope policy is ignored and the configured override slope policy is applied to the WRED queue. The specified *queue-id* must be a WRE-enabled queue to be successful.
- The **no** form of the command removes the slope policy override for the WRED queue on the egress SAP.
- Parameters** *slope-policy-name* — Overrides the SAP Egress QoS policy derived WRED slope policy for the specified queue-id. The defined slope policy must exist or the command will fail.

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 ingress 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 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 queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). 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 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 criteria below, 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, etc.), 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>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** *weight* — **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 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
- Default** 1
- cir-weight** *cir-weight* — The **cir-weight** keyword defines 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 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 — 100
- Default** 1

rate

- Syntax** **rate** *pir-rate* [**cir** *cir-rate*]
no rate
- Context** config>service>vpls>sap>egress>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 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 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 queues and schedulers to operate at their maximum rates.

The **no** form of this command returns all queues created with this *queue-id* by association with the QoS policy to the default PIR and CIR parameters.

Parameters *pir-rate* — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** is accepted. Any other value will result in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue's **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (*fc*) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue will be allowed to forward packets in a given second, thus shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queues PIR rate to the default value, that value must be specified as the PIR value.

Values 1 — 100000000, **max**

Default max

cir cir-rate — The **cir** parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue will operate. A value of 0 — 10000000 or the keyword **max** or **sum** is accepted. Any other value will result in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir pir-rate** is multiplied by the *cir cir-rate*. If the **cir** is set to max, then the CIR rate is set to infinity.

The SAP ingress context for CIR is dependent on the defined forwarding class (*fc*) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

Values 0 — 10000000, **max**, **sum**

Default sum

scheduler-policy

Syntax	scheduler-policy <i>scheduler-policy-name</i> 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 associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the config>qos>scheduler-policy <i>scheduler-policy-name</i> 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. 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 ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

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 created on associated SAPs.

Values Any existing valid scheduler policy name.

vlan-translation

Syntax	vlan-translation {vlan-id copy-outer} no vlan-translation
Context	config>service>vpls>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 the copy-outer keyword, the outer VLAN ID will be copied to the inner position on QinQ-encapsulated ports. The feature is not supported on default-dot1q SAPs (1/1/1:* and 1/1/1:0), as well as on TopQ SAPs. The no form of this command sets the default value, and no action will be taken.
Default	per default the preserved VLAN values will not be overwritten
Parameters	<i>vlan-id</i> — Specifies the to use the VLAN ID of the SAP. Values 0 — 4094 <i>copy-outer</i> — Specifies that the outer VLAN ID will be copied to the inner position on QinQ-encapsulated ports

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 13](#) defines the default behavior for Dot1P evaluation.

Table 13: 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.

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

Port / SAP Type	Existing Packet Tags	PBits Used for Match
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.

Table 14: 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

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

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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.

Note that 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**.

policer-control-policy

Syntax **policer-control-policy** *policy-name* [create]
no policer-control-policy

Context config>service>vpls>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. 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 thus 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 (in-profile / out-of-profile) 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.

In order 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* — 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.

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* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 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	<p>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.</p> <p>When the no collect-stats command is issued the statistics are still accumulated by the 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.</p>
Default	no collect-stats

VPLS Template Commands

template

Syntax	template
Context	config>service
Description	This is the node for service templates.

vpls-template

Syntax	vpls-template <i>name/id</i> create [no] vpls-template <i>name/id</i>
Context	config>service>template
Description	<p>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 VPLSes (no template configuration necessary) and the operator cannot change these values.</p> <p>vpn-id: none</p> <p>description: "Service <svc id> auto-generated by control VPLS <svc-id>"</p> <p>service-name: "Service <svc id>" (Auto-generated)</p> <p>shutdown: no shutdown</p> <p>Following existing attributes can be set by the user in the instantiated VPLSes:</p> <p>[no] sap</p> <p>All the other VPLS attributes are not supported.</p>
Parameters	<i>name/id</i> — Specifies the name in ASCII or the template ID.
Values	name: ASCII string
Values	ID: [1..2147483647]

vpls-sap-template

Syntax	vpls-sap-template <i>name/id</i> create [no] vpls-sap-template <i>name/id</i>
Context	config>service>template

ETH-CFM Service Commands

- 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..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. Note that 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 FIB 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 FIB 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 — 600

Provider Tunnel Commands

provider-tunnel

Syntax	provider-tunnel
Context	configure>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	configure>service>vpls>provider-tunnel
Description	<p>This 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-VPLS instance. The P2MP LSP is referred to, in this case, as the Inclusive Provider Multicast Service Interface (I-PMSI).</p> <p>When enabled, this feature relies on BGP Auto-Discovery (BGP-AD) or BGP-VPLS to discover the PE nodes participating in a given 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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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:</p> <pre>config>service>vpls [b-vpls]>provider-tunnel>inclusive>rsvp>lsp-template p2mp-lsp-template-name</pre> <p>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:</p> <pre>config>service>vpls [b-vpls]>provider-tunnel>inclusive>mldp</pre> <p>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.</p> <p>The user can specify if the node is both root and leaf in the VPLS instance:</p>

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. Note that 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-sdp's.

Note that 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 <i>seconds</i> no data-delay-interval
Context	configure>service>vpls>provider-tunnel>inclusive
Description	<p>This command configures the I-PMSI data delay timer.</p> <p>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.</p> <p>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. Note that 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>interface>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.</p> <p>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.</p> <p>The no version of this command re-instates the default value for this delay timer.</p>
Parameters	<i>seconds</i> — The delay time value in seconds.

ETH-CFM Service Commands

Values 3—180 seconds

Default 15 seconds

mldp

Syntax [no] mldp

Context configure>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-VPLS instance.

root-and-leaf

Syntax [no] root-and-leaf

Context configure>service>vpls>provider-tunnel>inclusive

Description This command configures the node to operate as both root and leaf of the I-PMSI in a given 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 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. Note that 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-sdp's..

The **no** version of this command re-instates the default value.

rsvp

Syntax [no] rsvp

Context configure>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 <i>p2mp-lsp-template-name</i> no lsp-template
Context	configure>service>vpls>provider-tunnel>inclusive>rsvp
Description	<p>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-PMIS).</p> <p>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.</p> <p>The no version of this command removes the P2MP LSP template from the I-PMIS configuration.</p>
Parameters	<i>p2mp-lsp-template-name</i> — The name of the P2MP LSP template. This is a string of 32 characters maximum.
	Default None

VPLS SDP Commands

mesh-sdp

Syntax	mesh-sdp <i>sdp-id</i> [: <i>vc-id</i>] [vc-type { ether vlan }] [root-leaf-tag leaf-ac] no mesh-sdp <i>sdp-id</i> [: <i>vc-id</i>]
Context	config>service>vpls
Description	<p>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.</p> <p>Note that 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.</p> <p>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 <i>sdp-id</i> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</p> <p>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.</p> <p>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.</p>
Default	No <i>sdp-id</i> is bound to a service.
Special Cases	VPLS — Several SDPs can be bound to a VPLS. Each SDP must be destined to a different router. If two <i>sdp-id</i> bindings terminate on the same router, an error occurs and the second SDP is binding is rejected.
Parameters	<p><i>sdp-id</i> — The SDP identifier.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — 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.</p> <p>Values 1 — 4294967295</p> <p>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 <i>draft-martini-l2circuit-trans-mpls</i>.</p> <ul style="list-style-type: none"> • 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 **ether** 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 **ether** 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.

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 Etree VPLS.

spoke-sdp

Syntax	spoke-sdp <i>sdp-id[:vc-id]</i> [vc-type { ether vlan }] [split-horizon-group <i>group-name</i>] endpoint [no-endpoint] [root-leaf-tag leaf-ac] no spoke-sdp <i>sdp-id[:vc-id]</i>
Context	config>service>vpls
Description	<p>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.</p> <p>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.</p> <p>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 <i>sdp-id</i> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</p> <p>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.</p>

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* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 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.

split-horizon-group *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.

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 Etree VPLS.

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 configures the egress SDP context.

qos

Syntax	qos <i>network-policy-id</i> port-redirect-group <i>queue-group-name</i> instance <i>instance-id</i> no qos [<i>network-policy-id</i>]
Context	configure>service>apipe>spoke-sdp>egress configure>service>cpipe>spoke-sdp>egress

```
configure>service>epipe>spoke-sdp>egress
configure>service>fpipe>spoke-sdp>egress
configure>service>ipipe>spoke-sdp>egress
config>service>vpls>spoke-sdp>egress
config>service>vpls>mesh-sdp>egress
config>service>pw-template>egress
config>service>vprn>interface>spoke-sdp>egress
config>service>ies>interface>spoke-sdp>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 or 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-SPD inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SPDs 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-SPD 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-SPD 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:

- a 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.
 - b 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.
4. 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, 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, but the DSCP is not modified by the policer operation.

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 — 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 — 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 <i>network-policy-id</i> fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i> no qos
Context	configure>service>apipe>spoke-sdp>ingress configure>service>cpipe>spoke-sdp>ingress configure>service>epipe>spoke-sdp>ingress configure>service>fpipe>spoke-sdp>ingress configure>service>ipipe>spoke-sdp>ingress config>service>vpls>spoke-sdp>ingress config>service>vpls>mesh-sdp>ingress config>service>pw-template>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>ies>interface>spoke-sdp>ingress
Description	<p>This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.</p> <p>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.</p> <p>Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:</p> <ol style="list-style-type: none"> 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 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. 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. 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. 5. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance. <p>The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:</p> <ol style="list-style-type: none"> 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. 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-SPD 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.
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:
 - a 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*.
 - b 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.
 4. 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.
 5. If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:
 - a 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.
 - b 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-SPD, 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 — 65535
- fp-redirect-group** *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.
- ingress-instance** *instance-id* — Specifies the identification of a specific instance of the queue-group.
- Values** 1 — 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 enables 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 <i>mda-id</i>
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	<i>mda-id</i> — Specifies an MFIB-allowed MDA destination.
Values	slot/mda slot: 1 — 10 mda: 1 — 2

vc-label

Syntax	vc-label <i>egress-vc-label</i> no vc-label [<i>egress-vc-label</i>]
Context	config>service>vpls>mesh-sdp>egress config>service>vpls>spoke-sdp>egress
Description	This command configures the egress VC label.
Parameters	<i>vc-label</i> — A VC egress value that indicates a specific connection.
Values	16 — 1048575

vc-label

Syntax	vc-label <i>ingress-vc-label</i> no vc-label [<i>ingress-vc-label</i>]
Context	config>service>vpls>mesh-sdp>ingress config>service>vpls>spoke-sdp>ingress
Description	This command configures the ingress VC label.
Parameters	<i>vc-label</i> — A VC ingress value that indicates a specific connection.
Values	2048 — 18431

static-mac

Syntax	[no] static-mac <i>ieee-mac-address</i>
Context	config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp

ETH-CFM Service Commands

Description	<p>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).</p> <p>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.</p> <p>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.</p> <p>Note that 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.</p> <p>Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.</p> <p>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.</p>
Default	none
Parameters	<p><i>ieee-mac-address</i> — Specifies the 48-bit MAC address for the static ARP in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i>, <i>bb</i>, <i>cc</i>, <i>dd</i>, <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.</p>

transit-policy

Syntax	transit-policy prefix <i>prefix-aasub-policy-id</i> no transit-policy
Context	config>service>vpls>spoke-sdp
Description	<p>This command assigns a transit policy id.</p> <p>The no form of the command removes the transit policy ID from the spoke SDP configuration.</p>
Default	no transit-policy
Parameters	<p><i>prefix-aasub-policy-id</i> — Specifies the transit policy ID.</p> <p>Values 1 — 65535</p>

vlan-vc-tag

Syntax	vlan-vc-tag <i>0..4094</i> no vlan-vc-tag [<i>0..4094</i>]
Context	config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp
Description	<p>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</p>

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..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

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
Description	This command enables fast leave. When IGMP or MLD fast leave processing is enabled, the SR OS router 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 thus 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 <i>vpls-id</i> 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	<i>vpls-id</i> — Specifies the MVR VPLS from which multicast channels should be copied into this SAP. Values <i>service-id:</i> 1 — 2147483648

group

Syntax	[no] group <i>grp-address</i>
Context	config>service>vpls>sap>igmp-snooping>static config>service>vpls>spoke-sdp>igmp-snooping>static config>service>vpls> >igmp-snooping>static

```
config>service>vpls>sap>mld-snooping>static
config>service>vpls>spoke-sdp>mld-snooping>static
config>service>vpls>mesh-sdp>mld-snooping>static
```

Description	This command adds a static multicast group as a (*, g). 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	<i>grp-address</i> — Specifies an IGMP or MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

group-policy

Syntax	group-policy <i>policy-name</i> no group-policy
Context	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	<i>policy-name</i> — 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, etc.), 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, see Enabling IGMP Group Membership Report Filtering in the SR OS Triple Play Guide.

fault-propagation-bmac

Syntax	fault-propagation-bmac [<i>mac-name</i> <i>ieee-address</i>] [create] no fault-propagation-bmac [<i>mac-name</i> <i>ieee-address</i>]
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).

VPLS Multicast Commands

- 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.

feature

- Syntax** [no] feature
- Context** config>service>vpls>sap
config>service>ies|vprn
- Description** This command enables feature.

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

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 Ether type inserted and used to determine the presence of a received VLAN tag for both VLAN tags is 0x8100. A different Ether type can be used for the outer VLAN tag by configuring the PW template with **use-provisioned-sdps** and setting the Ether type using the SDP **vlan-vc-etype** parameter (this Ether type value is then used for all mesh/spoke SDPs using that SDP).

The **no** form of this command sets 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
```

This command forces vc-vlan-type forwarding in the data path for spoke/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-label [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 LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the **ldp**, **rsvp-te**, or **mpls** options. This feature is not supported on a service bound to a GRE SDP. This feature is also not supported on multicast packets forwarded using RSVP P2MP LPS 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).

In order 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/VRPN 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 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 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 VRPN 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 enables the context to configure IGMP host tracking parameters.

disable-router-alert-check

Syntax	[no] disable-router-alert-check
Context	config>service>vpls>igmp-snooping 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 <i>expiry-time</i> no expiry-time
Context	config>service>vpls>igmp-snooping config>service>vpls>sap>igmp-snooping
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	<i>expiry-time</i> — Specifies the time, in seconds, that this system continues to track an inactive host.
Values	1 — 65535

import

Syntax	import <i>policy-name</i> no import
Context	config>service>vpls>sap>igmp-snooping

VPLS Multicast Commands

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	<i>policy-name</i> — Specifies the import policy name.

max-num-groups

Syntax	max-num-groups <i>max-num-groups</i> no max-num-groups
Context	config>service>vpls>sap>igmp-snooping
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	<i>max-num-groups</i> — Specifies the maximum number of multicast groups allowed to be tracked. Values 1 — 196607

max-num-sources

Syntax	max-num-sources <i>max-num-sources</i> no max-num-sources
Context	config>service>vpls>sap>igmp-host-traking 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	<i>max-num-sources</i> — Specifies the maximum number of multicast sources allowed per group. Values 1 — 1000

max-num-grp-sources

Syntax	max-num-grp-sources [1..32000] no max-num-grp-sources
Context	config>service>vpls>sap>igmp-host-traking 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..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
- 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** — No import policy is specified.
- Parameters** *policy-name* — 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, etc.), 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
- 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.
- Default** 10
- Parameters** *seconds* — Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 — 50

mcac

Syntax	mcac
Context	config>service>vpls>mesh-sdp>snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>sap>igmp-snooping
Description	This command configures multicast CAC policy and constraints for this interface.
Default	none

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>service>vpls>mesh-sdp>snooping>mcac config>service>vpls>spoke-sdp>snooping>mcac config>service>vpls>sap>igmp-snooping>mcac
Description	This command configures the multicast CAC policy name.
Parameters	<i>policy-name</i> — 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, etc.), the entire string must be enclosed within double quotes.

unconstrained-bw

Syntax	unconstrained-bw <i>bandwidth</i> mandatory-bw <i>mandatory-bw</i> no unconstrained-bw
Context	config>service>vpls>mesh-sdp>snooping>mcac config>service>vpls>spoke-sdp>snooping>mcac config>service>vpls>sap>igmp-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	<i>bandwidth</i> — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).
Values	0 — 2147483647

mandatory-bw *mandatory-bw* — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).

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 given 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 — 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
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 proper operation on mixed port-speed LAGs and can be used for non-mixed port-speed LAGs as well.
Default	no use-lag-port-weight — 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
Description	This command enables the context to configure multicast CAC constraints.
Default	none

level

Syntax	level <i>level-id</i> bw <i>bandwidth</i> no level <i>level-id</i>
Context	config>service>vpls>sap>igmp-snooping>mcac>mc-constraints
Description	This command configures levels and their associated bandwidth for multicast cac policy on this interface.

VPLS Multicast Commands

- Parameters** *level-id* — Specifies has an entry for each multicast CAC policy constraint level configured on this system.
- Values** 1 — 8
- bandwidth* — Specifies the bandwidth in kilobits per second (kbps) for the level.
- Values** 1 — 2147483647

number-down

- Syntax** **number-down** *number-lag-port-down*
no number-down
- Context** config>service>vpls>sap>igmp-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* — If 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 — 64 (for 64-link LAG)
1 — 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 — 1000

max-num-sources

Syntax	max-num-sources <i>max-num-sources</i> 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	<i>max-num-sources</i> — Specifies the maximum number of multicast sources allowed per group.
Values	1 — 1000

mrouter-port

Syntax	[no] mrouter-port
Context	config>service>vpls>sap>igmp-snooping config>service>vpls>spoke-sdp>igmp-snooping config>service>vpls>mesh-sdp>igmp-snooping config>service>vpls>allow-ip-int-bind>igmp-snp
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 reports generated by the system as a result of someone 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. Note that 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 can not be set.
Default	no mrouter-port

mvr

Syntax	mvr
Context	config>service>vpls>igmp-snooping

VPLS Multicast Commands

```
config>service>vpls>mld-snooping
config>service>vpls>sap>igmp-snooping
```

Description This command enables 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* — The time interval, in seconds, that the router transmits general host-query messages.

Values 2 — 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	<i>seconds</i> — Specifies the length of time to wait to receive a response to the host-query message from the host.
Values	1 — 1023

query-src-ip

Syntax	query-src-ip <i>ipv6-address</i> 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 <i>address</i> 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	<i>ip-address</i> — The source IP source address in transmitted IGMP reports.

robust-count

Syntax	robust-count <i>robust-count</i> no robust-count
Context	config>service>vpls>igmp-snooping

VPLS Multicast Commands

```
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 — 7
config>service>vpls>igmp-snooping: 1 — 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	<i>value</i> — Specifies the number of attributes accepted on a per BVPLS basis.
Values	1 — 4095 for MVRP

attribute-table-size

Syntax	[no] attribute-table-size <i>value</i>
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	<i>value</i> — Specifies the number of attributes accepted on a per BVPLS basis.
Values	SR-7/SR-12: 1 — 2047

attribute-table-high-wmark

Syntax	[no] attribute-table-high-wmark <i>high-water-mark</i>
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	<i>high-water-mark</i> — 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% — 100%

attribute-table-low-wmark

VPLS Multicast Commands

Syntax	[no] attribute-table-low-wmark <i>low-water-mark</i>
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	<i>low-water-mark</i> — Specifies utilization of the MRP attribute table of this service at which a table full alarm will be cleared by the agent. Values 1% — 100%

flood-time

Syntax	flood-time <i>flood-time</i> 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	<i>flood-time</i> — Specifies the MRP flood time, in seconds. Values 3 — 600

join-time

Syntax	[no] join-time <i>value</i>
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	<i>value</i> — Specifies the timer value in 10th of seconds for sending join-messages. Values 1 — 10 tenths of a second

leave-time

Syntax	[no] leave-time <i>value</i>
Context	config>service>vpls>sap>mrp config>service>vpls>spoke-sdp>mrp config>service>vpls>mesh-sdp>mrp
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>Refer to IEEE 802.1ak-2007 section 10.7.4.2.</p>
Default	30
Parameters	<i>value</i> — [30-60] tenths of a second

leave-all-time

Syntax	[no] leave-all-time <i>value</i>
Context	config>service>vpls>sap>mrp config>service>vpls>spoke-sdp>mrp config>service>vpls>mesh-sdp>mrp
Description	<p>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.</p>
Default	100
Parameters	<i>value</i> — [60-300] tenths of a second

mrp-policy

Syntax	[no] mrp-policy <i>name</i>
Context	config>service>vpls>sap>mrp config>service>vpls>spoke-sd>mrp config>service>vpls>mesh-sdp>mrp

VPLS Multicast Commands

Description	This command instructs MMRP to use the <i>mrp-policy</i> 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 <i>policy-name</i> refers to a non-existing <i>mrp-policy</i> the command should return error. Changes to a <i>mrp-policy</i> are allowed and applied to the SAP/SDPs under which the policy is referenced.
Default	no <i>mrp-policy</i> is defined
Parameters	<i>policy-name</i> — 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.

periodic-time

Syntax	[no] periodic-time <i>value</i>
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 Transmitting Timer is set to one second when it is started.
Default	10
Parameters	<i>value</i> — [10-100] tenths of a second

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 <i>policy-name</i> 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>spoke-sdp config>service>vpls>sap
Context	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 wildcard 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 <i>num-groups</i> no max-num-groups
Context	config>service>vpls>pim-snooping config>service>vpls>spoke-sdp>pim-snooping config>service>vpls>sap>pim-snooping
Description	This command configures the maximum groups for PIM snooping.
Parameters	<i>num-groups</i> — Specifies the maximum groups for PIM snooping.
Values	1 — 16000

oper-group

Syntax	[no] oper-group <i>name</i>
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 <i>name</i> . The oper-group <i>name</i> 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	<i>name</i> — A character string of maximum 32 ASCII characters identifying the group instance.

monitor-oper-group

Syntax	[no] monitor-oper-group <i>name</i>
Context	config>service>vpls>site

VPLS Multicast Commands

```
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 <i>name</i> 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	<i>name</i> — A character string of maximum 32 ASCII characters identifying the group instance.

hold-time

Syntax	hold-time <i>seconds</i> 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	<i>seconds</i> — Specifies the PIM snooping hold time, in seconds Values 0 — 300 Default 90

mode

Syntax	mode <i>mode</i>
Context	config>service>vpls>pim-snooping
Description	This command sets the PIM snooping mode to proxy or plain snooping.
Parameters	<i>mode</i> — Specifies PIM snooping mode. Values snooping, proxy Default proxy

precedence

Syntax	precedence <i>precedence-value</i> primary no precedence
Context	config>service>vpls>spoke-sdp
Description	This command configures the spoke SDP precedence.
Default	4
Parameters	<i>precedence-value</i> — Specify the spoke SDP precedence. Values 0 — 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 7x50 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 7x50 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 given 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

VPLS Multicast Commands

```
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*
- 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 (s,g) entry, to allow multicast traffic for a multicast group from a specified source. For a multicast group, more than one source address can be specified. Static (s,g) entries cannot be added, if a starg is previously created. The **no** form of the command removes the source from the configuration.
- Default** none
- Parameters** *ip-address* — Specifies the IPv4 or 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 to allow multicast traffic for the corresponding multicast group from any source. This command can only be enabled if no existing source addresses for this group are specified. The **no** form of the command removes 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* — Specify the IGMP or MLD version.

Values 1, 2, 3

to-sap

Syntax	to-sap <i>sap-id</i> 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 behaviour) but to another SAP. This command configures the SAP to which the multicast data needs to be copied.
Default	no to-sap
Parameters	<i>sap-id</i> — Specifies the SAP to which multicast channels should be copied. See Common CLI Command Descriptions on page 1271 for command syntax.

app-profile

Syntax	app-profile <i>app-profile-name</i> no app-profile
Context	config>service>vpls>sap
Description	This command configures the application profile name.
Parameters	<i>app-profile-name</i> — 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 enables the context to configure ARP host parameters.

host-limit

Syntax	host-limit <i>max-num-hosts</i> 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	<i>max-num-hosts</i> — specifies the maximum number of ARP hosts allowed on this SAP. Values 1 — 32767

min-auth-interval

Syntax	min-auth-interval <i>min-auth-interval</i> 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	<i>min-auth-interval</i> — Specifies the minimum authenticational interval, in minutes. Values 1 — 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 to 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.

VPLS Multicast Commands

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	not enabled
Parameters	<p>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.</p> <p>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.</p> <p>When arp-reply-agent is enabled with sub-ident:</p> <ul style="list-style-type: none">• 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	<p>Enabling force-l2pt-boundary will force that all SAPs managed by the given 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.</p> <p>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.</p>

If force-l2pt-boundary is enabled on a given 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 enables 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 <i>threshold</i> 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	<i>threshold</i> — Specifies the maximum length of a fragment to be transmitted.
Values	128 — 512

interleave

VPLS Multicast Commands

Syntax	interleave no interleave
Context	config>service>vpls>sap>frame-relay>frf.12
Description	<p>This command enables interleaving of high priority frames and low-priority frame fragments within a FR SAP using FRF.12 end-to-end fragmentation.</p> <p>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).</p> <p>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.</p> <p>The receive direction of the FR SAP supports both modes of operation concurrently, with and without fragment interleaving.</p> <p>The no form of this command restores the default mode of operation.</p>
Default	no interleave

scheduling-class

Syntax	scheduling-class <i>class-id</i> 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	<i>class-id</i> — Specifies the scheduling class.
Values	0 — 3

host-connectivity-verify

Syntax	host-connectivity-verify source-ip <i>ip-address</i> [source-mac <i>ieee-address</i>] [interval <i>interval</i>] [action { remove alarm }]
Context	config>service>vpls config>service>vpls>sap
Description	This command enables subscriber host connectivity verification on a given 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**
- source-ip** *ip-address* — Specify an unused IP address in the same network for generation of subscriber host connectivity verification packets.
 - source-mac** *ieee-address* — Specifies the source MAC address to be used for generation of subscriber host connectivity verification packets.
 - interval** *interval* — 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 — 6000
- Note that 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 given host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, etc.). 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.

Egress Multicast Group Commands

egress-multicast-group

Syntax	egress-multicast-group <i>egress-multicast-group-name</i> no egress-multicast-group <i>group-name</i>
Context	config>service
Description	This command creates an egress multicast group (EMG) context. An EMG is created as an object used to group VPLS SAPs that are allowed to participate in efficient multicast replication (EMR). EMR is a method to increase the performance of egress multipoint forwarding by sacrificing some destination-based features. Eliminating the requirement to perform unique features for each destination allows the egress forwarding plane to chain together multiple destinations into a batch replication process. In order to perform this batch replication function, similar characteristics are required on each SAP within the EMG.

Only SAPs defined on Ethernet access ports are allowed into an egress-multicast-group.

In order to understand the purpose of an egress-multicast-group, an understanding of the system's use of flooding lists is required. A flooding list is maintained at the egress forwarding plane to define a set of destinations to which a packet must be replicated. Multipoint services make use of flooding lists to enable forwarding a single packet to many destinations. Examples of multipoint services that use flooding lists are VPLS, IGMP snooping and IP multicast routing. Currently, the egress forwarding plane will only use efficient multicast replication for VPLS and IGMP snooping flooding lists.

In VPLS services, a unique flooding list is created for each VPLS context. The flooding list is used when a packet has a broadcast, multicast or unknown destination MAC address. From a system perspective, proper VPLS handling requires that a broadcast, multicast or unknown destined packet be sent to all destinations that are in the forwarding state. The ingress forwarding plane ensures the packet gets to all egress forwarding planes that include a destination in the VPLS context. It is the egress forwarding plane's job to replicate the packet to the subset of the destinations that are reached through its interfaces and each of these destinations are included in the VPLS context's flooding list.

For IGMP snooping, a unique flooding list is created for each IP multicast (s,g) record. This (s,g) record is associated with an ingress VPLS context and may be associated with VPLS destinations in the source VPLS instance or other VPLS instances (in the case of MVR). Again, the ingress forwarding plane ensures that an ingress IP multicast packet matching the (s,g) record gets to all egress forwarding planes that have a VPLS destination associated with the (s,g) record. The egress forwarding plane uses the flooding list owned by the (s,g) record to replicate the packet to all VPLS destinations in the flooding list. The IGMP Snooping function identifies which VPLS destinations should be associated with the (s,g) record.

With normal multicast replication, the egress forwarding plane examines which features are enabled for each destination. This includes ACL filtering, mirroring, encapsulation and queuing. The resources used to perform this per destination multicast processing are very expensive to the egress forwarding plane when high replication bandwidth is required. If destinations with similar egress functions can be grouped together, the egress forwarding plane can process them in a more efficient manner and maximize replication bandwidth.

The egress-multicast-group object is designed to allow the identification of SAPs with similar egress characteristics. When a SAP is successfully provisioned into an egress-multicast-group, the system is

ensured that it may be batched together with other SAPs in the same group at the egress forwarding plane for efficient multicast replication. A SAP that does not meet the common requirements is not allowed into the egress-multicast-group.

At the forwarding plane level, a VPLS flooding list is categorized into chainable and non-chainable destinations. Currently, the only chainable destinations are SAPs within an egress-multicast-group. The chainable destinations are further separated by egress-multicast-group association. Chains are then created following the rules below:

- A replication batch chain may only contain SAPs from the same egress-multicast-group
- A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

Further subcategories are created for an IGMP (s,g) flooding list. A Layer 2 (s,g) record is created in a specific VPLS instance (the instance the (s,g) flow ingresses). SAPs within that VPLS context that join the (s,g) record are considered native SAPs within the flooding list. SAPs that join the (s,g) flooding list through the multicast VPLS registration process (MVR) from another VPLS context using the **from-vpls** command are considered alien SAPs. The distinction between native and alien in the list is maintained to allow the forwarding plane to enforce or suspend split-horizon-group (SHG) squelching. When the source of the (s,g) matching packet is in the same SHG as a native SAP, the packet must not be replicated to that SAP. For a SAP in another VPLS context, the source SHG of the packet has no meaning and the forwarding plane must disregard SHG matching between the native source of the packet and the alien destination. Because the SHG squelch decision is done for the whole chain based on the first SAP in the chain, all SAPs in the chain must be all native or all alien SAPs. Chains for IGMP (s,g) flooding lists are created using the following rules:

1. A replication batch chain may only contain SAPs from the same egress-multicast-group.
2. A replication batch chain may only contain all alien or all native SAPs.
3. A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

When a packet associated with a flooding list is received by the egress forwarding plane, it processes the packet by evaluating each destination on the list sequentially in a replication context. If the current entry being processed in the list is a non-chained destination, the forwarding plane processes the packet for that destination and then moves on to process other packets currently in the forwarding plane before returning to process the next destination in the list. If the current entry being processed is a chained destination, the forwarding plane remains in the replication context until it has forwarded to each entry in that chain. Once the replication context finishes with the last entry in the chain, it moves on to process other packets waiting for egress processing before returning to the replication context. Processing continues in this manner until the packet has been forwarded to all destinations in the list.

Batch chain processing of a chain of SAPs improves replication efficiency by bypassing the functions that perform egress mirroring decisions on SAPs within the chain and making a single ACL filtering decision for the whole chain. Each destination in the chain may have a unique egress QoS policy and per destination queuing is still performed for each destination in the chain. Also, while each SAP in the chain must be on access ports with the same encap-type, if the encap-type is dot1q, each SAP may have a unique dot1q tag.

One caveat to each SAP having a unique egress QoS policy in the chain is that only the Dot1P marking decisions for the first SAP in the list is enforced. If the first SAP's QoS policy forwarding class action states that the packet should not be remarked, none of the replicated packets in the chain will have the dot1P bits remarked. If the first SAP's QoS policy forwarding class action states that the

Egress Multicast Group Commands

packet should be remarked with a specific dot1P value, all the replicated packets for the remaining SAPs in the chain will have the same dot1P marking.

While the system supports 32 egress multicast groups, a single group would usually suffice. An instance where multiple groups would be needed is when all the SAPs requiring efficient multicast replication cannot share the same common requirements. In this case, an egress multicast group would be created for each set of common requirements. An egress multicast group may contain SAPs from many different VPLS instances. It should be understood that an egress multicast group is not equivalent to an egress forwarding plane flooding list. An egress multicast group only identifies which SAPs may participate in efficient multicast replication. As stated above, entries in a flooding list are populated due to VPLS destination creation or IGMP snooping events.

The **no** form of the command removes a specific egress multicast group. Deleting an egress multicast group will only succeed when the group has no SAP members. To remove SAP members, use the **no multicast-group** *group-name* command under each SAP's egress context.

Parameters	<i>group-name</i> — Multiple egress multicast groups may be created on the system. Each must have a unique name. The egress-multicast-group-name is an ASCII string up to 16 characters in length and follows all the naming rules as other named policies in the system. The group's name is used throughout the system to uniquely identify the Egress Multicast Group and is used to provision a SAP into the group.
Default	None, each egress multicast group must be explicitly configured.
Values	Up to 32 egress multicast groups may be created on the system.

description

Syntax	description <i>description-string</i> no description
Context	config>service>egress-multicast-group
Description	This command defines an ASCII string associated with egress-multicast-group-name. The no form of the command removes an existing description string from egress-multicast-group.
Default	none
Parameters	<i>description-string</i> — The description command accepts a description-string parameter. The description-string parameter is an ASCII string of up to 80 characters in length. Only printable 127 bit ASCII characters are allowed. If the string contains spaces, the string must be specified with beginning and ending quotes.
Values	An ASCII string up to 80 characters in length.

dest-chain-limit

Syntax	dest-chain-limit <i>destinations per pass</i> no dest-chain-limit
Context	config>service>egress-multicast-group
Description	<p>This command defines the maximum length of an egress forwarding plane efficient multicast replication chain for an egress-multicast-group. Varying the maximum length of chains created for an egress multicast group has the effect of efficient multicast batched chain replication on other packets flowing through the egress forwarding plane. While replicating for the SAPs within a replication chain, other packets are waiting for the forwarding plane to finish. As the chain length increases, forwarding latency for the other waiting packets may increase. When the chain length decreases, a loss of efficiency in the replication process will be observed.</p> <p>The no form of the command restores the default value.</p>
Default	16
Parameters	<p><i>destinations per pass</i> — This parameter must be specified when executing the dest-chain-limit command. When executed, the command will use the number-of-destinations parameter to reorganize all efficient multicast SAP chains that contain members from the egress-multicast-group.</p> <p>The <i>destinations per pass</i> parameter can be modified at any time. Be aware that when changing the maximum chain length, the system will rebuild the chains according to the new limit. When this happens, it is possible that packets will not be replicated to a destination while it is being reorganized in the flooding list's chains. Only the chains associated with the egress-multicast-group context the command is executed in will be affected by changing the parameter.</p> <p>It is expected that the optimal replication chain length will be between 10 and 16. Since so many variables affect efficient multicast (i.e. ingress packet rate, number of chains, size of replicated packets), only proper testing in the environment that replication will be performed will identify the best dest-chain-limit value for each Egress Multicast Group.</p> <p>Setting the <i>destinations per pass</i> parameter to a value of 0 has the effect of removing from all egress forwarding planes all chains with members from the egress-multicast-group. Replication to each destination SAP from the group is performed using the normal method (non-efficient replication). The value 0 is not considered a normal value for dest-chain-limit and is provided for debugging purposes only. Setting the value to 0 is persistent between reboots of the system.</p> <p>Setting the <i>destinations per pass</i> parameter to a value of 1 has the effect of placing each egress-multicast-group member SAP into a chain with a single SAP. The value 1 is not considered a normal value for the dest-chain-limit and is provided for debugging purposes only. Setting the value to 1 is persistent between reboots of the system.</p>
Values	1 — 30

sap-common-requirements

Syntax	sap-common-requirements
Context	config>service>egress-multicast-group
Description	This command configures the common SAP parameter requirements. The SAP common requirements are used to evaluate each SAP for group membership. If a SAP does not meet the

Egress Multicast Group Commands

specified requirements, the SAP is not allowed into the egress-multicast-group. Once a SAP is a member of the group, attempting to change the parameters on the SAP will fail.

egress-filter

Syntax	egress-filter [ip <i>ip-filter-id</i>] egress-filter [ipv6 <i>ipv6-filter-id</i>] egress-filter [mac <i>mac-filter-id</i>] no egress-filter [ip <i>ip-filter-id</i>] [ipv6 <i>ipv6-filter-id</i>] [mac <i>mac-filter-id</i>]
Context	config>service>egress-multicast-group>sap-common-requirements
Description	<p>This command identifies the type of filter and actual filter ID that must be provisioned on the SAP prior to the SAP being made a member of the egress-multicast-group. If the SAP does not have the specified filter applied, the SAP cannot be provisioned into the group. It is important that the egress filter applied to each SAP within the egress-multicast-group be the same since the batch replication process on an efficient multicast replication chain will apply the first SAP's ACL decision to all other SAPs on the chain. Once the SAP is made a member of the egress-multicast-group, the SAP's egress filter cannot be changed on the SAP.</p> <p>Changing the egress-filter parameters within the sap-common-requirements node automatically changes the egress filter applied to each member SAP. If the filter cannot be changed on the SAP due to resource constraints, the modification will fail.</p> <p>The specified egress-filter does not contain an entry that is defined as an egress mirror-source. Once the filter is associated with the egress-multicast-group, attempting to define one of its entries as an egress mirror source will fail.</p> <p>The no form of the command removes the egress-filter removes the egress filter from each member SAP. The no egress-filter command specifies that an egress filter (IP, IPv6 or MAC) is not applied to a new member SAP within the egress-multicast-group.</p>
Default	no filter. The egress filter ID must be defined with the associated ip or mac keyword. If an egress-filter is not specified or the no egress-filter command is executed in the sap-common-requirements node, a new member SAP does not have an egress IP or MAC filter defined.
Parameters	<p>ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.</p> <p>Values 1 — 65535</p> <p>ipv6 <i>ipv6-filter-id</i> — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.</p> <p>Values 1 — 65535</p> <p>mac <i>mac-filter-id</i> — 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.</p> <p>Values 1 — 65535</p>

encap-type

Syntax	encap-type {dot1q null} no encap-type
Context	config>service>egress-multicast-group>sap-common-requirements
Description	<p>This command specifies the encapsulation type that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group. The config>port>ethernet>access>encap-type command is used to define the encapsulation type for the Ethernet port. The allowed encapsulation type values are dot1q and null. If the SAP does not exist on a port with the specified encap-type, it will not be allowed into the egress-multicast-group.</p> <p>If at least one SAP is currently a member of the efficient-multicast-group, the encap-type cannot be changed within the sap-common-requirements node. If the efficient-multicast-group does not contain any member SAPs, the encap-type may be changed at any time.</p> <p>There is no interaction between an efficient-multicast-group and the corresponding access ports associated with its members since all SAPs must be deleted from a port before its encap-type can be changed. When the SAPs are deleted from the port, they are also automatically deleted from the efficient-multicast-group.</p> <p>The no form of the command returns the egress-multicast-group required encapsulation type for SAPs to dot1q. If the current encap-type is set to null, the command cannot be executed when SAPs exist within the egress-multicast-group.</p>
Default	dot1q — For an egress-multicast-group. null — If member SAPs are on a null encapsulated access port.
Parameters	<p>null — The null keyword is mutually exclusive with the dot1q keyword. When the encap-type within the sap-common-requirements is specified to be null, the encapsulation type for the access ports associated with all SAPs within the egress-multicast-group must be set to null.</p> <p>dot1q — The dot1q keyword is mutually exclusive with the null keyword. When the encap-type within the sap-common-requirements is specified to be dot1q, the encapsulation type for the access ports associated with all SAPs within the egress-multicast-group must be set to dot1q.</p>

qinq-etype

Syntax	qinq-etype [0x0600..0xffff] no qinq-etype
Context	config>service>egress-multicast-group>sap-common-requirements
Description	This command specifies the EtherType used for QinQ encapsulation.
Default	no qinq-etype
	<i>etherType</i> — Defines the dot1q EtherType that must be associated with a SAP's access port when the encap-type is set to dot1q. Any valid EtherType may be specified.
Values	[0x0600 — 0xffff]: [1536 — 65535] in decimal or hex

qinq-fixed-tag-value

Syntax	qinq-fixed-tag-value <i>tag-value</i> no qinq-fixed-tag-value
Context	config>service>egress-multicast-group>sap-common-requirements
Description	This command configures the fixed tag value used for QinQ encapsulation.
Default	no qinq-fixed-tag-value
Parameters	<i>tag-value</i> — Specifies the provisioned common value of the fixed 802.1Q tag of all the QinQ SAP's in this egress multicast group. The value 0 is used to indicate that the actual value of the fixed tag will be defined implicitly by the corresponding tag of the first SAP added to this egress multicast group.
Values	0, 1 — 4094

dot1q-etype

Syntax	dot1q-etype [0x0600..0xffff] no dot1q-etype
Context	config>service>egress-multicast-group>sap-common-requirements
Description	This command specifies the dot1q EtherType that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group. The config>port>ethernet>access>dot1q-etype command is used to define the EtherType used when encapsulating a packet with a dot1q tag on the Ethernet port. Any valid EtherType is allowed on the port. If the current encap-type for the egress-multicast-group is set to null, the dot1q-etype EtherType is ignored when evaluating SAP membership in the group. If the encap-type is set to dot1q (the default), a member SAP's access port must be configured with the same dot1q-etype EtherType as the egress-multicast-group. If at least one SAP is currently a member of the efficient-multicast-group, the dot1q-etype value cannot be changed within the sap-common-requirements node. If the efficient-multicast-group does not contain any member SAPs, the dot1q-etype value may be changed at any time. If an access port currently has SAPs associated with it that are defined within an egress-multicast-group and the port is currently set to encap-type dot1q, the dot1q-etype value defined on the port cannot be changed. The no form of the command returns the egress-multicast-group dot1q EtherType to the default value of 0x8100. If the current encap-type is set to a value other than 0x8100, the command cannot be executed when SAPs exist within the egress-multicast-group.
Default	The default dot1q-etype is 0x8100 for an egress-multicast-group.
Parameters	<i>ethertype</i> — Defines the dot1q EtherType that must be associated with a SAP's access port when the encap-type is set to dot1q. Any valid EtherType may be specified.
Values	0x0600 — 0xffff 1536 — 65535 in decimal or hex
Default	0x8100

BGP Auto-Discovery Commands

bgp

Syntax	bgp
Context	config>service>vpls
Description	This command enables 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 enables the context to configure the BGP-VPLS parameters and addressing.

max-ve-id

Syntax	max-ve-id <i>value</i> no max-ve-id
Context	config>service>vpls>bgp-vpls
Description	<p>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.</p> <p>Also upon reception of a higher VE-id in an NLRI imported in this VPLS instance (RT = configured import RT) the following action must be taken:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <p>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.</p> <p>The actions taken for other max-ve-id values are described below:</p> <ul style="list-style-type: none"> • max-ve-id value higher than all VE-IDs (local and received) is allowed and there are no effects.

BGP Auto-Discovery Commands

- 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.
 - Note that 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	<i>value</i> — Specifies the allowed range of [1-value] for the VE-id. The configured value must be bigger than the existing VE-ids.
Values	1-65535

ve-name

Syntax	ve-name <i>name</i> 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. Command “no shutdown” cannot be used if there is no ve-name configured.
Default	no ve-name
Parameters	<i>name</i> — A character string identifying the VPLS Edge instance.
Values	32 ASCII chars max

ve-id

Syntax	ve-id <i>ve-id-value</i> 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. Command “no shutdown” cannot be used if there is no ve-id configured.

Default no ve-id

Parameters *value* — A two bytes 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 1-65535

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

BGP Auto-Discovery Commands

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 — 2147483647

split-horizon-group *group-name* — The specified group-name overrides the split horizon group template settings.

import-rt *ext-comm* — Specify 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 — 65535
2byte-asnumber	0 — 65535
ext-comm-val	0 — 4294967295
4byte-asnumber	0 — 4294967295

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>service>vpls>bgp>pw-template-bindin

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 given 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 bfd-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* — 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

BGP Auto-Discovery Commands

- 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 & no RT is configured under BGP node - RT = 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* — Specify communities allowed to be sent to remote PE neighbors.
- import** *ext-community* — Specify 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-addr | a.b.c.d |
| | comm-val | 0 — 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 enables 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

BGP Auto-Discovery Commands

Syntax	route-distinguisher <i>rd</i> route-distinguisher auto-rd no route-distinguisher												
Context	config>service>vpls>bgp												
Description	<p>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.</p> <p>If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:</p> <ul style="list-style-type: none">• if BGP AD VPLS-id is configured & 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 <p>Values and format (6 bytes, other 2 bytes of type will be automatically generated)</p> <p>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.</p>												
Parameters	<p><i>ip-addr:comm-val</i> — Specifies the IP address.</p> <table><tr><td>Values</td><td>ip-addr</td><td>a.b.c.d</td></tr><tr><td></td><td>comm-val</td><td>0 — 65535</td></tr></table> <p><i>as-number:ext-comm-val</i> — Specifies the AS number.</p> <table><tr><td>Values</td><td>as-number</td><td>1 — 65535</td></tr><tr><td></td><td>ext-comm-val</td><td>0 — 4294967295</td></tr></table> <p>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.</p>	Values	ip-addr	a.b.c.d		comm-val	0 — 65535	Values	as-number	1 — 65535		ext-comm-val	0 — 4294967295
Values	ip-addr	a.b.c.d											
	comm-val	0 — 65535											
Values	as-number	1 — 65535											
	ext-comm-val	0 — 4294967295											

vsi-import

Syntax	vsi-import <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no vsi-import
Context	config>service>vpls>bgp-ad>vsi-id config>service>vpls>bgp
Description	<p>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.</p> <p>The policy name list is handled by the SNMP agent as a single entity.</p>

bgp-evpn

Syntax	[no] bgp-evpn
Context	config>service>vpls
Description	This command enables 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 learnt 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 <i>num-moves</i> window <i>minutes</i>
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 given 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..10 minutes
	Default 3 minutes

retry

BGP Auto-Discovery Commands

Syntax	retry <i>minutes</i> no retry
Context	config>service>vpls>bgp-evpn>mac-duplication
Description	<p>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.</p> <p>If no retry is configured, this implies that, once 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.</p>
Default	9 minutes
Parameters	<i>minutes</i> — I.
	Values 2 — 60 minutes

unknown-mac-route

Syntax	[no] unknown-mac-route
Context	config>service>vpls>bgp-evpn
Description	<p>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 learnt 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. Note that, although the 7x50 can be configured to generate and advertise the unknown-mac-route, the 7x50 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.</p>
Default	no unknown-mac-route

vxlan

Syntax	vxlan
Context	config>service>vpls>bgp-evpn
Description	<p>This command enables the context to configure the VXLAN parameters when BGP EVPN is used as the control plane.</p>

shutdown

Syntax	<code>[no] shutdown</code>
Context	<code>config>service>vpls>bgp-evpn.vxlan</code>
Description	This command enables/disables the automatic creation of VXLAN auto-bindings by BGP-EVPN.
Default	shutdown

VPLS Show Commands

egress-label

Syntax	egress-label <i>egress-label1</i> [<i>egress-label2</i>]
Context	show>service
Description	<p>This command displays service information using the range of egress labels.</p> <p>If only the mandatory <i>egress-label1</i> parameter is specified, only services using the specified label are displayed.</p> <p>If both <i>egress-label1</i> and <i>egress-label2</i> parameters are specified, the services using the range of labels X where <i>egress-label1</i> <= X <= <i>egress-label2</i> are displayed.</p> <p>Use the show router ldp bindings command to display dynamic labels.</p>
Parameters	<p><i>egress-label1</i> — The starting egress label value for which to display services using the label range. If only <i>egress-label1</i> is specified, services only using <i>egress-label1</i> are displayed.</p> <p>Values 0, 2049 — 131071</p> <p><i>egress-label2</i> — The ending egress label value for which to display services using the label range.</p> <p>Default The <i>egress-label1</i> value.</p> <p>Values 2049 — 131071</p>

fdb-info

Syntax	fdb-info
Context	show>service
Description	Displays global FDB usage information.
Output	Show FDB-Info Command Output — The following table describes show FDB-Info command output.

Label	Description
Service ID	The value that identifies a service.
Mac Move	Indicates the administrative state of the MAC movement feature associated with the service.

Label	Description (Continued)
Mac Move Rate	The maximum rate at which MAC's can be re-learned in this TLS 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 MAC's. The rate is computed as the maximum number of re-learns allowed in a 5 second interval. The default rate of 10 re-learns per second corresponds to 50 re-learns in a 5 second period.
Mac Move Timeout	Indicates 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.
Table Size	The maximum number of learned and static entries allowed in the FDB.
Total Count	The current number of entries (both learned and static) in the FDB of this service.
Learned Count	The current number of learned entries in the FDB of this service.
Static Count	The current number of static entries in the FDB of this service.
Local Age	The seconds used to age out FDB entries learned on local SAPs.
High WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is raised by the agent.
Low WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is cleared by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled in this service.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded in this service.
MAC Aging	Specifies whether the MAC aging process is enabled in this service.
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of MAC's supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.
Total Service FDB	The current number of service FDBs configured on this node.
Total FDB Configured Size	The sum of configured FDBs.
Total FDB Entries In Use	The total number of entries (both learned and static) in use.

Sample Output

fdb-mac

- Syntax** `fdb-mac ieee-address [expiry]`
- Context** `show>service`
- Description** This command displays the FDB entry for a given MAC address.
- Parameters** *ieee-address* — The 48-bit MAC address for which to display the FDB entry 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** **Show FDB-MAC Command Output** — The following table describes the 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. P — Indicates the MAC is protected by the MAC protection feature.

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
-----
```

Show, Clear, Debug Commands

```
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 shows whether restrict-protected-src or restrict-unprotected-dst are enabled on SDPs.

```
*A:ian1# show service id 1 sdp 1:1 detail
```

```
=====
Service Destination Point (Sdp Id : 1:1) Details
=====
-----
Sdp Id 1:1  -(1.1.1.2)
-----
...
Flags          : RxProtSrcMac
...
Restr MacProt Src : Enabled                Restr MacUnpr Dst : Disabled
```

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 which to display services using the label range. If only *start-label* is specified, services only using *start-label* are displayed.
- Values** 0, 2048 — 131071
- end-label* — The ending ingress label value for which to display services using the label range.
- Default** The *start-label* value.
- Values** 2049 — 131071

Output **Show Service Ingress-Label** — The following table describes show service ingress-label output fields.

Label	Description
Svc ID	The service identifier.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is spoke.
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

```

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.
filter *filter-id* — The ingress or egress filter policy ID for which to display matching SAPs.

Values 1 — 65535

dyn-script — Displays dynamic service SAPs information.

sap-id — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 1271](#) for command syntax.

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 **Show Service SAP** — The following table describes 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.
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.
Etree 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.

Sample Output

```
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
=====
```

sdp

Syntax **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 or 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 — 17407

far-end ip-addr — Displays only SDPs matching with the specified system IP address of the far-end destination or router 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 **Show Service SDP** — The following table describes show service-id SDP output fields.

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke
VC Type	Displays the VC type, ether or vlan.

Label	Description (Continued)
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 MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: 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.
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.

Label	Description (Continued)
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.

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
=====
```

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 — 17407
vc-id — The virtual circuit identifier.
Values 1 — 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 bindings 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** **Show Service SDP Using** — The following table describes service-using output fields.

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.
Type	Specifies the type of SDP: Spoke
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.
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.

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
-----
2          300:2      Spok 10.0.0.13     Up        131070  131070
-----
Number of SDPs :
-----
*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
=====
```

service-using

Syntax **service-using** [**epipe**] [**vpls**] [**mirror**] [**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.
vpls — Displays matching VPLS instances.
mirror — Displays matching mirror services.
customer *customer-id* — Displays services only associated with the specified customer ID.

Default Services associated with a customer.

Values 1 — 2147483647

etree — Specifies matching of all VPLS services configured as E-Tree.

Output **Show Service Service-Using** — The following table describes show service service-using output 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.

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#

```

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.

```

id

Syntax	id <i>service-id</i>
Context	show>service
Description	This command displays information for a particular service-id.
Parameters	<i>service-id</i> — The unique service identification number that identifies the service in the service domain. Values service-id: 1 — 214748364 svc-name: A string up to 64 characters in length. all — Display detailed information about the service. base — Display basic service information. endpoint — Display service endpoint information. fdb — Display FDB entries. labels — Display labels being used by this service. sap — Display SAPs associated to the service. sdp — Display SDPs associated with the service. stp — Display STP information.

all

- Syntax** all
- Context** show>service>id
- Description** This command displays detailed information for all aspects of the service.
- Output** **Show service ID all output** — The following table describes the command output fields.

Label	Description
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	Specifies 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.
SAP Count	The number of SAPs specified for this service.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Last Changed	The date and time of the most recent change to this customer.
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.
Ingress qos-pol- icy	The SAP ingress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.

Label	Description (Continued)
High priority offered	The packets or octets count of the high priority traffic for the SAP.
Managed by Service	Specifies the service-id of 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 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.
Prune state	Specifies the STP state inherited from the management VPLS.

Sample Output

```
*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 Qinq-Vc : Disabled
Force Vlan-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

```

Show, Clear, Debug Commands

```
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
```

VPLS Show Commands

```

Port Role           : N/A
Port Number         : 0
Port Path Cost      : 10
Admin Edge          : Disabled
Link Type           : Pt-pt
Root Guard          : Disabled
Last BPDU from      : N/A
Designated Bridge   : N/A

Port State          : Forwarding
Port Priority        : 128
Auto Edge           : Enabled
Oper Edge           : N/A
BPDU Encap          : Dot1d
Active Protocol     : N/A

Designated Port Id: 0

Fwd Transitions     : 0
Cfg BPDUs rcvd      : 0
TCN BPDUs rcvd      : 0
TC bit BPDUs rcvd   : 0
RST BPDUs rcvd      : 0

Bad BPDUs rcvd      : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 0
TC bit BPDUs tx      : 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
QinQ Dot1p          : Default
Description           : (Not Specified)
Admin 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
Split Horizon Group: (Not Specified)

Encap                 : qinq
Oper State            : Up
QinQ Ethertype        : 0x8100

Etree Root Leaf Tag: Disabled
Etree Leaf AC         : Disabled
Max Nbr of MAC Addr: No Limit
Learned MAC Addr      : 0
OAM MAC Addr          : 0
Host MAC Addr         : 0
SPB MAC Addr          : 0
BGP EVPN Addr         : 0
Admin MTU              : 1522
Ingr IP Fltr-Id       : n/a
Ingr Mac Fltr-Id      : n/a
Ingr IPv6 Fltr-Id     : n/a
tod-suite              : None

Etree Leaf Tag        : 0
Total MAC Addr        : 0
Static MAC Addr       : 0
DHCP MAC Addr         : 0
Intf MAC Addr         : 0
Cond MAC Addr         : 0
EVPN Static Addr     : 0
Oper MTU              : 1522
Egr IP Fltr-Id        : n/a
Egr Mac Fltr-Id       : n/a
Egr IPv6 Fltr-Id     : n/a
qinq-pbit-marking    : both
Egr Agg Rate Limit: max
Limit Unused BW       : Disabled
Host Conn Verify      : Disabled
Discard Unkwn Srce: Disabled
Mac Pinning           : Disabled

Q Frame-Based Acct   : Disabled
ARP Reply Agent       : Disabled
Mac Learning          : Enabled
Mac Aging             : Enabled
BPDU Translation     : Disabled
L2PT Termination     : Disabled
Vlan-translation     : None

Acct. Pol             : None
Collect Stats         : Disabled

```

Show, Clear, Debug Commands

Anti Spoofing : None
Avl Static Hosts : 0
Calling-Station-Id : n/a

Dynamic Hosts : Enabled
Tot Static Hosts : 0

Application Profile: None
Transit Policy : None

Oper Group : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)
Cflowd : Disabled
MCAC Policy Name :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Use LAG port weight: no
Restr MacProt Src : Disabled
Auto Learn Mac Prot: Disabled
Time to RetryReset : never
Mac Move : Blockable
Egr MCast Grp :
Auth Policy : None

Monitor Oper Grp : (none)
MCAC Const Adm St : Enable
MCAC Max Mand BW : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited
Restr MacUnpr Dst : Disabled
RestProtSrcMacAct : Disable
Retries Left : 3
Blockable Level : Tertiary

----- ETH-CFM SAP specifics -----

Tunnel Faults : accept
MC Prop-Hold-Timer : n/a
Squelch Levels : None

AIS : Disabled
V-MEP Filtering : Disabled

----- Stp Service Access Point specifics -----

Stp Admin State : Up
Core Connectivity : Down
Port Role : N/A
Port Number : N/A
Port Path Cost : 10
Admin Edge : Disabled
Link Type : Pt-pt
Root Guard : Disabled
Last BPDUs from : N/A
CIST Desig Bridge : N/A

Stp Oper State : Down
Port State : Forwarding
Port Priority : 128
Auto Edge : Enabled
Oper Edge : N/A
BPDU Encap : Dot1d
Active Protocol : N/A
Designated Port : N/A

Forward transitions: 0
Cfg BPDUs rcvd : 0
TCN BPDUs rcvd : 0
TC bit BPDUs rcvd : 0
RST BPDUs rcvd : 0
MST BPDUs rcvd : 0

Bad BPDUs rcvd : 0
Cfg BPDUs tx : 0
TCN BPDUs tx : 0
TC bit BPDUs tx : 0
RST BPDUs tx : 0
MST BPDUs tx : 0

----- ARP host -----

Admin State : outOfService
Host Limit : 1

Min Auth Interval : 15 minutes

----- QOS -----

Ingress qos-policy : 2
Ingress FP QGrp : (none)

Egress qos-policy : 2
Egress Port QGrp : (none)

VPLS Show Commands

```

Ing FP QGrp Inst      : (none)
Shared Q plcy        : n/a
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
DHCP Snooping        : Down
Lease Populate       : 0
Action               : Keep

```

```

Proxy Admin State    : Down
Proxy Lease Time     : N/A
Emul. Server Addr    : Not Configured

```

Subscriber Management

```

Admin State          : Down
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
MAC DA Hashing       : False

```

```

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

Show, Clear, Debug Commands

```
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
```

VPLS Show Commands

```

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 Id	Oper Stat	MRtr Port	Pim Port	Send Qrys	Max Grps	Max Srcs	Max Grp	MVR From-VPLS	Num Grps
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 Id	Oper State	MRtr Port	Send Queries	Max Num Groups	MVR From-VPLS	Num 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/Provided	Arp Reply Agent	Info Option	Admin 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

Show, Clear, Debug Commands

```
-----
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          Precedence         : 4
Endpoint         : N/A
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

```

Show, Clear, Debug Commands

```
Peer Vccv CC Bits : mplsRouterAlertLabel

Application Profile: None
Transit Policy      : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr   : 0
OAM MAC Addr       : 0
Host MAC Addr      : 0
SPB MAC Addr       : 0
BGP EVPN Addr      : 0

Total MAC Addr     : 0
Static MAC Addr    : 0
DHCP MAC Addr      : 0
Intf MAC Addr      : 0
Cond MAC Addr      : 0
EVPN Static Addr   : 0

MAC Learning       : Enabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False
Oper Group         : (none)
Rest Prot Src Mac  : Disabled
Auto Learn Mac Prot : Disabled
Ing. Vlan Trans.   : 0

Discard Unkwn Srce: Disabled

Block On Mesh Fail: False
Monitor Oper Grp   : (none)

RestProtSrcMacAct : Disable

Ingress Qos Policy : (none)
Ingress FP QGrp    : (none)
Ing FP QGrp Inst   : (none)

Egress Qos Policy  : (none)
Egress Port QGrp   : (none)
Egr Port QGrp Inst : (none)

KeepAlive Information :
Admin State         : Disabled
Hello Time         : 10
Max Drop Count     : 3

Oper State          : Disabled
Hello Msg Len      : 0
Hold Down Time     : 10

Statistics          :
I. Fwd. Pkts.      : 0
I. Fwd. Octs.      : 0
E. Fwd. Pkts.      : 0
I. Dro. Pkts.      : 0
I. Dro. Octs.      : 0
E. Fwd. Octets     : 0

-----
Control Channel Status
-----
PW Status           : disabled
Peer Status Expire  : false
Request Timer       : <none>
Acknowledgement     : false

Refresh Timer       : <none>

MCAC Policy Name    :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0

MCAC Max Mand BW   : no limit
MCAC Avail Mand BW: unlimited
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)
```

Show, Clear, Debug Commands

```
Etree Root Leaf Tag: Disabled
Etree Leaf AC      : Disabled
Max Nbr of MAC Addr: No Limit
Learned MAC Addr  : 0
OAM MAC Addr      : 0
Host MAC Addr     : 0
SPB MAC Addr      : 0
BGP EVPN Addr     : 0
Admin MTU         : 1522
Ingr IP Fltr-Id   : n/a
Ingr Mac Fltr-Id  : n/a
Ingr IPv6 Fltr-Id: n/a
tod-suite         : None

Q Frame-Based Acct : Disabled
ARP Reply Agent    : Disabled
SHCV pol IPv4      : None
Mac Learning       : Enabled
Mac Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination  : Disabled
Vlan-translation  : None

Acct. Pol         : None

Anti Spoofing     : None
Avl Static Hosts  : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy    : None

Oper Group        : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)
Cflowd           : Disabled
MCAC Policy Name  :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Use LAG port weight: no
Restr MacProt Src : Disabled
Auto Learn Mac Prot: Disabled
Time to RetryReset : never
Mac Move          : Blockable
Egr MCast Grp    :
Auth Policy       : None

Etree Leaf Tag    : 0
Total MAC Addr    : 0
Static MAC Addr   : 0
DHCP MAC Addr     : 0
Intf MAC Addr     : 0
Cond MAC Addr     : 0
EVPN Static Addr  : 0
Oper MTU          : 1522
Egr IP Fltr-Id   : n/a
Egr Mac Fltr-Id  : n/a
Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max
Limit Unused BW   : Disabled
Host Conn Verify  : Disabled

Discard Unkwn Srce: Disabled
Mac Pinning       : Disabled

Collect Stats     : Disabled

Dynamic Hosts     : Enabled
Tot Static Hosts  : 0

Monitor Oper Grp  : (none)

MCAC Const Adm St : Enable
MCAC Max Mand BW  : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited

Restr MacUnpr Dst : Disabled
RestProtSrcMacAct : Disable
Retries Left      : 3
Blockable Level   : Tertiary

-----
ETH-CFM SAP specifics
-----
Tunnel Faults      : accept
MC Prop-Hold-Timer : n/a
Squelch Levels     : None

AIS                : Disabled
V-MEP Filtering    : Disabled

-----
Stp Service Access Point specifics
-----
Stp Admin State    : Up
Core Connectivity  : Down
Port Role          : N/A

Stp Oper State     : Down
Port State         : Forwarding
```

VPLS Show Commands

```

Port Number          : N/A
Port Path Cost       : 10
Admin Edge           : Disabled
Link Type            : Pt-pt
Root Guard           : Disabled
Last BPDU from       : N/A
CIST Desig Bridge    : N/A

Port Priority        : 128
Auto Edge           : Enabled
Oper Edge            : N/A
BPDU Encap           : Dot1d
Active Protocol      : N/A
Designated Port     : N/A

Forward transitions: 0
Cfg BPDUs rcvd      : 0
TCN BPDUs rcvd      : 0
TC bit BPDUs rcvd   : 0
RST BPDUs rcvd      : 0
MST BPDUs rcvd      : 0

Bad BPDUs rcvd      : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx        : 0
TC bit BPDUs tx     : 0
RST BPDUs tx        : 0
MST BPDUs tx        : 0
-----
ARP host
-----
Admin State          : outOfService
Host Limit           : 1
Min Auth Interval   : 15 minutes
-----
QOS
-----
Ingress qos-policy  : 2
Ingress FP QGrp     : (none)
Ing FP QGrp Inst    : (none)
Shared Q plcy       : n/a
I. Sched Pol        : (Not Specified)
E. Sched Pol        : (Not Specified)
I. Policer Ctl Pol  : (Not Specified)
E. Policer Ctl Pol  : (Not Specified)

Egress qos-policy   : 2
Egress Port QGrp    : (none)
Egr Port QGrp Inst : (none)
Multipoint shared   : Disabled
-----
DHCP
-----
Description          : (Not Specified)
Admin State          : Down
DHCP Snooping        : Down
Lease Populate       : 0
Action               : Keep

Proxy Admin State    : Down
Proxy Lease Time     : N/A
Emul. Server Addr    : Not Configured
-----
Subscriber Management
-----
Admin State          : Down
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

MAC DA Hashing      : False

Subscriber Limit     : 1
Single-Sub-Parameters
  Prof Traffic Only  : False
  Non-Sub-Traffic    : N/A

Static host management
MAC learn options    : N/A
-----

```

Show, Clear, Debug Commands

```
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
```

VPLS Show Commands

```

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 Id	Oper Stat	MRtr Port	Pim Port	Send Qrys	Max Grps	Max Srcs	Max Grp	MVR From-VPLS	Num Grps
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 Id	Oper State	MRtr Port	Send Queries	Max Num Groups	MVR From-VPLS	Num 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

Show, Clear, Debug Commands

```

-----
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.

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* — All IP addresses.
mac 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 sap-id — Displays SAP information for the specified SAP ID.

interface — Specifies matching service ARP entries associated with the IP interface.

ip-address — The IP address of the interface for which to display matching ARP entries.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display matching ARPs.

Output **Show Service-ID ARP** — The following table describes show service-id ARP output 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.

base

- Syntax** **base**
- 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.
- Output** **Show Service-ID Base** — The following table describes 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	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

fdb

- Syntax** `fdb [sap sap-id [expiry]] | [mac ieee-address [expiry]] | [detail] [expiry]`
- Context**
`show>service>id`
`show>service>fdb-mac`
- Description** This command displays FDB entries for a given MAC address.
- Parameters**
sap *sap-id* — Specifies the physical port identifier portion of the SAP. See [Common CLI Command Descriptions on page 1271](#) for command syntax.
detail — Displays detailed information.
expiry — Displays time until MAC is aged out.
- Show FDB Information** — The following table describes service FDB output 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 MAC's 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 reenabling the SAP/SDP.
Table Size	Specifies the maximum number of learned and static entries allowed in the FDB of this service.
Total Count	Displays the total number of learned entries in the FDB of this service.

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-learned Count	Displays the current number of OAM entries in the FDB of 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.
Mac Aging	Indicates whether the MAC aging process is enabled.
Relearn Only	Displays, that when enabled, either the FDB table of this service is full, or that the maximum system-wide number of MA's supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.
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. L — Learned - Dynamic entries created by the learning process. OAM — Entries created by the OAM process. Static — Statically configured.
Last Change	Indicates the time of the most recent state changes.

Sample Output

isid-using

Syntax	isid-using [<i>ISID</i>]
Context	show>service
Description	This command displays services using an ISID.
Parameters	<i>ISID</i> — Specifies a 24 bit (0..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 — 16777215

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	Show Service-ID Labels — The following table describes show service-id labels output fields:

Label	Description
Svc Id	The service identifier.
Sdp Id	The SDP identifier.

Show, Clear, Debug Commands

Label	Description
Type	Indicates whether the SDP is spoke.
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.

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#
```

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	Show L2PT Fields — The following table describes show L2PT output fields:

Label	Description
Service id	Displays the 24 bit (0..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 PDU's 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.

```
A:ALA-48>show>service>id# l2pt
```

```
=====
```

Show, Clear, Debug Commands

```
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-termination          Admin Bpdu-translation  Oper Bpdu-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#

```

mac-move

Syntax	mac-move
Context	show>service>id
Description	This command displays MAC move related information about the service.

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

```

Show, Clear, Debug Commands

```
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

```
*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
```

mrollers

Syntax **mrollers [detail]**

Context show>service>id>mld-snooping

Description This command displays all multicast routers.

provider-tunnel

Syntax **provider-tunnel**

Description This command displays the service provider tunnel information.

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
-----
-----

```

provider-tunnel

Syntax provider-tunnel

Context show>service>id

Description This command displays 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-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>id

Description This command displays the proxy-ARP entries existing for a particular service. A 7x50 receiving an ARP request from a SAP or SDP-binding will perform a lookup in the proxy-arp table for the service. If the 7x50 finds a match, it will reply to the ARP and will not let the ARP be flooded in the VPLS service. If the 7x50 does not find a match, the ARP will be flooded within the service. The command allows for a specific IP addresses to be shown.

The "detail" modifier allows the user to display all the entries. An individual ip-address entry can also be shown.

Output 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
=====
```

Show, Clear, Debug Commands

A:PE71(1)#

proxy-nd

- Syntax** `proxy-nd [ip-address ip-address] [detail]`
- Context** `show>service>id`
- Description** This command displays the information about the proxy-nd settings configured in a given service. The "detail" modifier allows the user to display all the entries. An individual ip-address entry can also be shown.

Output **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)#

```

vxlan

- Syntax** `vxlan`
- Context** `show>service>id`
`show>service`
- Description** This command displays the VXLAN bindings auto-created in a given 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** **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
-----

```

Show, Clear, Debug Commands

60

60

Up

sap

- Syntax** `sap sap-id [filter]`
- 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 sap-id` — The ID that displays SAPs for the service in the `slot/mdalport[.channel]` form. See [Common CLI Command Descriptions on page 1271](#) for command syntax.
- `detail` — Displays detailed information for the SAP.
- Output** **Show Service-ID SAP** — The following table describes show service 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 Ether type 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, L2OperDown, RelearnLimitExceeded, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, 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.
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.

Label	Description (Continued)
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.
Queueing 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.
Queueing Stats (Egress QoS Policy)	
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.

Label	Description (Continued)
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.

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
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)
```

Show, Clear, Debug Commands

```
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
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
Designated Port   : N/A
CIST Desig Bridge : N/A

Forward transitions: 0
Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 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    : (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       : test2
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

 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

Show, Clear, Debug Commands

```

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#

```

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 — 17407
far-end *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** **Show Service-ID SDP** — The following table describes show service-id SDP output fields.

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is spoke.
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.

Label	Description (Continued)
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.
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.

Label	Description (Continued)
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 signalling 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

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
```

VPLS Show Commands

```

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

```

Show, Clear, Debug Commands

```
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
=====

```

Show, Clear, Debug Commands

```
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. Octets : 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

```

Show, Clear, Debug Commands

```
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)

```

Show, Clear, Debug Commands

```
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
```

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.

Values 32 chars max

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
-----
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
```

```
-----  
-----  
Associations  
-----  
SAP                1/1/3:1  
SDP                108:1  
SDP                109:1  
-----  
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 — 4094
- Output** **Show Service-ID STP Output** — The following table describes show service-id STP output fields:

Label	Description
Bridge-id	Specifies the MAC address used to identify this bridge in the network.
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.
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.

Label	Description (Continued)
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.
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 towards the spanning tree root which include this port.
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.

Sample Output

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** **Show All Service-ID** — The following table describes the show all service-id command output 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.
Num Groups	Specifies the actual number of multicast groups that can be joined on this SAP.

mfib

- Syntax** **mfib** [**brief** | **statistics**] [**ip** | **mac**] **brief**
mfib [**group** *grp-address* | *]
- Context** show>service>id
- Description** This command displays the multicast FIB on the VPLS service.
- Parameters** **brief** — Displays a brief output.
group *grp grp-address* — Displays the multicast FIB for a specific multicast group address.
- Output** **Show Output** — The following table describes the command output fields:

Label	Description
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 be forwarded/blocked. Local means that the multicast stream will be forwarded/blocked to a SAP or SDP local to the service.

The following example shows the MFIB for an EVPN-VXLAN service*A:PE1# show service id 1 mfib

```
=====
Multicast FIB, Service 1
=====
Source Address  Group Address      Sap/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
=====
```

mrouters

Syntax	mrouters [detail]
Context	show>service>id>igmp-snooping
Description	This command displays all multicast routers.
Parameters	detail — Displays detailed information.

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	<p>group grp-ip-address — Displays the IGMP snooping port database for a specific multicast group address.</p> <p>sap sap-id — Displays the IGMP snooping port database for a specific SAP. See Common CLI Command Descriptions on page 1271 for command syntax.</p> <p>sdp 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.</p> <p>Values 1 — 17407</p> <p>vc-id — The virtual circuit ID on the SDP ID for which to display information.</p> <p>Default For mesh SDPs only, all VC IDs.</p> <p>Values 1 — 4294967295</p>
Output	Show Output — The following table describes the show output fields:

Label	Description
Group Address	The IP multicast group address for which this entry contains information.

Label	Description
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. In exclude' mode, reception of packets sent to the given 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.
Compatibility mode	Specifies the IGMP mode. This is 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.
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** `group` *grp-ip-address* — Displays the IGMP snooping proxy reporting database for a specific multicast group address.
- Output** **Show Output** — The following table describes the show output 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. In the “exclude” mode, reception of packets sent to the given 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** **Show Output** — The following table describes the show output fields:

Label	Description
SAP Id	Specifies the SAP ID of the service.

Label	Description (Continued)
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. See [Common CLI Command Descriptions on page 1271](#) for command syntax.

`sdp sdp-id` — Displays the IGMP snooping source groups for a specific spoke or mesh SDP.

Values 1 — 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 — 4294967295

Output **Show Output** — The following table describes the show output 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.

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#
```

statistics

Syntax	statistics [sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i> vxlan vtep <i>ip-address vni vni</i>]
Context	show>service>id>igmp-snooping
Description	This command displays IGMP snooping statistics for the VPLS service.
Parameters	<p>sap <i>sap-id</i> — Displays IGMP snooping statistics for a specific SAP. See Common CLI Command Descriptions on page 1271 for command syntax.</p> <p>sdp <i>sdp-id</i> — Displays the IGMP snooping statistics for a specific spoke or mesh SDP.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID for which to display information.</p> <p>Default For mesh SDPs only, all VC IDs.</p> <p>Values 1 — 4294967295</p> <p>Values 1 — 16777215</p>

bgp-evpn

Syntax	bgp-evpn
Context	show>service>id
Description	This command displays the bgp-evpn configured parameters for a given 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	<p>Sample Output</p> <pre>*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 ----- =====</pre>

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

Parameters **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
=====
```

VPLS Clear Commands

id

Syntax	id <i>service-id</i>
Context	clear>service clear>service>statistics
Description	This command clears commands for a specific service.
Parameters	<i>service-id</i> — The ID that uniquely identifies a service.
Values	service-id: 1 — 214748364 svc-name: A string up to 64 characters in length.

statistics

Syntax	statistics
Context	clear>service>stats
Description	This command clears session statistics for this service.

fdb

Syntax	fdb { all mac <i>ieee-address</i> sap <i>sap-id</i>] mesh-sdp <i>sdp-id[:vc-id]</i> spoke-sdp <i>sdp-id:vc-id</i> }
Context	clear>service>id
Description	This command clears FDB entries for the service.
Parameters	all — Clears all FDB entries. mac <i>ieee-address</i> — Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers. <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1271 for command syntax. 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.

Show, Clear, Debug Commands

sdp-id — The SDP ID for which to clear associated FDB entries.

vc-id — The virtual circuit ID on the SDP ID for which to clear associated FDB entries.

Values	sdp-id[:vc-id]	<i>sdp-id</i>	1 — 17407
		<i>vc-id</i>	1 — 4294967295
	sdp-id:vc-id	<i>sdp-id</i>	1 — 17407
		<i>vc-id</i>	1 — 4294967295

mesh-sdp

Syntax	mesh-sdp <i>sdp-id[:vc-id]</i> ingress-vc-label
Context	clear>service>id
Description	This command clears and resets the mesh SDP bindings for the service.
Parameters	<i>sdp-id</i> — The mesh SDP ID to be reset. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset. Default All VC IDs on the SDP ID. Values 1 — 4294967295

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* — The spoke SDP ID to be reset.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 — 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.

spoke-sdp

Syntax	spoke-sdp <i>sdp-id</i> [: <i>vc-id</i>] { all counters stp l2pt mrp }
Context	clear>service>statistics>id
Description	This command clears statistics for the spoke SDP bound to the service.
Parameters	<i>sdp-id</i> — The spoke SDP ID for which to clear statistics. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset. Values 1 — 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.

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 <i>sap-id</i> }
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
Parameters	all — Clears all detected protocol statistics. <i>sap-id</i> — Clears the specified lease state SAP information. See Common CLI Command Descriptions on page 1271 for command syntax.

port-db

Syntax	port-db [sap <i>sap-id</i>] [group <i>grp-address</i> [source <i>ip-address</i>]] port-db sdp <i>sdp-id:vc-id</i> [group <i>grp-address</i> [source <i>ip-address</i>]]
Context	clear>service>id>igmp-snooping

Show, Clear, Debug Commands

Description	This command clears the information on the IGMP snooping port database for the VPLS service.
Parameters	<p>sap <i>sap-id</i> — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See Common CLI Command Descriptions on page 1271 for command syntax.</p> <p><i>sdp-id</i> — 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.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID for which to clear information.</p> <p>Default For mesh SDPs only, all VC IDs.</p> <p>Values 1 — 4294967295</p> <p>group <i>grp-address</i> — Clears IGMP snooping statistics matching the specified group address.</p> <p>source <i>ip-address</i> — Clears IGMP snooping statistics matching the specified particular source.</p>

querier

Syntax	querier
Context	clear>service>id>igmp-snooping
Description	This command clears the information on the IGMP snooping queriers for the VPLS service.

VPLS Debug Commands

id

Syntax	id <i>service-id</i>
Context	debug>service
Description	This command debugs commands for a specific service.
Parameters	<i>service-id</i> — The ID that uniquely identifies a service.
Values	service-id: 1 — 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.

Show, Clear, Debug Commands

mac

Syntax	[no] mac <i>ieee-address</i>
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 <i>sap-id</i>
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 <i>sdp-id:vc-id</i>
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	<i>sdp-id</i> — 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 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to display information. Values 1 — 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 — 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 <i>sap-id</i>
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 <i>sdp-id:vc-id</i>
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	<i>sdp-id</i> — Displays only MLD entries associated with the specified mesh SDP or spoke SDP. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to display information. Values 1 — 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

Show, Clear, Debug Commands

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.

mrrp-mac

Syntax **[no] mrrp-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 zeroes)

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	<i>sap-id</i> — See Common CLI Command Descriptions on page 1271 for command syntax.

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	<i>sdp-id</i> — Displays only MLD entries associated with the specified mesh SDP or spoke SDP. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to display information. Values 1 — 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.

enables the debug of the proxy-arp function for a given service. Alternatively, the debug can be enabled only for certain entries given by their IP or MAC addresses

enables the debug of the proxy-nd function for a given 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
Description	This command enables debugging for a particular SAP.
Parameters	<i>sap-id</i> — 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

Show, Clear, Debug Commands

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. See [Common CLI Command Descriptions on page 1271](#) for command syntax.

sdp

Syntax `[no] sdp sdp-id:vc-id`

Context `debug>service>stp`

Description This command enables STP debugging for a specific SDP.

provider-tunnels

Syntax `provider-tunnels type`

Context `tools>dump>service>vpls`

Description This command dumps the inclusive provider tunnels based on type.

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
```

Show, Clear, Debug Commands

```
=====
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	<pre>*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</pre>

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	<pre>*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</pre>

vxlan

Syntax	vxlan [clear]
Context	tools>dump>service
Description	<p>This command displays the number of times a service could not add a VXLAN binding or <VTEP, Egress VNI> due to the following limits:</p> <ul style="list-style-type: none"> - 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. <p>The command adds a [clear] option to clear the above statistics.</p>
Output	<pre>*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</pre>

dup-vtep-egrvni

Syntax	dup-vtep-egrvni [clear]
Context	tools>dump>service>vxlan

Show, Clear, Debug Commands

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 command provides a 'clear' option.

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

```
*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
```

IEEE 802.1ah Provider Backbone Bridging

In This Chapter

This chapter provides information about Provider Backbone Bridging (PBB), process overview, and implementation notes.

Topics in this chapter include:

- [IEEE 802.1ah Provider Backbone Bridging \(PBB\) Overview on page 888](#)
- [PBB Features on page 889](#)
 - [Integrated PBB-VPLS Solution on page 889](#)
 - [PBB Technology on page 891](#)
 - [PBB Mapping to Existing VPLS Configurations on page 892](#)
 - [SAP and SDP Support on page 894](#)
 - [PBB Packet Walkthrough on page 896](#)
 - [IEEE 802.1ak MMRP for Service Aggregation and Zero Touch Provisioning on page 921](#)
 - [MMRP Support Over B-VPLS SAPs and SDPs on page 923](#)
 - [PBB and BGP-AD on page 928](#)
 - [PBB ELINE Service on page 928](#)
 - [MAC Flush on page 931](#)
 - [Access Multi-Homing for Native PBB \(B-VPLS over SAP Infrastructure\) on page 936](#)
 - [PBB and IGMP/MLD Snooping on page 948](#)
 - [PBB QoS on page 949](#)
 - [PBB OAM on page 965](#)
- [Configuration Examples on page 967](#)

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.

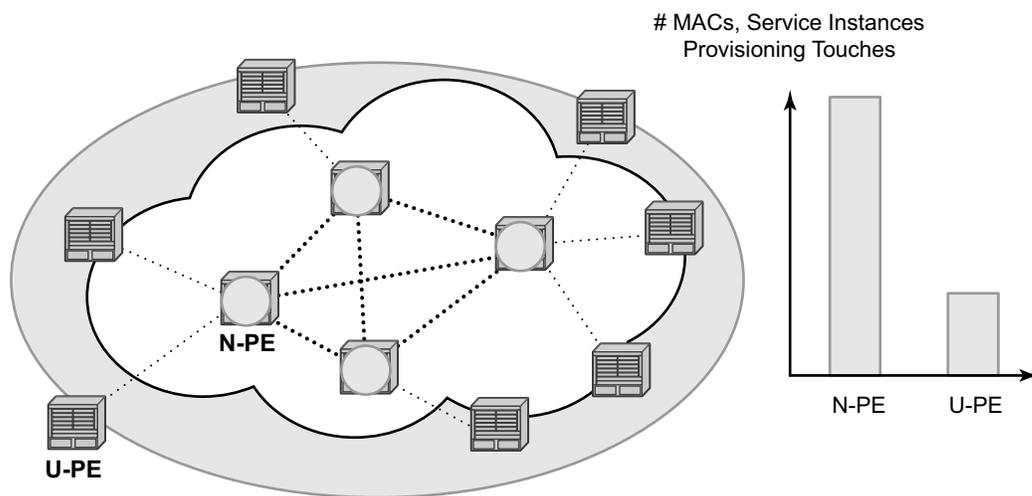
Virtual Private LAN Service (VPLS), RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, provides a solution for extending Ethernet LAN services using MPLS tunneling capabilities through a routed, traffic-engineered MPLS backbone without running (M)STP across the backbone. As a result, VPLS has been deployed on a large scale in service provider networks.

Alcatel-Lucent's 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.

PBB Features

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 82](#).

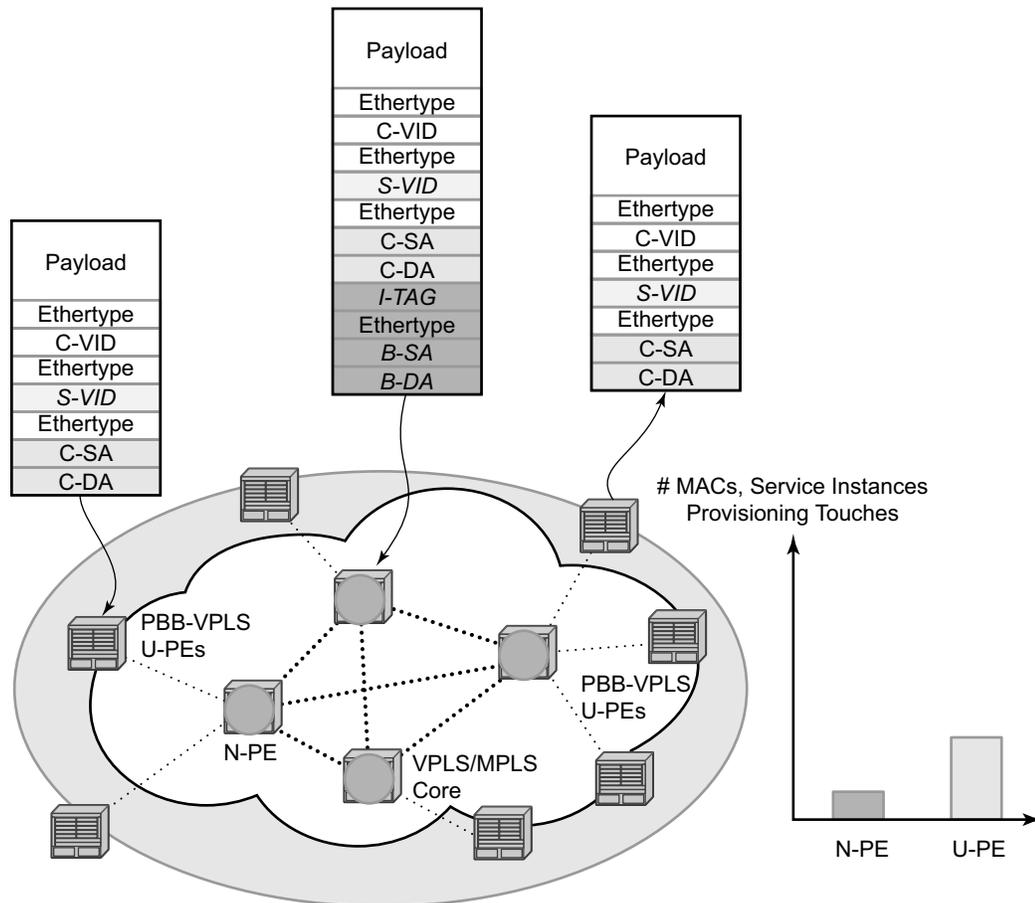


OSSG190

Figure 82: 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 83](#).



OSSG191

Figure 83: 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.

Alcatel-Lucent's 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.

PBB Technology

IEEE 802.1ah specification encapsulates the customer or QinQ payload in a provider header as shown in [Figure 84](#).

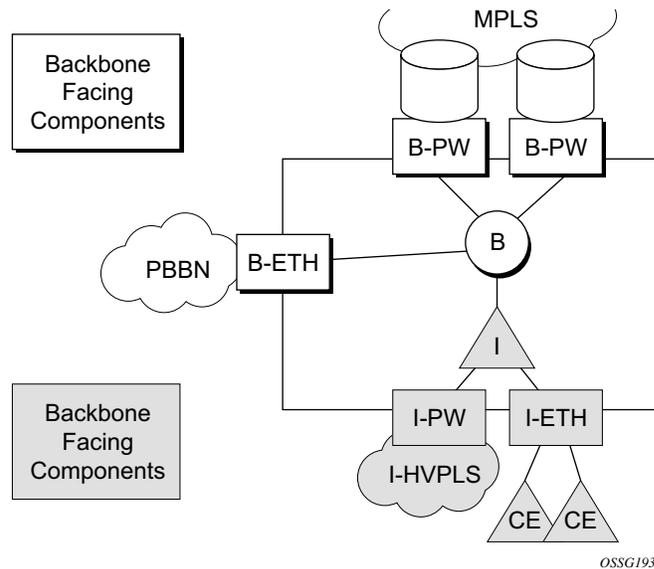


Figure 84: QinQ Payload in Provider Header Example

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 pseudo-wire (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.

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.

Alcatel-Lucent’s 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 85](#).

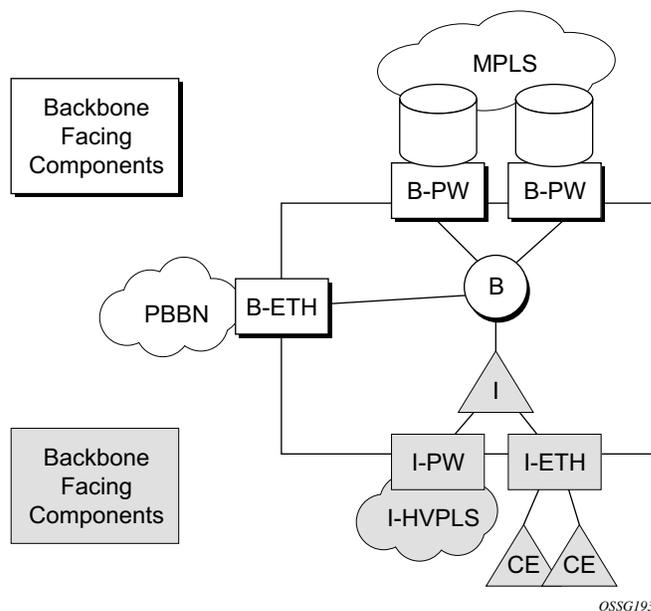


Figure 85: PBB Mapping to VPLS Constructs

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. Alcatel-Lucent’s 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 ELAN service (refer to the triangle marked “I” in [Figure 85](#)). 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.

In order to support PBB ELINE (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.

SAP and SDP Support

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 to 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.

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.
- Note that 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).

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 86](#).

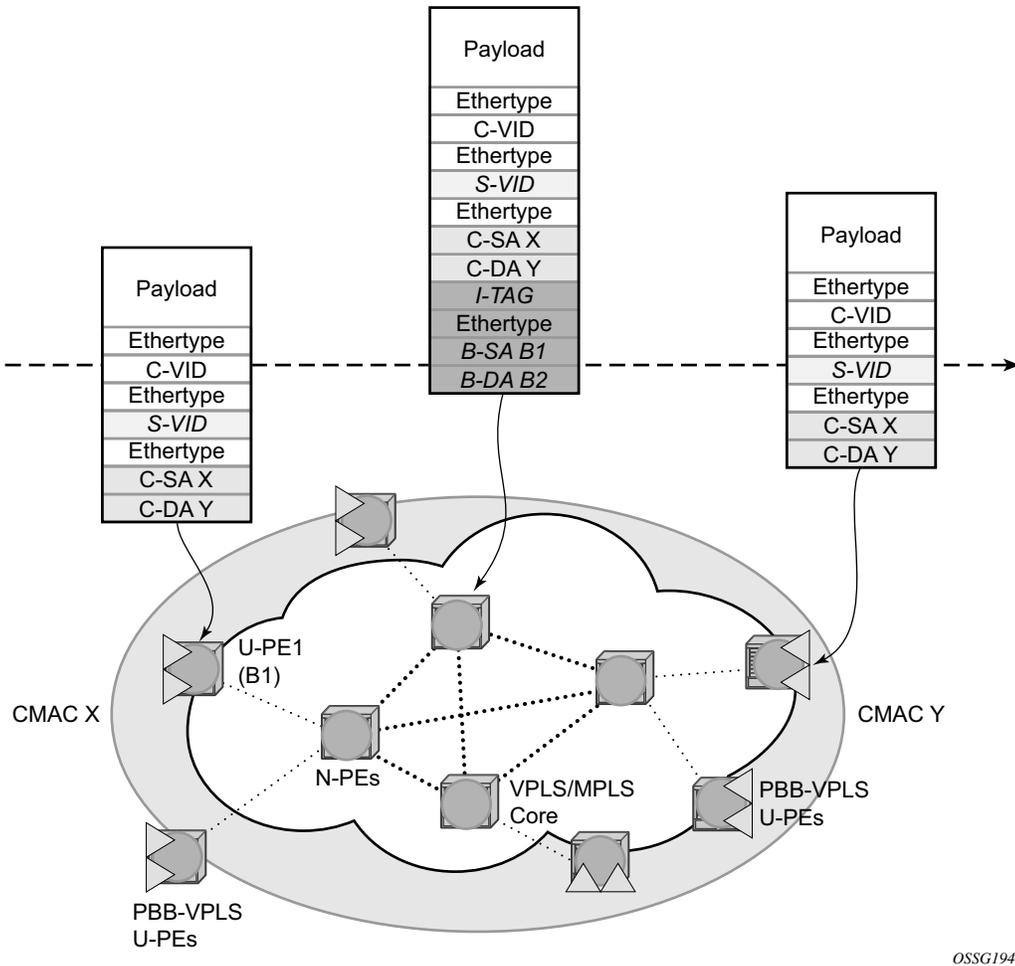


Figure 86: PBB Packet Walkthrough

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 FIB on U-PE1 is consulted to determine the destination B-MAC 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.

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 \(SPBM\) on page 898](#)
- Rapid Spanning Tree Protocol (RSTP) optionally with MMRP (SAPs and spoke SDPs); see [MMRP Support Over B-VPLS SAPs and SDPs on page 923](#).
- Multiple Spanning Tree Protocol (MSTP) optionally with MMRP (SAPs and spoke SDPs); see the *Layer 2 Service Guide* 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 on page 921](#).

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.

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. SROS 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.

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 15: 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.

SPB for B-VPLS

The SROS 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 15](#)). 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 SROS 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 SROS SPB allowing high scalability of PBB Epipes. I-VPLS I-SIDs are distributed in SROS 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 SROS 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.

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 SROS 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 been configured the same between neighboring switches. The FID for a given instance of SPB between two neighbor switches must be the same. The independence of VID encapsulation is inherent to SROS 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 SROS SPB.

Figure 87 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 used but in SROS 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 87 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 share the same topology but they may use different paths for data traffic; see [Shortest Path and Single Tree on page 902](#).

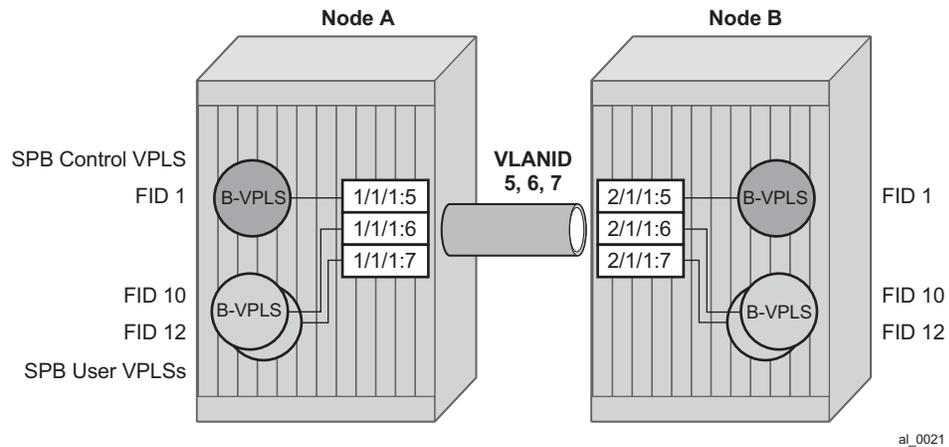


Figure 87: Control and User B-VPLS with FIDs

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 FIB. If I-VPLS services are to be supported on a B-VPLS that B-VPLS must have an independent FID.

Shortest Path and Single Tree

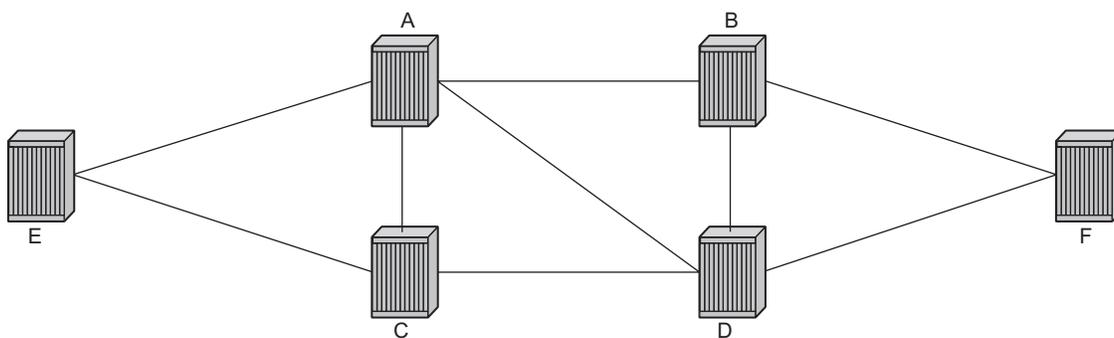
IEEE 802.1aq standard SPB uses a source specific tree model. The standard model is more computationally intensive for multicast traffic since in addition to the SPF algorithm for unicast and multicast from a single node, an all pairs shortest path needs to be computed for other nodes in the network. 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 SROS SPB design is optimized for unicast delivery using shortest paths. Ethernet supporting unicast and multicast services are commonly deployed in Ethernet transport networks. SROS SPB Single tree multicast (also called shared tree or *,G) operates similarly today. The difference is that SPB multicast never floods unknown traffic.

The SROS 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 88](#).



al_0022

Figure 88: 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 89](#) 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. Note that any nodes that are directly connected in this network have only one path between them (not shown for simplicity).

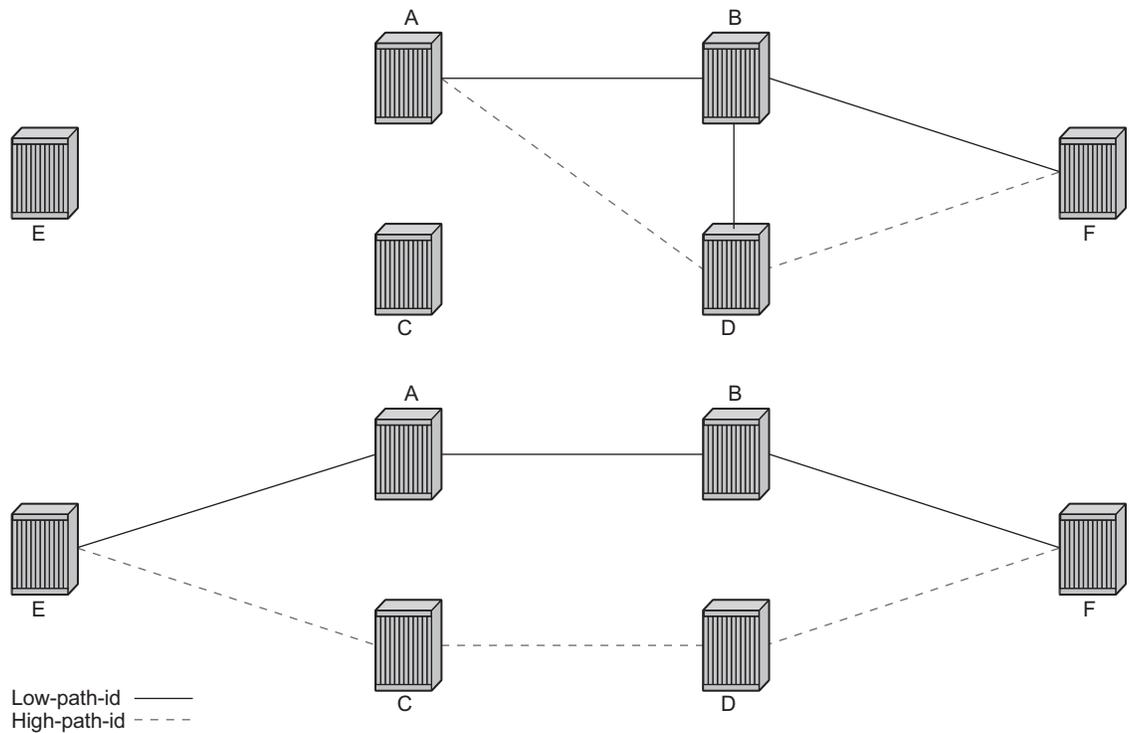


Figure 89: Unicast Paths for Low-path-id and High-path-id

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 90](#) illustrates the multicast paths for low-path-id and high-path-id algorithm.

Shortest Path Bridging MAC Mode (SPBM)

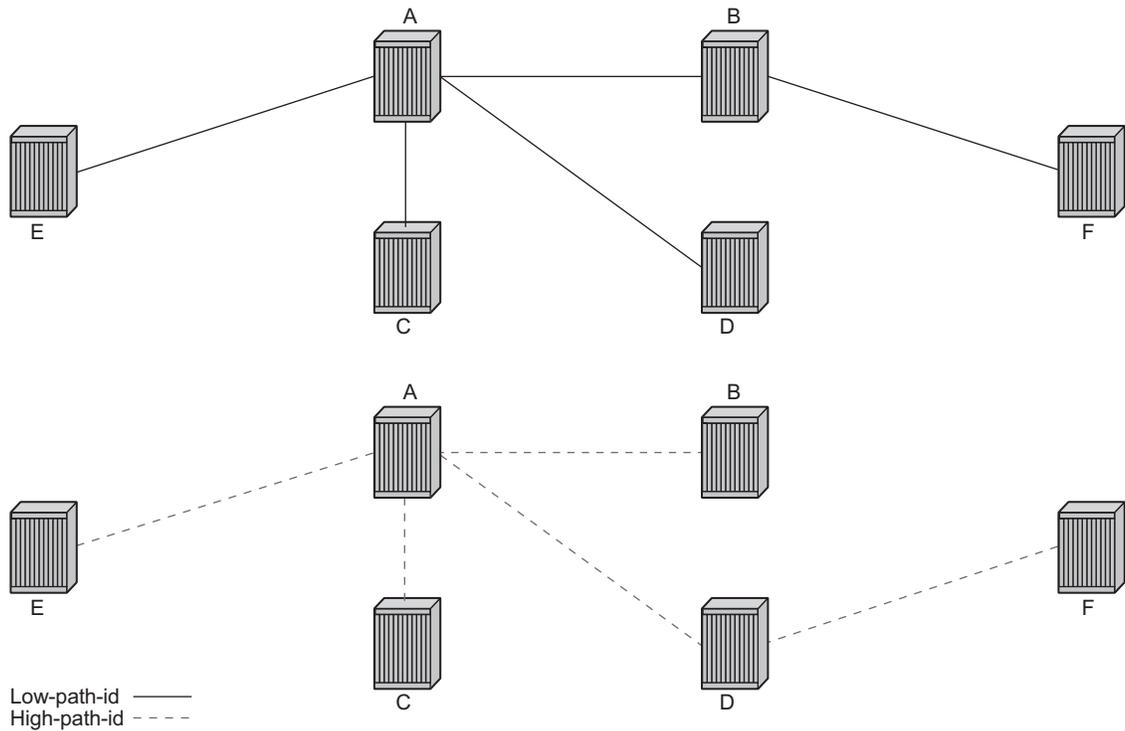


Figure 90: Multicast Paths for Low-path-id and High-path-id

All nodes in this network use one of these trees. Note that 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.

Data Path and Forwarding

The implementation of SPB on SROS 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 SROS. 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 SROS.) In addition, a single tree reduces the amount of computation required when there is topology change.

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 SROS supports both unicast and multicast forwarding, but with no learning and unicast and multicast may take different paths. In addition, SROS 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 given 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..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).

Shortest Path Bridging MAC Mode (SPBM)

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 16: 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

SPB Levels

Levels are part of IS-IS. SPB supports Level 1 within a control B-VPLS. Future enhancements may make use of levels.

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.

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.

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.

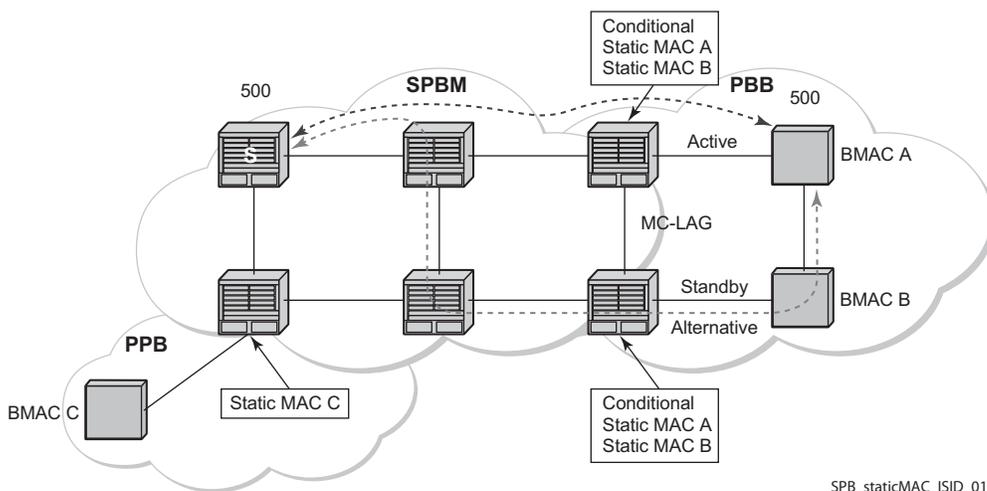


Figure 91: Static MACs Example

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.

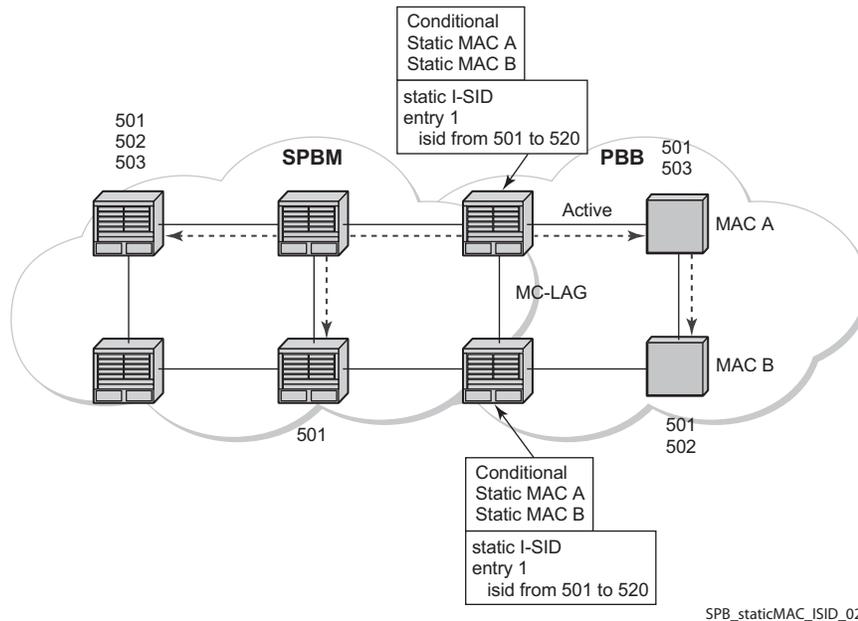


Figure 92: Static ISIDs Example

SPBM ISID Policies

Note that 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 17](#).

Table 17: 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.

ISID Policy Control

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 B-MAC used is the source-bmac of the Control/User VPLS. Typically remote B-MACs (the ultimate source-bmac) and the associated ISIDs are configured as static under the SPBM interface. This allows all remote B-MACs and all remote ISIDs can be configured once per interface.

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 B-MAC 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.

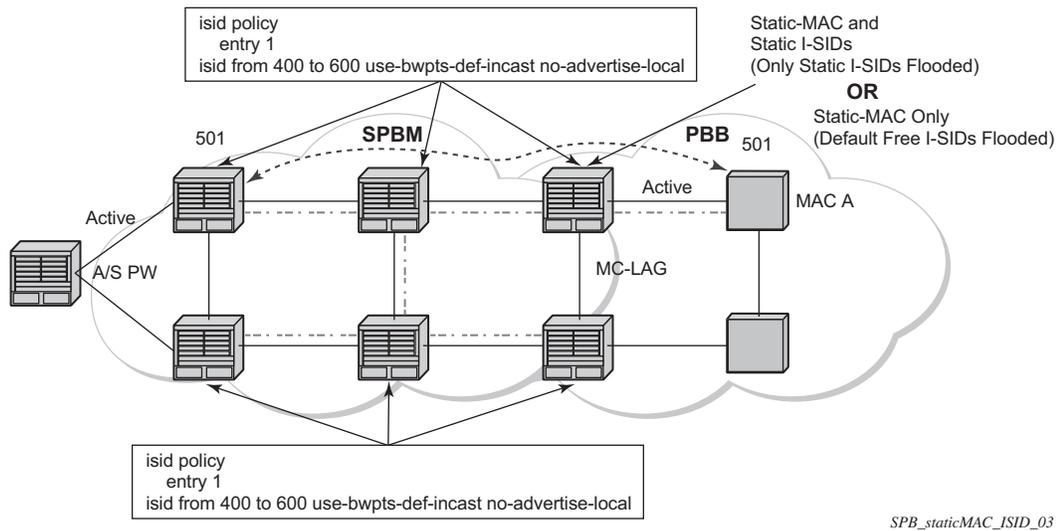


Figure 93: ISID Policy Example

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 desired, 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.

Example Network Configuration

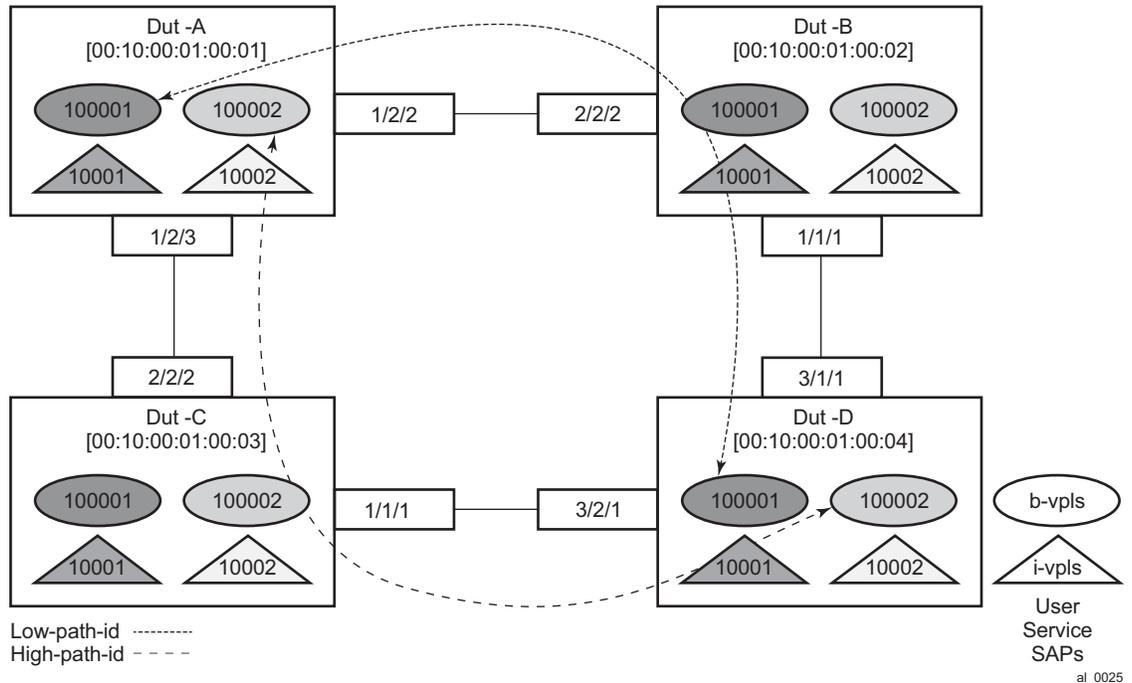


Figure 94: Sample Network

Figure 94 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 100001 to provide forwarding. It fate shares the control topology. In Figure 94, the control B-VPLS uses the low-path-id algorithm and the user B-VPLS uses high-path-id algorithm. Note that 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.

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
```

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. Note that 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
-----
```

Example Network Configuration

```

No. of MAC Routes: 12
=====
ISID Route Table
=====
Fid  ISID                               NextHop If                               SysID                               Ver.
-----
1    10001                               sap:1/2/2:1.1                             Dut-B                               10
                                    sap:1/2/3:1.1                             Dut-C
100  10002                               sap:1/2/2:1.1                             Dut-B                               10
                                    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. Note that 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
=====
```

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
-

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
-

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

IEEE 802.1ah 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). Alcatel-Lucent's 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 95](#).

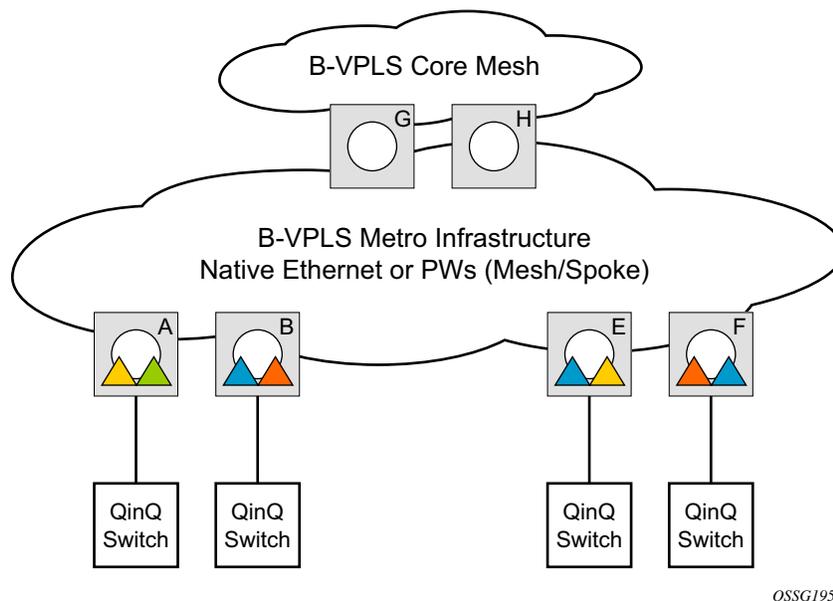


Figure 95: Customer Services Transported in 1 B-VPLS (M:1 Model)

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 96](#) 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.

Alcatel-Lucent’s 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 96](#).

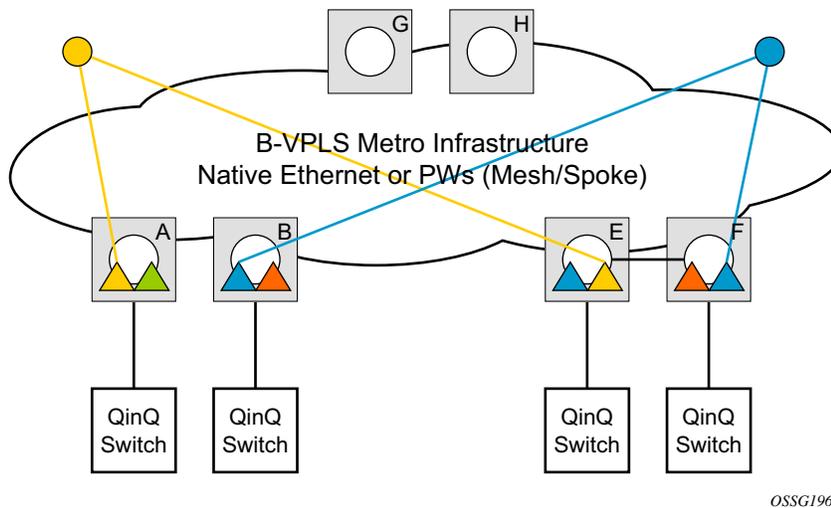


Figure 96: Flood Containment Requirement in M:1 Model

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 to 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.

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.

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.

Alcatel-Lucent's 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 FIB 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.

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.

In order to minimize the convergence time, Alcatel-Lucent's 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 the **configure>service>vpl>bvpls>mrp>flood-time** *value* command where *value* represents the amount of time in seconds that flooding will be enabled. Refer to the [PBB Command Reference on page 975](#) 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

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.

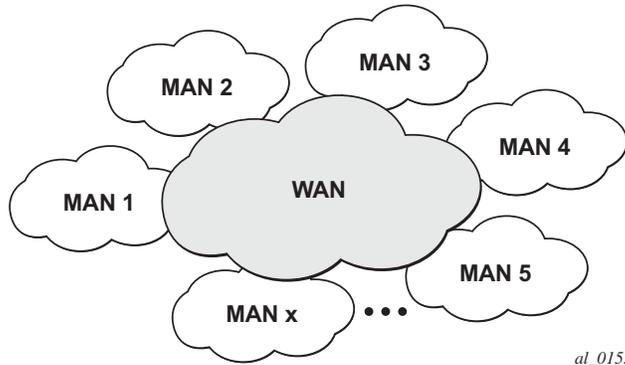


Figure 97: Inter-Domain Topology

Figure 97 depicts 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. In order 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 98. The case of native PBB metro domains inter-connected via a MPLS core is used in this example. Other technology combinations are possible.

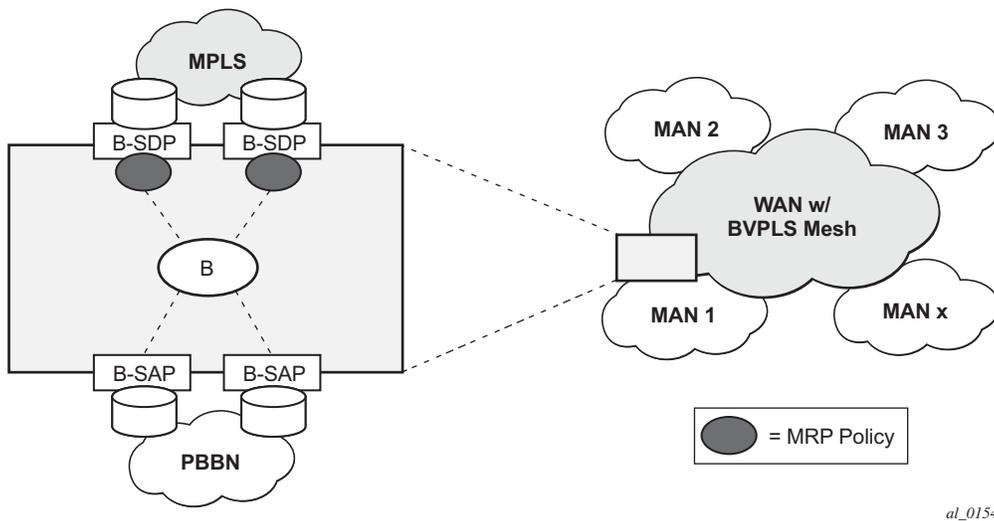


Figure 98: Limiting the Scope of MMRP Advertisements

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 criterium 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**.

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 *Layer 2 Service Guide*.

PBB ELINE Service

ELINE service is defined in PBB (IEEE 802.1ah) as a point-to-point service over the B-component infrastructure. Alcatel-Lucent's implementation offers support for PBB ELINE through the mapping of multiple Epipe services to a Backbone VPLS infrastructure.

The use of Epipe scales the ELINE 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. Note that 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 ELINE service support the configuration of a SAP or non-redundant spoke-SDP.

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.

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 IXCMs or XMASs.
- 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 ELINE (Epipe) and ELAN (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.

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. Note that 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 ageing timeout (assuming that the new periodic sending interval is less than the ageing 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-VPLSes (to avoid any traffic flooding), all above cases could experience an increased traffic interruption, or a permanent loss of traffic, as only traffic towards 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.

In order 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.

MAC Flush

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.

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 FIB 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 a 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.

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-VPLSes.

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 I-VPLS 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.

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

In order 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.

Alcatel-Lucent 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 FIBs throughout the core PEs. Subsequently, the receiving PBB VPLS PEs uses the BMAC and ISID list information to identify the specific I-VPLS FIBs 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 99](#) for the case when a Topology Change Notification (TCN) is received on PBB PE 2 from a QinQ access in the I-VPLS domain.

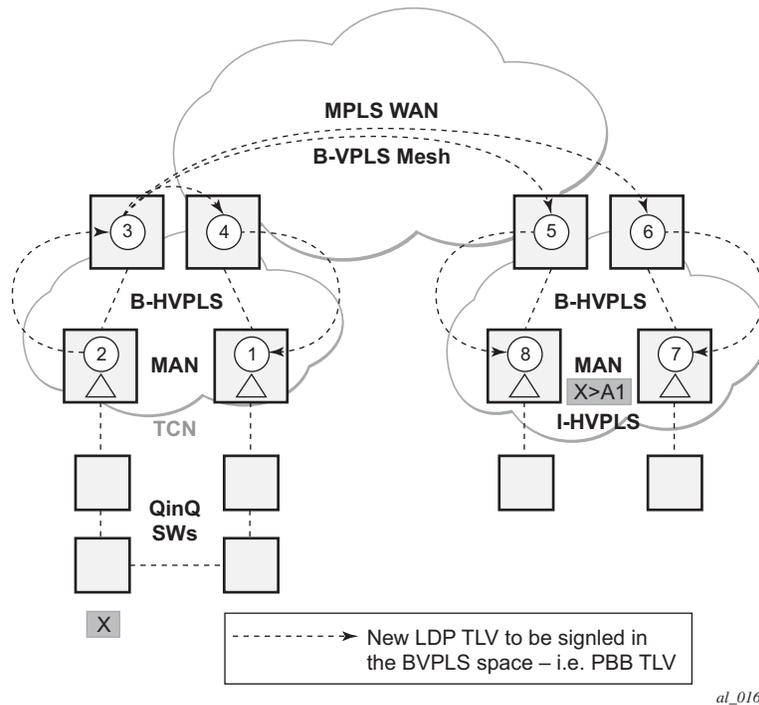


Figure 99: TCN Triggered PBB Flush-All-But-Mine Procedure

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 1-8 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 FIB 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 FIBs 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 B-MAC. 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**.

Access Multi-Homing for Native PBB (B-VPLS over SAP Infrastructure)

Alcatel-Lucent 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 100](#). MC-LAG is used to multi-home a number of edge switches running QinQ to PBB BEBs.

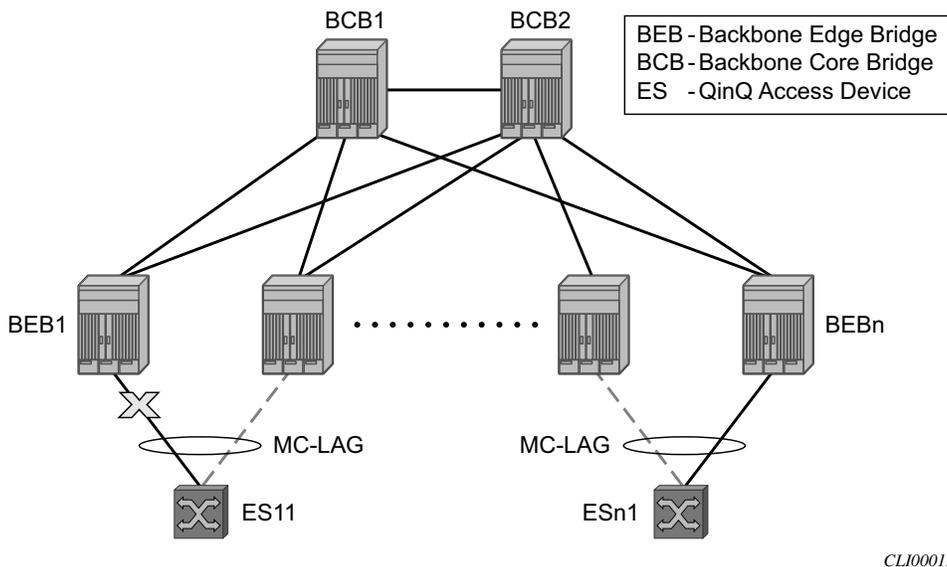


Figure 100: Access Dual-Homing into PBB BEBs - Topology View

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.

Alcatel-Lucent implementation provides a solution for both PBB ELINE (Epipe) and ELAN (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.

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 101](#) 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.

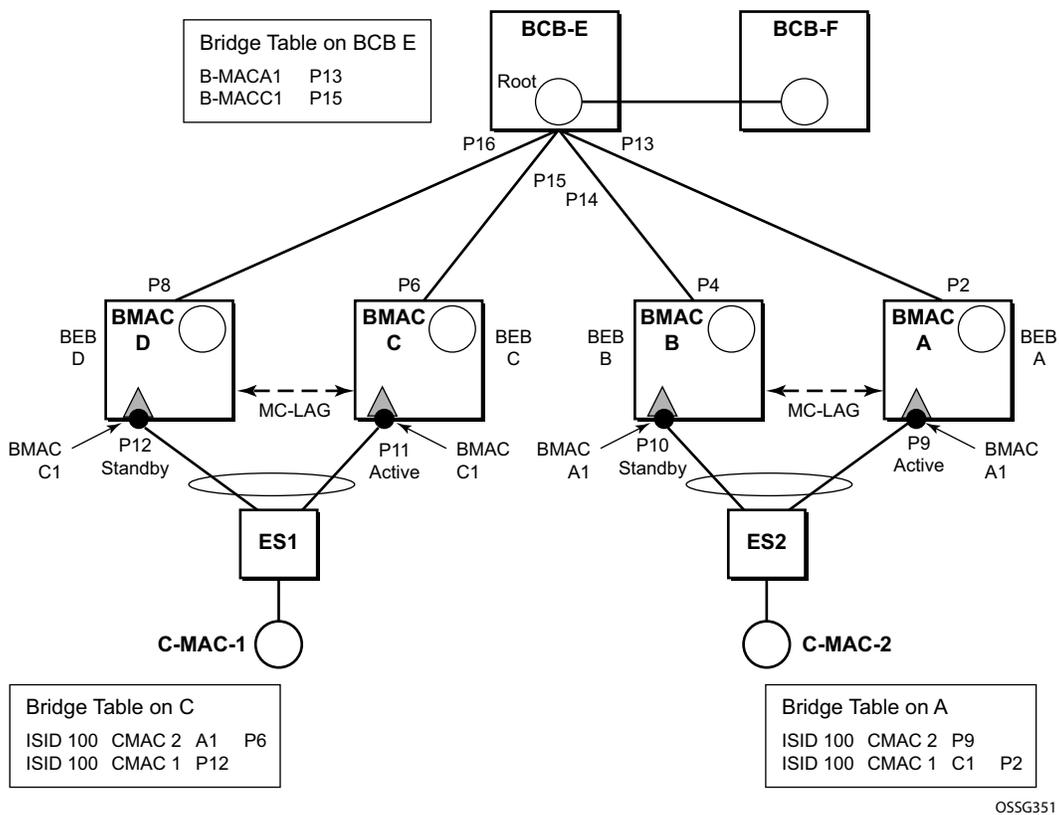


Figure 101: PBB Active Topology and Access Multi-Homing

Access Multi-Homing for Native PBB (B-VPLS over SAP Infrastructure)

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 given 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 FIB entries in BCB-E and BEB A as depicted in [Figure 101](#).

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 102](#) depicts the case of access link failure.

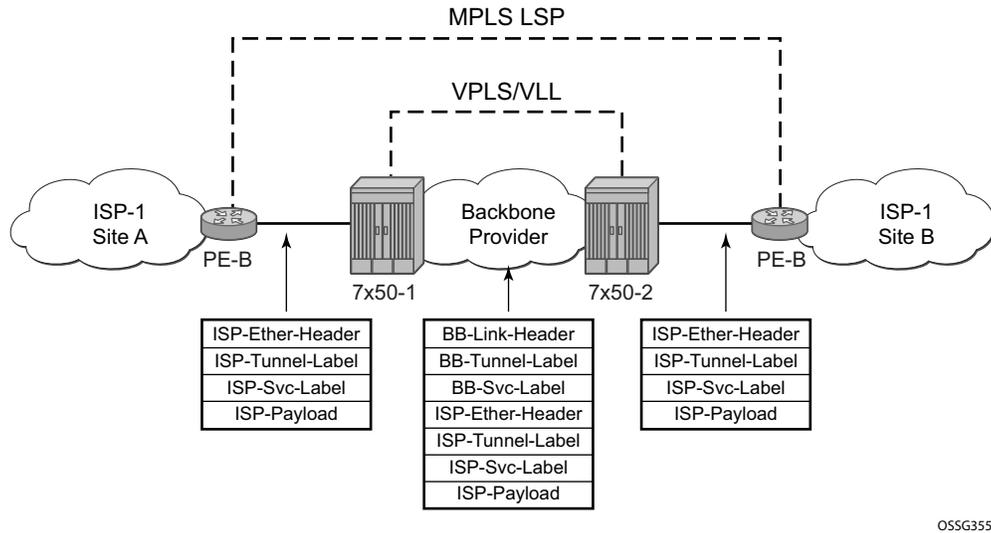


Figure 102: Access Multi-Homing - Link Failure

On failure of the active link P11 on BEB C the following processing steps apply:

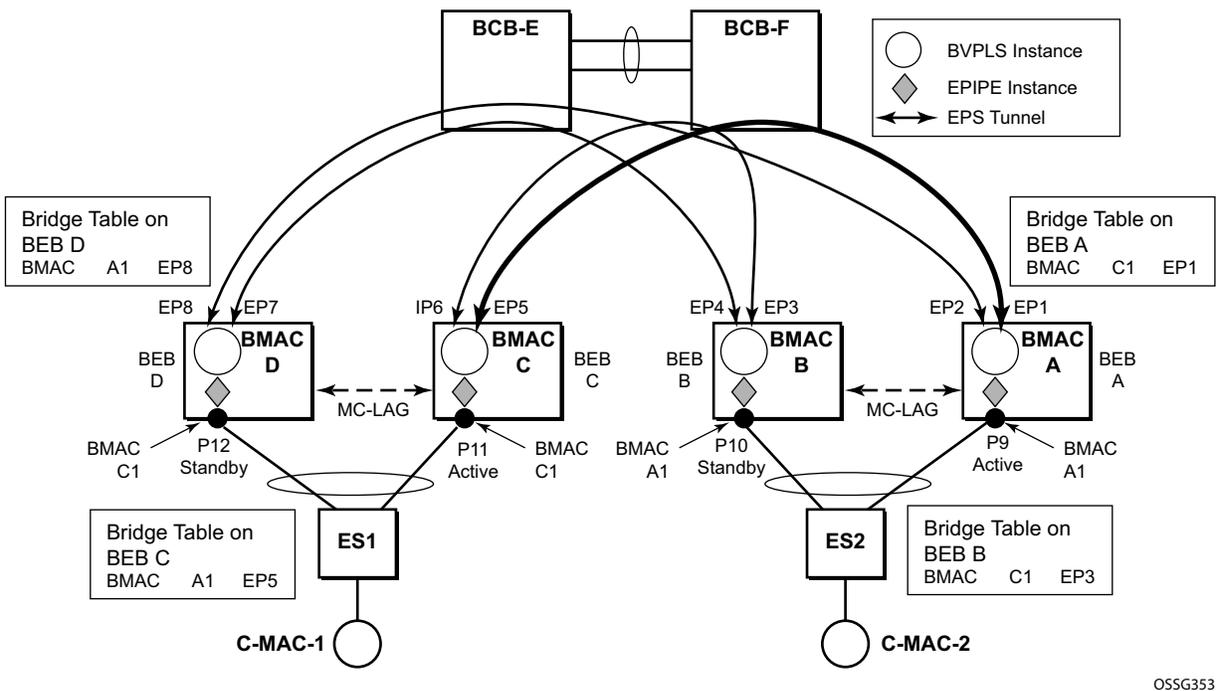
- MC-LAG protocol activates the standby link P12 on the pair BEB D.
- BMAC C1 becomes active on BEB D and any traffic received on BEB D with destination BMAC C1 is forwarded on the corresponding I-VPLS SAPs on P12.
- BEB D determines the related B-VPLS instance(s) associated with all the I-VPLS SAP(s) 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 Alcatel-Lucent OUI.
- As a result, all the FIB entries in BCBs or BEBs along the path will be automatically updated to reflect the move of BMAC C1 to BEB D.
- Note that 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 FIB 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 linecard.

Identical procedure is used when the whole BEB C fails.

Solution Description for PBB Epipe over G.8031 Ethernet Tunnels

This section discusses the Access Multi-Homing solution for PBB ELINE 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 ELINE service and the related BVPLS infrastructure are depicted in [Figure 103](#).



OSSG353

Figure 103: Access Multi-Homing Solution for PBB Epipe

The ELINE 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 ELAN 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 given 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 FIB entries in BEB C and BEB A as depicted in Figure 103. 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 104. The same procedure applies for the link failure case.

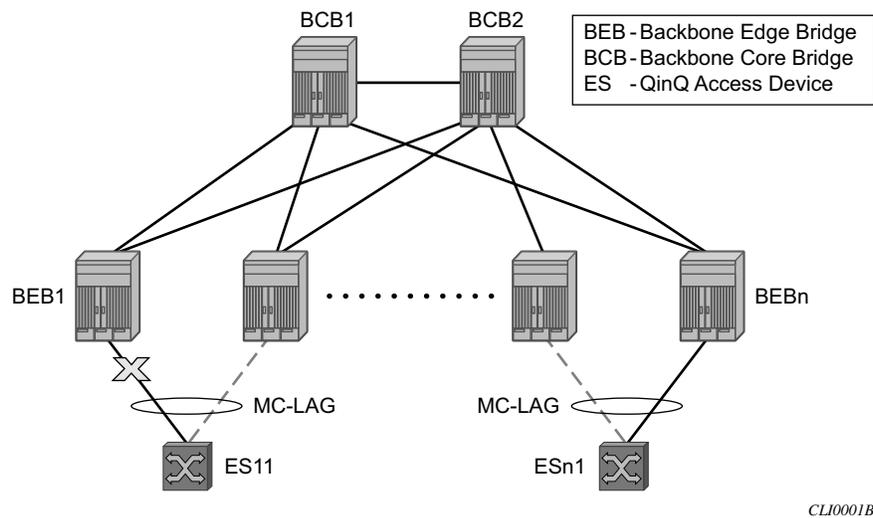


Figure 104: Access Dual-Homing for PBB ELINE - BEB Failure

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.

Access Multi-Homing for Native PBB (B-VPLS over SAP Infrastructure)

- As a result, the FIB 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.

Note that the same process is executed for all the MC-LAGs affected by BEB C failure so BEB failure will be the worst case scenario.

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 105](#).

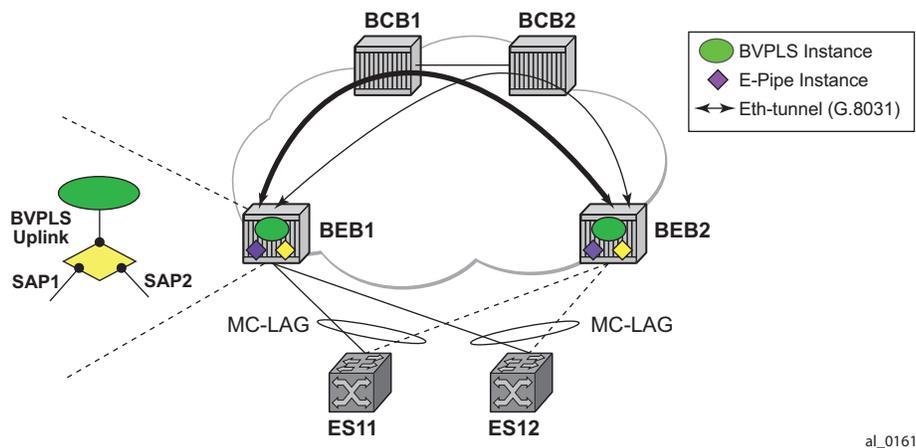


Figure 105: Solution for Access Dual-Homing with Local Switching for PBB Eline/epipe

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 105](#))

- 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.

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 the *Layer 2 Services 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-bypls-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 on page 932](#).

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.

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 106](#).

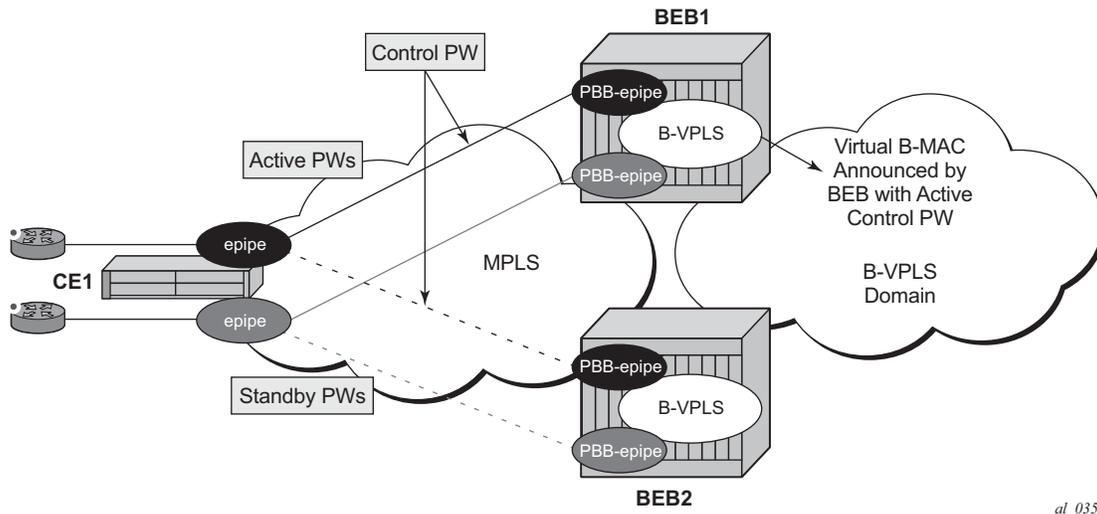


Figure 106: Active/Standby PW into PBB Epipes

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 106](#), 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 desired 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. Note that 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 given 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 106](#)):

- If any non-control pseudowires 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 desired, 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 7x50 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 pseudowires to be part of a multi-homed site.
- Network-domains.
- Support for the following tunneling technologies
 - RFC 3107
 - GRE
 - L2TPv3

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 B-MAC address used in the PBB header for the query message. This is the B-MAC 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` CLI command to specify the mac-name of the destination B-MAC 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 or the maximum number of multicast sources allowed per group

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). Note that 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.

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 107](#).

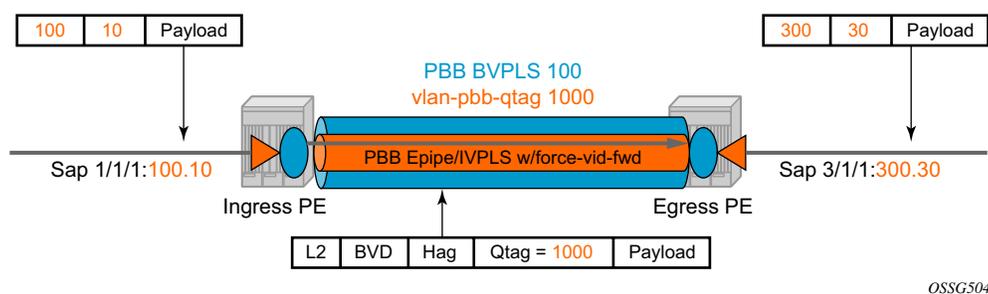


Figure 107: PCP, DE Bits Transparency in PBB

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 dot1p 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.

Once 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.

Configuration Examples

This section gives usage examples for the new commands under PBB Epipe or IVPLS instances.

PBB IVPLS usage:

```
configure service vpls 100 ivpls
  sap 1/1/1:101
  pbb
    backbone-vpls 10 isid 100
    force-qtag-forwarding
```

PBB Epipe Usage:

```
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
```

Details Solution Description

Figure 107 depicts a specific use case. Keeping the same topology - an ingress PBB PE, a PBB core and an egress PBB PE - let us 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. lastly, 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.

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.

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. Note that 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 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
```

Note that 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.

Note that 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 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.

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.

Egress Queue Scheduling

Figure 108 displays an example of egress queue scheduling.

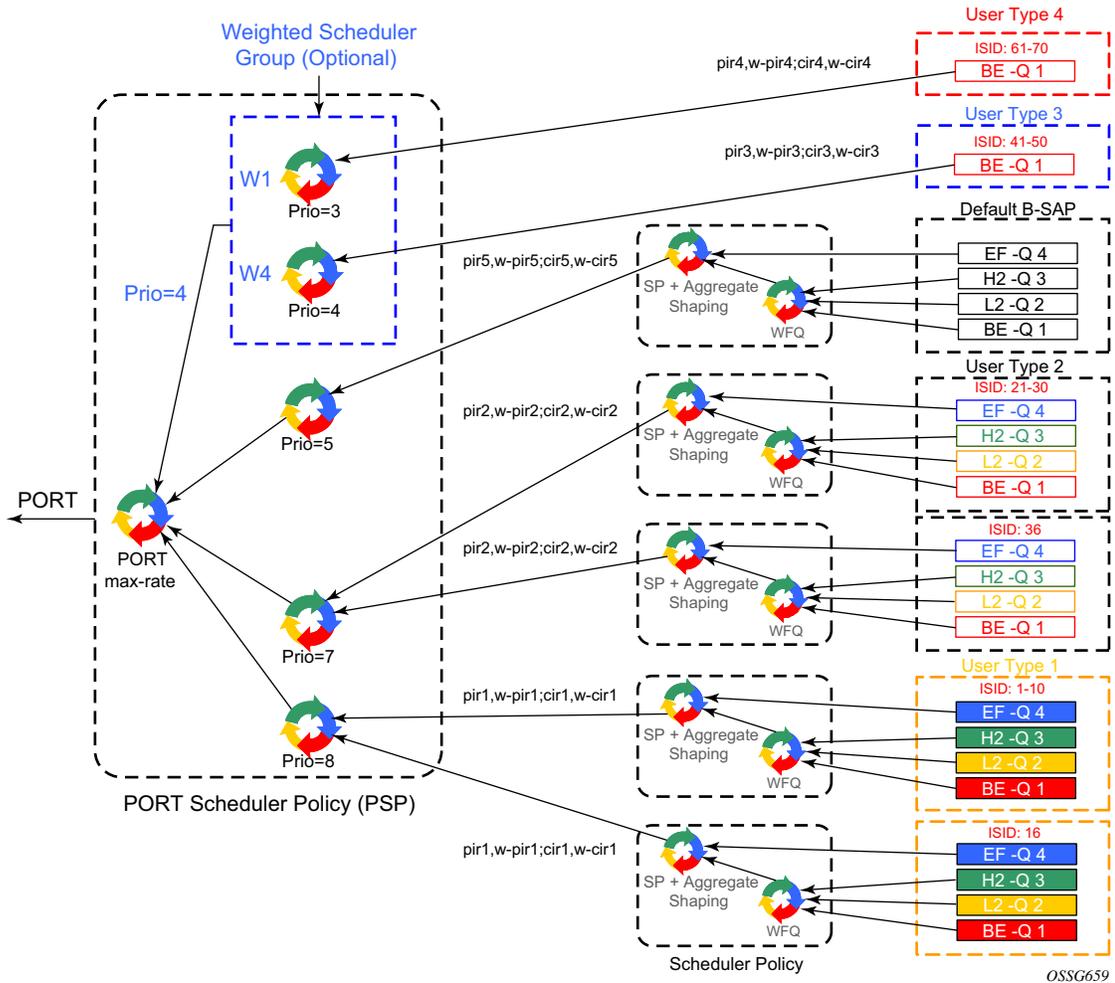


Figure 108: Egress Queue Scheduling

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 108 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.

In order 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 but the packet-byte-offset value is still in effect as explained above.

B-SAP per-ISID Shaping Configuration Example

The following CLI configuration for B-SAP per-ISID shaping achieves the specific use case shown in [Figure 108 on page 960](#).

```

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

```

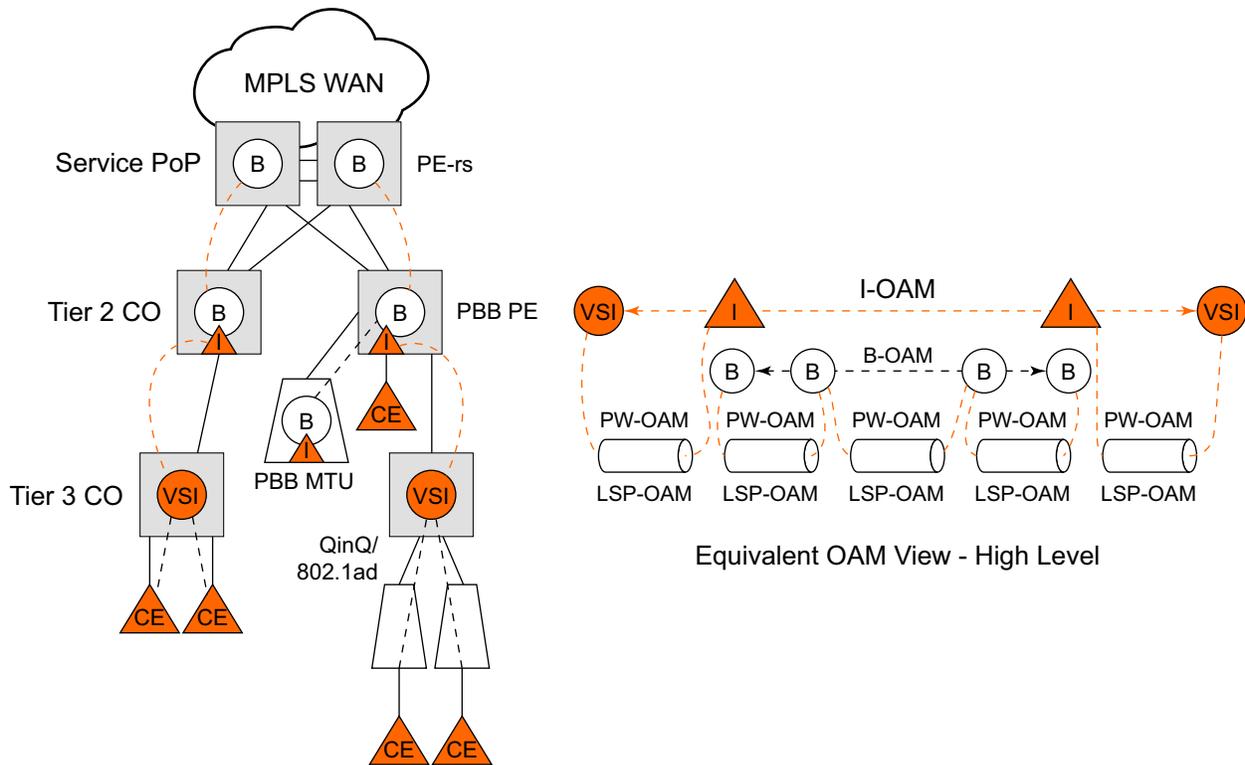
Egress B-SAP per ISID Shaping

```
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
exit
```

PBB OAM

Alcatel-Lucent's PBB implementation support 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 109](#).



OSSG200

Figure 109: 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 signalling 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.

Mirroring

There are no restrictions for mirroring in I-VPLS or B-VPLS.

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 towards 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 (towards 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)
-

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.

Configuration Examples

Use the CLI syntax displayed to configure PBB.

PBB using G.8031 Protected Ethernet Tunnels

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
    encaps-type dot1q
  exit
  ccm-hold-time down 5 up 10 // 50 ms down, 1 sec up
  lag-emulation
    access adapt-gos 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
    encaps-type dot1q
```

Configuration Examples

```
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

```

CLI Syntax:

```

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

```

Configuration Examples

```
                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
```

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)
```

Access Multi-Homing over MPLS for PBB Epipes

This section gives an example configuration for BEB1 from [Figure 106](#).

```
*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

```

PBB Command Reference

Command Hierarchies

- [Global Commands on page 975](#)
- [SAP Commands on page 977](#)
- [Mesh SDP Commands on page 977](#)
- [Spoke SDP Commands on page 977](#)
- [Show Commands on page 979](#)
- [Clear Commands on page 979](#)
- [Debug Commands on page 980](#)

Global Commands

```

config
  — service
    — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls|i-vpls] [create]
      — [no] spb instance [fid value] [create]
        — [no] shutdown
        — level level-number
          — bridge-priority value
          — ect-algorithm name fid-range fid-range
          — forwarding-tree-topology[st|spf]
        — lsp-lifetime seconds
        — no lsp-lifetime
        — lsp-wait lsp-wait [lsp-initial-wait] [lsp-second-wait]
        — overload [timeout seconds]
        — no overload
        — overload-on-boot [timeout seconds]
        — no overload-on-boot
        — [no] spf-wait spf-wait [spf-initial-wait] [spf-second-wait]
      — spbm-control-vpls mgmt vpls svc id fid val
      — 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

```

Command Hierarchies

```
— source-bmac ieee-address
— no source-bmac
  — backbone-smac ieee-address
  — no backbone-smac

config
— service
  — [no] vpls service-id [customer customer-id] [b-vpls] [create]
    — pbb
      — backbone-vpls service-id[:isid]
      — no backbone-vpls
      — [no] force-qtag-forwarding
      — source-bmac ieee-address
      — no source-bmac
      — [no] use-sap-bmac
      — mac-notification
        — [no] count value
        — [no] interval value
        — renotify value
        — no renotify

config
— service
  — [no] vpls service-id [customer customer-id] [i-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
      — [no] send-bvpls-flush {[all-from-me] | [all-but-mine]}

config
— service
  — service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching]
  — no service-id
    — pbb
      — [no] force-qtag-forwarding
```

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]
      — 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 milliseconds
          — 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

```

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

```

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** *milliseconds*
 - **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**

BGP-MH for I-VPLS Commands

Note: Refer to the *Layer 2 Services Guide* for information about BGP-MH for I-VPLS commands.

- ```

config
 — service
 — vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls] [create]
 — no vpls service-id
 — site name [create]
 — 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-id value
 — no site-id
 — split-horizon-group group-name
 — no split-horizon-group
 — spoke-sdp sdp-id:vc-id
 — no spoke-sdp

```

## Show Commands

```

show
 — eth-cfm
 — association [ma-index] [detail]
 — cfm-stack-table
 — cfm-stack-table port [{all-ports | all-sdps | all-virtuals}][level 0..7][direction up|down]
 — cfm-stack-table port-id [vlan qtag [qtag]] [level 0..7] [direction up|down]
 — cfm-stack-table sdp sdp-id[:vc-id] [level 0..7][direction up|down]
 — cfm-stack-table virtual service-id [level 0..7]
 — cfm-stack-table facility [{all-ports|all-lags|all-lag-ports|all-tunnel-meps| all-router-inter-
 faces}][level 0..7] [direction up|down]
 — cfm-stack-table facility collect-lmm-stats
 — cfm-stack-table facility lag id [tunnel 1..4094] [level 0..7] [direction up|down]
 — cfm-stack-table facility port id [level 0..7] [direction up|down]
 — cfm-stack-table facility router-interface ip-int-name [level 0..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]

```

## Clear Commands

```

clear
 — service
 — statistics
 — id service-id
 — counters

```

## Command Hierarchies

- **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}

## Debug Commands

- debug**
  - **service**
    - **id** *service-id*
      - [no] **mrp**
        - **all-events**
        - [no] **applicant-sm**
        - [no] **leave-all-sm**
        - [no] **mmp-mac** *ieee-address*
        - [no] **mrpdu**
        - [no] **periodic-sm**
        - [no] **registrant-sm**
        - [no] **sap** *sap-id*
        - [no] **sdp** *sdp-id:vc-id*
        - [no] **spb**
          - [no] **adjacency** {sap *sap-id* | spoke-sdp *sdp-id:vc-id* | nbr-system-id}
          - [no] **interface** { sap <sap-id> | spoke-sdp <sdp-id:vc-id> }
          - [no] **l2db**
          - [no] **lsdb** {system-id | lsp-id }
          - [no] **packet** { ptop-hello l1-hello l1-psnp l1-csnp l1-lsp }  
**detail**
          - [no] **spf** { system-id }

---

## PBB Service Commands

---

### VPLS Service Commands

#### vpls

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>vpls</b> <i>service-id</i> <b>customer</b> <i>customer-id</i> <b>vpn</b> <i>vpn-id</i> [ <b>m-vpls</b> ] [ <b>b-vpls</b>   <b>i-vpls</b> ] [ <b>create</b> ]<br><b>vpls</b> <i>service-id</i><br><b>no vpls</b> <i>service-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command creates or edits a Virtual Private LAN Services (VPLS) instance. The <b>vpls</b> command is used to create or maintain a VPLS service. If the <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> exists, the context for editing the service is entered.</p> <p>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.</p> <p>When a service is created, the <b>create</b> keyword must be specified if the <b>create</b> command is enabled in the <b>environment</b> context. When a service is created, the <b>customer</b> keyword and <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> must already exist having been created using the <b>customer</b> 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.</p> <p>Once a service is created, the use of the <b>customer</b> <i>customer-id</i> is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect <i>customer-id</i> specified will result in an error.</p> <p>More than one VPLS service may be created for a single customer ID.</p> <p>By default, no VPLS instances exist until they are explicitly created.</p> <p>The <b>no</b> form of this command deletes the VPLS service instance with the specified <i>service-id</i>. 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.</p> <p><i>service-id</i> — 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 <i>service-id</i> must be the same number used for every SR OS router on which this service is defined.</p> <p><b>Values</b>     1 — 2147483648</p> <p><b>customer</b> <i>customer-id</i> — 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.</p> <p><b>Values</b>     1 — 2147483647</p> |

## VPLS Service Commands

**vpn** *vpn-id* — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

**Values** 1 — 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

## 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** 1-64

## 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 once created. To change a FID the SPB component would have to be shutdown, deleted and recreated with a new FID.

|                   |                                                                           |
|-------------------|---------------------------------------------------------------------------|
| <b>Default</b>    | no spb                                                                    |
| <b>Parameters</b> | <i>instance-id</i> — Specifies the instance ID for an SPB IS-IS instance. |
|                   | <b>Values</b> 1024–2047 (4 available)                                     |
|                   | <b>Default</b> 1024                                                       |
|                   | <i>FID</i> — Specifies FID value.                                         |
|                   | <b>Values</b> 1-4095                                                      |
|                   | <b>Default</b> 1                                                          |

Note: SPB operates with disable-learning, disable aging and discard-unknown. The state of these commands is ignored when SPB is configured.

## spb

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] spb [create]</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>vpls b-vpls>sap>spb><br>config>service>vpls b-vpls>spoke-sdp>spb>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | 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.<br><br>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 <b>config&gt;service&gt;vpls b-vpls&gt;spb</b> context. |
| <b>Default</b>     | no spb                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## spbm-control-vpls

|                    |                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>spbm-control-vpls service-id fid fid</b><br><b>no spbm-control-vpls</b>                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>vpls <i>service-id</i> b-vpls>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | 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.<br><br>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 Service Commands

VPLS, SAPs or SDP bindings associated with the FID. The FID must be independent from the FID assigned to other services.

**Default** none

## 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 LSP's every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from this formula:  $\text{lsp-lifetime}/2$

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 — 65535

## lsp-wait

**Syntax** **lsp-wait lsp-wait [lsp-initial-wait [lsp-second-wait]]**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>vpls b-vpls>spb                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command is used to customize the throttling of SPB LSP-generation. 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 <i>lsp-wait</i> timer until a maximum value is reached. This is a control B-VPLS command.                                                                                                                                                      |
| <b>Parameters</b>  | <p><i>lsp-max-wait</i> — Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated.</p> <p><b>Values</b> 1 — 120</p> <p><b>Default</b> 5</p> <p><i>lsp-initial-wait</i> — Specifies the initial LSP generation delay in seconds.</p> <p><b>Values</b> 0 — 100</p> <p><b>Default</b> 0</p> <p><i>lsp-second-wait</i> — Specifies the hold time in seconds between the first and second LSP generation.</p> <p><b>Values</b> 1 — 100</p> <p><b>Default</b> 1</p> |

## overload

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>overload [timeout seconds]</b><br><b>no overload</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>vpls b-vpls>spb                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>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 be used for other transit traffic.</p> <p>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.</p> <p>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.</p> <p>The <b>no</b> form of the command causes the router to exit the overload state.</p> |
| <b>Default</b>     | no overload                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>seconds</i> — The time, in seconds, that this router must operate in overload state.</p> <p><b>Values</b> 60 — 1800</p> <p><b>Default</b> Infinity (overload state maintained indefinitely)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## overload-on-boot

## VPLS Service Commands

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>overload-on-boot</b> [timeout <i>seconds</i> ]<br><b>no overload-on-boot</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>vpls b-vpls>spb>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>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 bootup in the overload state until one of the following events occur:</p> <ul style="list-style-type: none"><li>• The timeout timer expires.</li><li>• A manual override of the current overload state is entered with the <b>config&gt;service&gt;vpls instance&gt;b-vpls&gt;spb&gt;no overload</b> command.</li></ul> <p>The <b>no</b> form of the command does not affect the overload-on-boot function.</p> <p>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:</p> <pre>L1 LSDB Overload : Manual on boot (Indefinitely in overload)</pre> <p>This state can be cleared with the <b>config&gt;service&gt;vpls instance&gt;b-vpls&gt;spb&gt;no overload</b> command.</p> <p>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:</p> <pre>L1 LSDB Overload : Manual on boot (Overload Time Left : 17)</pre> <p>The overload state can be cleared before the timeout expires with <b>config&gt;service&gt;vpls instance&gt;b-vpls&gt;spb&gt;no overload</b> command.</p> <p>The <b>no</b> form of the command removes the overload-on-boot functionality from the configuration.</p> |
| <b>Default</b>     | no overload-on-boot                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>seconds</i> — The time, in seconds, that this router must operate in overload state.<br><b>Values</b> 60 — 1800<br><b>Default</b> Infinity (overload state maintained indefinitely)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## spf-wait

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] spf-wait</b> <i>spf-wait</i> [ <i>spf-initial-wait</i> [ <i>spf-second-wait</i> ]]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>vpls b-vpls>spb>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command defines the maximum interval between two consecutive SPF calculations in seconds. Timers that determine when to initiate the first, second and subsequent SPF calculations after a topology change occurs can be controlled with this command.</p> <p>Subsequent SPF runs (if required) occur at exponentially increasing intervals of the <i>spf-second-wait</i> interval. For example, if the <i>spf-second-wait</i> 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 <i>spf-wait</i> value. The SPF interval remains at the <i>spf-wait</i> value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval drops back to <i>spf-initial-wait</i>.</p> |

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no spf-wait                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b> | <p><i>spf-wait</i> — Specifies the maximum interval in seconds between two consecutive spf calculations.</p> <p><b>Values</b> 1 — 120</p> <p><b>Default</b> 10</p> <p><i>spf-initial-wait</i> — Specifies the initial SPF calculation delay in milliseconds after a topology change.</p> <p><b>Values</b> 10 — 100000</p> <p><b>Default</b> 1000</p> <p><i>spf-second-wait</i> — Specifies the hold time in milliseconds between the first and second SPF calculation.</p> <p><b>Values</b> 1 — 100000</p> <p><b>Default</b> 1000</p> |

## level

|                    |                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>level</b> <i>level-number</i>                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>vpls b-vpls>spb                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>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.</p> <p>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.</p> |
| <b>Default</b>     | level 1                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <p><i>level-number</i> — The SPB level number.</p> <p><b>Values</b> 1, 2</p>                                                                                                                                                                                                                                                                                                               |

## bridge-priority

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>bridge-priority</b> <i>value</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>vpls b-vpls>spb>level <i>level-number</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>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.</p> <p>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.</p> |

|                |               |        |
|----------------|---------------|--------|
|                | <b>Values</b> | 0 — 15 |
| <b>Default</b> |               | 8      |

## ect-algorithm

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ect-algorithm</b> <i>name</i> <b>fid-range</b> <i>fid-range</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>vpls b-vpls>spb>level level-number                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures the ect-algorithm associated with a FID. Names are:</p> <ul style="list-style-type: none"> <li>• low-path-id</li> <li>• high-path-id</li> </ul> <p>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).</p> <p>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).</p> <p>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.</p> <p>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.</p> |
| <b>Default</b>     | low-path-id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <p><i>name</i> — low-path-id, high-path-id</p> <p><i>fid-range</i> — Range of Forwarding Identifier values.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|                    | <b>Values</b> 1 — 4095                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## forwarding-tree-topology

|                    |                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>forwarding-tree-topology unicast</b> [st spf]                                                                                                                                                                                    |
| <b>Context</b>     | config>service>vpls b-vpls>spb>level level-number                                                                                                                                                                                   |
| <b>Description</b> | <p>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.</p> |

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

|                    |                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsp-pacing-interval</b> <i>milliseconds</i><br><b>no lsp-pacing-interval</b>                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>vpls b-vpls>sap>spb><br>config>service>vpls b-vpls>spoke-sdp>spb>                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | This command configures the interval between SPB LSP PDUs sent from this interface. This command is valid only for interfaces on control B-VPLS.<br><br>To avoid bombarding adjacent neighbors with excessive data, pace the Link State Protocol Data Units (LSP's). If a value of zero is configured, no LSP's are sent from the interface.<br><br>The no form of the command reverts to the default value. |
| <b>Default</b>     | 100 — LSPs are sent in 100 millisecond intervals.                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>milliseconds</i> — The interval in milliseconds that SPB IS-IS LSP's can be sent from the interface expressed as a decimal integer.<br><br>0 — 65535                                                                                                                                                                                                                                                      |

## retransmit-interval

|                    |                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>retransmit-interval</b> <i>seconds</i><br><b>no retransmit-interval</b>                                                                                                                                                           |
| <b>Context</b>     | config>service>vpls b-vpls>sap>spb><br>config>service>vpls b-vpls>spoke-sdp>spb>                                                                                                                                                     |
| <b>Description</b> | 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.<br><br>The no form of the command reverts to the default value. |
| <b>Default</b>     | 100                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>seconds</i> — The interval in seconds that SPB IS-IS LSPs can be sent on the interface.<br><br><b>Values</b> 1 — 65535                                                                                                            |

## metric

**Syntax**    **metric** *value*  
**No metric**

## VPLS Service Commands

|                    |                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>vpls b-vpls>sap>spb>level<br>config>service>vpls b-vpls>spoke-sdp>spb>level                                                                                       |
| <b>Description</b> | This configures metric for this SPB interface SAP/spoke-sdp. This command is valid only for interfaces on control B-VPLS.<br><b>Values</b> 1 — 16,777,215<br><b>Default</b> 1000 |

## hello-interval

|                    |                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hello-interval</b> <i>seconds</i><br><b>no hello-interval</b>                                                                                                                                                                           |
| <b>Context</b>     | config>service>vpls b-vpls>sap>spb>level<br>config>service>vpls b-vpls>spoke-sdp>spb>level                                                                                                                                                 |
| <b>Description</b> | 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.<br>The no form of the command to reverts to the default value. |
| <b>Default</b>     | 3 — Hello interval default for the designated intersystem.<br>9 — Hello interval default for non-designated intersystems.                                                                                                                  |
| <b>Parameters</b>  | <i>seconds</i> — The hello interval in seconds expressed as a decimal integer.<br><b>Values</b> 1 — 20000                                                                                                                                  |

## hello-multiplier

|                    |                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hello-multiplier</b> <i>multiplier</i><br><b>no hello-multiplier</b>                                                                                                                                                            |
| <b>Context</b>     | config>service>vpls b-vpls>sap>spb>level<br>config>service>vpls b-vpls>spoke-sdp>spb>level                                                                                                                                         |
| <b>Description</b> | 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.<br>The no form of the command reverts to the default value. |
| <b>Default</b>     | 3 — SPB can miss up to 3 hello messages before declaring the adjacency down.                                                                                                                                                       |
| <b>Parameters</b>  | <i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer.<br><b>Values</b> 2 — 100                                                                                                                 |

## mrp

|                    |                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mrp</b>                                                                                                      |
| <b>Context</b>     | config>service>vpls<br>config>service>vpls>mesh-sdp<br>config>service>vpls>sap<br>config>service>vpls>spoke-sdp |
| <b>Description</b> | This command configures Multiple Registration Protocol (MRP) parameters. MRP is valid only under B-VPLS.        |

## attribute-table-size

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] attribute-table-size</b> <i>value</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>vpls>mrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | 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.<br><br>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. |
| <b>Default</b>     | maximum number of attributes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>value</i> — [1-2048] for 7450 ESS-6, 7450 ESS-7, 7450 ESS-12, 7750 SR-7, or 7750 SR-12<br>[1-1023] for 7450 ESS-1 or 7750 SR-1                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## attribute-table-high-wmark

|                    |                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] attribute-table-high-wmark</b> <i>high-water-mark</i>                                                  |
| <b>Context</b>     | config>service>vpls>mrp                                                                                        |
| <b>Description</b> | This command specifies the percentage filling level of the MMRP attribute table where logs and traps are sent. |
| <b>Default</b>     | 95%                                                                                                            |
| <b>Parameters</b>  | <i>high-water-mark</i> — 1%-100%                                                                               |

## attribute-table-low-wmark

|                |                                                             |
|----------------|-------------------------------------------------------------|
| <b>Syntax</b>  | <b>[no] attribute-table-low-wmark</b> <i>low-water-mark</i> |
| <b>Context</b> | config>service>vpls>mrp                                     |

## VPLS Service Commands

|                    |                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | 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. |
| <b>Default</b>     | 90%                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>low-water-mark</i> — 1%-100%                                                                                                                                                                                                               |

### flood-time

|                    |                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>flood-time</b> <i>flood-time</i><br><b>no flood-time</b>                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>vpls>mrp                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | 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. When “no flood-time” is executed, flooding behavior is disabled. |
| <b>Default</b>     | no flood-time                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <i>flood-time</i> — Specifies the MRP flood time, in seconds.<br><b>Values</b> 3 — 600                                                                                                                                                                                                                            |

### join-time

|                    |                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] <b>join-time</b> <i>value</i>                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>vpls>sap>mrp<br>config>service>vpls>spoke-sdp>mrp<br>config>service>vpls>mesh-sdp>mrp                                                                                                                                                                                    |
| <b>Description</b> | 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. |
| <b>Default</b>     | 2                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>value</i> — [1-10] tenths of a second                                                                                                                                                                                                                                                |

### leave-time

|               |                                     |
|---------------|-------------------------------------|
| <b>Syntax</b> | [no] <b>leave-time</b> <i>value</i> |
|---------------|-------------------------------------|

**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* — [30-60] tenths of a second

## 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  $\text{LeaveAllTime} < T < 1.5 * \text{leave-all-time}$  when it is started. Refer to IEEE 802.1ak-2007 section 10.7.4.3.

**Default** 100

**Parameters** *value* — [60-300] tenths of a second

## 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 Transmitting Timer is set to one second when it is started.

## VPLS Service Commands

|                   |                                            |
|-------------------|--------------------------------------------|
| <b>Default</b>    | 10                                         |
| <b>Parameters</b> | <i>value</i> — [10-100] tenths of a second |

### periodic-timer

|                    |                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] periodic-timer</b>                                                                           |
| <b>Context</b>     | config>service>vpls>sap>mrp<br>config>service>vpls>spoke-sdp>mrp<br>config>service>vpls>mesh-sdp>mrp |
| <b>Description</b> | This command enables or disables the Periodic Transmission Timer.                                    |
| <b>Default</b>     | disabled                                                                                             |

### send-flush-on-failure

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] send-flush-on-failure</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>service>vpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>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 “oper-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.</p> <p>This feature cannot be enabled on management VPLS.</p> |
| <b>Default</b>     | no send-flush-on-failure                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

### pbb

|                    |                                                              |
|--------------------|--------------------------------------------------------------|
| <b>Syntax</b>      | <b>pbb</b>                                                   |
| <b>Context</b>     | config>service<br>config>service>vpl<br>config>service>epipe |
| <b>Description</b> | This command configures global PBB parameters.               |

### mac-name

|               |                                                 |
|---------------|-------------------------------------------------|
| <b>Syntax</b> | <b>mac-name</b> <i>name</i> <i>ieee-address</i> |
|---------------|-------------------------------------------------|

**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-smac

- Syntax** **backbone-smac** *ieee-address*
- Context** config>service>pbb>source-bmac
- Description** This command configures the backbone source MAC address used for PBB. This command allows a per B-VPLS control of the B-SMAC and the B-Mcast MAC. All I-VPLS provisioned under this B-VPLS will share the provisioned value.
- Default** backbone-smac address is chassis MAC address
- Parameters** *ieee-address* — Specifies the backbone source MAC address.

## 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.

## VPLS Service Commands

- 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 — 16777215

### force-qtag-forwarding

- Syntax** **[no] force-qtag-forwarding**
- Context** config>service>vpls ivpls>pbb
- Description** This command forces the addition of a IEEE 802.1q tag after the Customer MAC (CMAC) address 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 as 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.
- The **no** form of this command disables the command.

### source-bmac

- Syntax** **source-bmac** *ieee-address*
- Context** config>service>vpls bvpls>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 bvpls>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

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mac-notification</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>vpls bvpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command controls the settings for the MAC notification message.</p> <p>The mac-notification message must be generated under the following events:</p> <ol style="list-style-type: none"> <li>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.</li> <li>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.</li> <li>3. 1st SAP on an active MC-LAG is associated (via IVPLS/Epipe) with the BVPLS</li> <li>4. The link between IVPLS/Epipe and BVPLS is configured and there are I-SAPs configured on an active MC-LAG link.</li> </ol> <p>The MAC notification is not sent for the following events:</p> <ol style="list-style-type: none"> <li>1. Change of source-bmac or source-bmac-lsb</li> <li>2. On changes of use-sap-bmac parameter</li> <li>3. If MC-LAG peering is not (initialized and in sync).</li> </ol> |

## interval

|                    |                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] interval <i>value</i></b>                                                                                                                                                                                 |
| <b>Context</b>     | config>service>vpls>pbb>mac-notification                                                                                                                                                                          |
| <b>Description</b> | This command controls the frequency of subsequent MAC notification messages.                                                                                                                                      |
| <b>Default</b>     | Inherits the chassis level configuration from config>service>mac-notification                                                                                                                                     |
| <b>Parameters</b>  | <p><i>value</i> — Specifies the frequency of subsequent MAC notification messages.</p> <p><b>Values</b> 100 ms – 10 sec, in increments of 100 ms up to 1 sec and then in increments of 1 second up to 10 sec.</p> |

## renotify

|                    |                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>renotify <i>value</i></b><br><b>no renotify</b>                                                                                                                                                                                               |
| <b>Context</b>     | config>service>vpls>pbb>mac-notification                                                                                                                                                                                                         |
| <b>Description</b> | 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. |

## VPLS Service Commands

**Default** no renotify

**Parameters** *value* — Specifies the time interval between re-notification in seconds.

**Values** 240—840 seconds

### 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—10

**Default** Inherits the chassis level configuration from config>service>mac-notification

### shutdown

**Syntax** [no] shutdown

**Context** config>service>vpls bvpls

**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..2147483648

*isid* — Specifies the ISID.

**Values** 0..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** onfig>service>vpls>pbb>bvpls>igmp-snooping  
 onfig>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** on [page 994](#).

**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

## VPLS Service Commands

config>service>vpls>pbb>bvpls>sdp>mld-snooping

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command specifies whether a multicast router is attached behind this SAP or spoke-SDP.</p> <p>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.</p> <p>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.</p> <p>Note that 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.</p> <p>If the <b>send-queries</b> command is enabled on this SAP or spoke-SDP, the <b>mrouter-port</b> parameter can not be set.</p> |
| <b>Default</b>     | no mrouter-port                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## sdp

|                    |                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] sdp sdp-id:vc-id</b>                                                                                                          |
| <b>Context</b>     | config>service>vpls>pbb>backbone-vpls                                                                                                 |
| <b>Description</b> | This command configures attributes of a SDP binding on the B-VPLS service.                                                            |
| <b>Parameters</b>  | <i>sdp-id</i> — Specifies the SDP ID.<br><b>Values</b> 1..17407<br><i>vc-id</i> — Specifies the VC ID.<br><b>Values</b> 1..4294967295 |

## stp

|                    |                                                             |
|--------------------|-------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] stp</b>                                             |
| <b>Context</b>     | config>service>Vpls>pbb>backbone-vpls                       |
| <b>Description</b> | This command enables or disable STP through B-VPLS service. |

## force-qtag-forwarding

|                |                                   |
|----------------|-----------------------------------|
| <b>Syntax</b>  | <b>[no] force-qtag-forwarding</b> |
| <b>Context</b> | config>service>vpls ivpls>pbb     |

```
config>service>epipe>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

## 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 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

## 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.

## mac-notification

**Syntax** **mac-notification**

**Context** config>service>pbb

**Description** This command controls the settings for the MAC notification messages.

## interval

|                    |                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] interval</b> <i>value</i>                                                                                   |
| <b>Context</b>     | config>service>pbb>mac-notification                                                                                 |
| <b>Description</b> | This command controls the frequency of subsequent MAC notification messages.                                        |
| <b>Default</b>     | 100 ms                                                                                                              |
| <b>Parameters</b>  | <i>value</i> — Specifies the frequency of subsequent MAC notification messages.                                     |
|                    | <b>Values</b> 100 ms – 10 sec, in increments of 100 ms up to 1 sec and then in increments of 1 second up to 10 sec. |

## count

|                    |                                                                                     |
|--------------------|-------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] count</b> <i>value</i>                                                      |
| <b>Context</b>     | config>service>pbb>mac-notification                                                 |
| <b>Description</b> | This command configures how often MAC notification messages are sent.               |
| <b>Parameters</b>  | <i>value</i> — Specifies, in seconds, how often MAC notification messages are sent. |
|                    | <b>Values</b> 1-10                                                                  |
|                    | <b>Default</b> 3                                                                    |

## epipe

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>epipe</b> <i>service-id</i> <b>customer</b> <i>customer-id</i> [ <b>vpn</b> <i>vpn-id</i> ] [ <b>vc-switching</b> ] [ <b>create</b> ]<br><b>epipe</b> <i>service-id</i><br><b>no epipe</b> <i>service-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>service                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>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 .</p> <p>No MAC learning or filtering is provided on an Epipe.</p> <p>When a service is created, the <b>customer</b> keyword and <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> must already exist having been created using the <b>customer</b> 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.</p> <p>Once a service is created, the use of the <b>customer</b> <i>customer-id</i> is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect <i>customer-id</i> specified will result in an error.</p> <p>By default, no epipe services exist until they are explicitly created with this command.</p> |

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.

*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 on which this service is defined.

**Values** 1 — 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 — 2147483647

**create** — Keyword used to create the service instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.



---

## PBB Show Commands

### eth-cfm

|                    |                                                |
|--------------------|------------------------------------------------|
| <b>Syntax</b>      | <b>eth-cfm</b>                                 |
| <b>Context</b>     | show                                           |
| <b>Description</b> | This command displays 802.1ag CFM information. |

### association

|                    |                                                          |
|--------------------|----------------------------------------------------------|
| <b>Syntax</b>      | <b>association</b> [ <i>ma-index</i> ] [ <b>detail</b> ] |
| <b>Context</b>     | show>eth-cfm                                             |
| <b>Description</b> | Shows association information.                           |
| <b>Parameters</b>  | <i>ma-index</i> — Specifies the MA index value.          |

**Values** 1 — 4294967295

**detail** — Displays all association detail.

|               |                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Output</b> | <pre>*A:alcag1-R6# show eth-cfm association ===== CFM Association Table ===== Md-index  Ma-index  Name                CCM-interval  Bridge-id ----- 1          1         ivpls                1              5000 ===== *A:alcag1-R6#</pre> |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## cfm-stack-table

**Syntax** **cfm-stack-table**  
**cfm-stack-table** [{all-ports|all-sdps|all-virtuals}] [level 0..7] [direction up|down]  
**cfm-stack-table port** *port-id* [vlan *qtag*[. *qtag*]] [level 0..7] [direction up|down]  
**cfm-stack-table sdp** *sdp-id[:vc-id]* [level 0..7] [direction up|down]  
**cfm-stack-table virtual** *service-id* [level 0..7]  
**cfm-stack-table facility** [{all-ports|all-lags|all-lag-ports|all-tunnel-meps|all-router-interfaces}] [level 0..7] [direction up|down]  
**cfm-stack-table facility collect-lmm-stats**  
**cfm-stack-table facility lag** *id*[tunnel 1..4094] [level 0..7] [direction up|down]  
**cfm-stack-table facility port** *id* [level 0..7] [direction up|down]  
**cfm-stack-table facility router-interface** *ip-int-name* [level 0..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** **port** *port-id* — Displays the bridge port or aggregated port on which MEPs or MHFs are configured.  
**vlan** *vlan-id* — Displays the associated VLAN ID.  
**level** — Display the MD level of the maintenance point.  
**Values** 0 — 7

**direction up (U)| down (D)** — Displays the direction in which the MP faces on the bridge port.

**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** *sdp-id[:vc-id]* — Displays CFM stack table information for the specified SDP.

**service** *service-id* — Displays CFM stack table information for the specified SDP.

### 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

=====
CFM SAP Stack Table
=====
Sap Lvl Dir Md-index Ma-index MepId Mac-address Defect

1/1/6:20.0 4 B 14 803 MIP d8:1c:01:01:00:06 -----
1/1/6:3000.1001 4 B 14 800 MIP 00:00:00:00:00:28 -----
1/1/6:2000.1002 4 B 14 802 MIP d8:1c:01:01:00:06 -----
1/1/6:0.* 4 B 14 805 MIP d8:1c:01:01:00:06 -----
1/1/9:300 2 U 12 300 28 00:00:00:00:00:28 -----
1/1/9:401 2 U 12 401 28 00:00:00:00:00:28 -----
1/1/9:600 2 U 12 600 28 00:00:00:00:00:28 -----
```

```

1/1/9:600 5 B 15 666 MIP 00:10:11:00:00:1c -----
1/1/10:4.* 2 U 12 4 28 00:00:00:00:00:28 --C----
1/1/10:1000.* 5 U 15 1000 28 00:00:00:00:00:28 -----
1/1/10:1001.* 5 U 15 1001 28 00:00:00:00:00:28 -----
=====

```

CFM Ethernet Tunnel Stack Table

```

Eth-tunnel Lvl Dir Md-index Ma-index MepId Mac-address Defect

```

No Matching Entries

CFM Ethernet Ring Stack Table

```

Eth-ring Lvl Dir Md-index Ma-index MepId Mac-address Defect

```

No Matching Entries

CFM Facility Port Stack Table

```

Port Tunnel Lvl Dir Md-index Ma-index MepId Mac-address Defect

```

```

1/2/4 0 0 D 10 1 28 00:00:00:00:00:28 -----
=====

```

CFM Facility LAG Stack Table

```

Lag Tunnel Lvl Dir Md-index Ma-index MepId Mac-address Defect

```

No Matching Entries

CFM Facility Tunnel Stack Table

```

Port/Lag Tunnel Lvl Dir Md-index Ma-index MepId Mac-address Defect

```

No Matching Entries

CFM Facility Interface Stack Table

```

Interface Lvl Dir Md-index Ma-index MepId Mac-address Defect

```

```

v28-v33 1 D 11 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

```

```

1/1/6:20.*

```

## PBB Show Commands

```

 21 4 B 14 804 MIP d8:1c:01:01:00:06 -----
=====

CFM SDP Stack Table
=====
Sdp Lvl Dir Md-index Ma-index MepId Mac-address Defect

1:1000 4 D 14 1000 28 00:00:00:00:00:28 -----
2:777 4 D 14 777 28 d8:1c:ff:00:00:00 -----
400:800 4 B 14 800 MIP 00:00:00:00:01:28 -----
=====

CFM Virtual Stack Table
=====
Service Lvl Dir Md-index Ma-index MepId Mac-address Defect

No Matching Entries
=====

```

## 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 — 4294967295
  - ma-index* — Specifies the MA index value.
    - Values** 1 — 4294967295
  - all-associations** — Displays information all maintenance associations.
  - detail** — Displays detailed information.

### Sample Output

```

*A:alcag1-R6# show eth-cfm domain
=====
CFM Domain Table
=====
Md-index Level Name Format

1 4 ivpls charString
=====
*A:alcag1-R6#

*A:alcag1-R6# show eth-cfm mep 51 domain 1 association 1

Mep Information

Md-index : 1 Direction : Up

```

## PBB Show Commands

```
Ma-index : 1
MepId : 51
IfIndex : 38043648
FngState : fngReset
LowestDefectPri : allDef
Defect Flags : None
Mac Address : 00:ae:ae:ae:ae:ae
CcmTx : 775
CcmLastFailure Frame:
 None
XconCcmFailure Frame:
 None
*A:alcag1-R6#
Admin : Enabled
CCM-Enable : Enabled
PrimaryVid : 5
HighestDefect : none
CcmLtmPriority : 7
CcmSequenceErr : 0
```

## 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 — 8191
  - md-index* — Specifies the maintenance domain (MD) index value.
    - Values** 1 — 4294967295
  - ma-index* — Specifies the MA index value.
    - Values** 1 — 4294967295
  - loopback** — Displays loopback information for the specified MEP.
  - linktrace** — Displays linktrace information for specified MEP.

### Sample Output

```
*A:alcag1-R6# oam eth-cfm loopback 00:af:af:af:af:af mep 51 domain 1 association 1
eth-cfm Loopback Test Initiated: Mac-Address: 00:af:af:af:af:af, out sap: 1/2/9:5
Sent 1 packets, received 1 packets [0 out-of-order, 0 Bad Msdu] -- OK
*A:alcag1-R6#

*A:alcag1-R6# oam eth-cfm linktrace 00:af:af:af:af:af mep 51 domain 1 association 1
Index Ingress Mac Egress Mac Relay Action

1 00:00:00:00:00:00 00:AF:AF:AF:AF:AF rlyHit terminate

No more responses received in the last 5 seconds.
*A:alcag1-R6#
```

## 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.

### 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**

### 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.

### Sample

```

*A:Dut-B# show service pbb mac-name
=====
MAC Name Table
=====
MAC-Name MAC-Address

test 00:03:03:03:03:02

```

## PBB Show Commands

```
=====
*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

|                    |                                                             |
|--------------------|-------------------------------------------------------------|
| <b>Syntax</b>      | <b>id service-id</b>                                        |
| <b>Context</b>     | show>service                                                |
| <b>Description</b> | This command displays information on a specific service ID. |

### 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** \*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..65535

## mmp

**Syntax** **mmp 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.

## PBB Show Commands

**Parameters** *ieee-address* — Hex string: xx:xx:xx:xx:xx:xx. OR xx-xx-xx-xx-xx-xx

**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

|                    |                                                  |
|--------------------|--------------------------------------------------|
| <b>Syntax</b>      | <b>adjacency [detail]</b>                        |
| <b>Context</b>     | show>service>id>spb                              |
| <b>Description</b> | This command displays SPB adjacency information. |
| <b>Parameters</b>  | <i>detail</i> — Show detailed information.       |
| <b>Output</b>      | <b>Sample Ouput</b>                              |

```

=====
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

|                    |                                             |
|--------------------|---------------------------------------------|
| <b>Syntax</b>      | <b>base</b>                                 |
| <b>Context</b>     | show>service>id>spb                         |
| <b>Description</b> | This command displays SPB base information. |
| <b>Output</b>      | <b>Sample Ouput</b>                         |

```

*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

```

## PBB Show Commands

```
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** **Sample Ouput**

```
*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** **Sample Ouput**

```
*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 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.

**Output** **Sample Ouput**

```
*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 SPBcontrol service FID information.

**Parameters** *fid* — A user service FID may be specified. All user service FIDs are displayed if the FID is not specified.

**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** — Specifies user VPLS Shortest Path Bridging (SPB) multicast forwarding data-base (Mfib) information.

**mfib**

**group-mac** *ieee-address* — Specifies the 48-bit IEEE 802.3 group MAC address.

**isid** *isid* — Specifies the value of ISID of the group MAC address of this entry.

**Output** **Sample Ouput**

```
*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 User SvcId SdpBind

```

## PBB Show Commands

```

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** 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

|                    |                                       |
|--------------------|---------------------------------------|
| <b>Syntax</b>      | <b>interface</b>                      |
| <b>Context</b>     | show>service>id>spb                   |
| <b>Description</b> | This command displays SPB interfaces. |
| <b>Output</b>      | <b>Sample Ouput</b>                   |

```
*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

|                    |                                                                   |
|--------------------|-------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mfib [group-mac <i>ieee-address</i>][<i>isis isid</i>]</b>     |
| <b>Context</b>     | show service id <svclid> spb                                      |
| <b>Description</b> | This command displays multicast forwarding data-base information. |
| <b>Parameters</b>  | <i>group-mac</i> — Optional IEEE group MAC format:                |

mac-address: xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

*isid* — Optional I-SID.

Format: 0..16777215

|               |                     |
|---------------|---------------------|
| <b>Output</b> | <b>Sample Ouput</b> |
|---------------|---------------------|

```
*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

|                |                     |
|----------------|---------------------|
| <b>Syntax</b>  | <b>routes</b>       |
| <b>Context</b> | show>service>id>spb |

# PBB Show Commands

**Description** This command displays SPB route information.

**Output** **Sample Ouput**

```
*A:Dut-A# show service id 100001 spb routes
=====
MAC Route Table
=====
Fid MAC NextHop If SysID Ver. Metric

Fwd Tree: unicast

1 00:10:00:01:00:02 sap:1/2/2:1.1 Dut-B 10 10
1 00:10:00:01:00:03 sap:1/2/3:1.1 Dut-C 10 10
1 00:10:00:01:00:04 sap:1/2/2:1.1 Dut-B 10 20
100 00:10:00:02:00:02 sap:1/2/2:1.1 Dut-B 10 10
100 00:10:00:02:00:03 sap:1/2/3:1.1 Dut-C 10 10
100 00:10:00:02:00:04 sap:1/2/3:1.1 Dut-C 10 20

Fwd Tree: multicast

1 00:10:00:01:00:02 sap:1/2/2:1.1 Dut-B 10 10
1 00:10:00:01:00:03 sap:1/2/3:1.1 Dut-C 10 10
1 00:10:00:01:00:04 sap:1/2/2:1.1 Dut-B 10 20
100 00:10:00:02:00:02 sap:1/2/2:1.1 Dut-B 10 10
100 00:10:00:02:00:03 sap:1/2/3:1.1 Dut-C 10 10
100 00:10:00:02:00:04 sap:1/2/3:1.1 Dut-C 10 20

No. of MAC Routes: 12
=====

ISID Route Table
=====
Fid ISID NextHop If SysID Ver.

1 10001 sap:1/2/2:1.1 Dut-B 10
 sap:1/2/3:1.1 Dut-C
100 10002 sap:1/2/2:1.1 Dut-B 10
 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

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**    **Sample Ouput**

```

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**

## PBB Show Commands

**Context** show>service>id>spb  
**Description** This command displays SPF Log information.  
**Output** **Sample Ouput**

```
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** **Sample Ouput**

```
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** **Sample Ouput**

```
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
=====
```

---

## PBB Clear Commands

### counters

|                    |                                                                                |
|--------------------|--------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>counters</b>                                                                |
| <b>Context</b>     | clear>service>statistics>id                                                    |
| <b>Description</b> | This command clears all traffic queue counters associated with the service ID. |

### mesh-sdp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mesh-sdp</b> <i>sdp-id[:vc-id]</i> { <b>all</b>   <b>counters</b>   <b>stp</b>   <b>mrp</b> }                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | clear>service>statistics>id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command clears the statistics for a particular mesh SDP bind.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>sdp-id</i> — Specifies the SDP ID for which to display information.</p> <p><b>Default</b> All SDPs.</p> <p><b>Values</b> 1 — 17407</p> <p><i>vc-id</i> — Displays information about the virtual circuit identifier.</p> <p><b>Values</b> 1 — 4294967295</p> <p><b>all</b> — Clears all queue statistics and STP statistics associated with the SDP.</p> <p><b>counters</b> — Clears all queue statistics associated with the SDP.</p> <p><b>stp</b> — Clears all STP statistics associated with the SDP.</p> <p><b>mrp</b> — Clears all MRP statistics associated with the SDP.</p> |

### mrp

|                    |                                                            |
|--------------------|------------------------------------------------------------|
| <b>Syntax</b>      | <b>mrp</b>                                                 |
| <b>Context</b>     | clear>service>statistics>id                                |
| <b>Description</b> | This command clears all MRP statistics for the service ID. |

### spoke-sdp

|                |                                                                                                                 |
|----------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>spoke-sdp</b> <i>sdp-id[:vc-id]</i> { <b>all</b>   <b>counters</b>   <b>stp</b>   <b>l2pt</b>   <b>mrp</b> } |
| <b>Context</b> | clear>service>statistics>id                                                                                     |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command clears statistics for the spoke SDP bound to the service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <i>sdp-id</i> — The spoke SDP ID for which to clear statistics.<br><b>Values</b> 1 — 17407<br><i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset.<br><b>Values</b> 1 — 4294967295<br><b>all</b> — Clears all queue statistics and STP statistics associated with the SDP.<br><b>counters</b> — Clears all queue statistics associated with the SDP.<br><b>stp</b> — Clears all STP statistics associated with the SDP.<br><b>l2pt</b> — Clears all L2PT statistics associated with the SDP.<br><b>mrp</b> — Clears all MRP statistics associated with the SDP. |

## sap

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sap <i>sap-id</i> {all   counters   stp  l2pt   mrp}</b>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | clear>service>statistics>id                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command clears statistics for the SAP.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>sap-id</i> — The SAP ID for which to clear statistics.<br><b>all</b> — Clears all queue statistics and STP statistics associated with the SAP.<br><b>counters</b> — Clears all queue statistics associated with the SAP.<br><b>stp</b> — Clears all STP statistics associated with the SAP.<br><b>l2pt</b> — Clears all L2PT statistics associated with the SAP.<br><b>mrp</b> — Clears all MRP statistics associated with the SAP. |

## stp

|                    |                                                         |
|--------------------|---------------------------------------------------------|
| <b>Syntax</b>      | <b>stp</b>                                              |
| <b>Context</b>     | clear>service>statistics>id                             |
| <b>Description</b> | Clears all spanning tree statistics for the service ID. |

## PBB Debug Commands

### mrp

|                    |                                                    |
|--------------------|----------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mrp</b>                                    |
| <b>Context</b>     | debug>service>id                                   |
| <b>Description</b> | This command enables and configures MRP debugging. |

### all-events

|                    |                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>all-events</b>                                                                                                                                                   |
| <b>Context</b>     | debug>service>id>mrp                                                                                                                                                |
| <b>Description</b> | 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

|                    |                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] applicant-sm</b>                                                                                                                               |
| <b>Context</b>     | debug>service>id>mrp                                                                                                                                   |
| <b>Description</b> | This command enables debugging of the applicant state machine.<br>The <b>no</b> form of the command disables debugging of the applicant state machine. |

### leave-all-sm

|                    |                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] leave-all-sm</b>                                                                                                                               |
| <b>Context</b>     | debug>service>id>mrp                                                                                                                                   |
| <b>Description</b> | This command enables debugging of the leave all state machine.<br>The <b>no</b> form of the command disables debugging of the leave all state machine. |

### mmp-mac

|                |                                         |
|----------------|-----------------------------------------|
| <b>Syntax</b>  | <b>[no] mmp-mac <i>ieee-address</i></b> |
| <b>Context</b> | 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 zeroes)

## 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* — See [Common CLI Command Descriptions on page 1469](#) for command syntax.

## sdp

|                    |                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] sdp</b> <i>sdp-id:vc-id</i>                                                                                                                                                                                                        |
| <b>Context</b>     | debug>service>id>mrp                                                                                                                                                                                                                       |
| <b>Description</b> | This command filters debug events and only shows events for the particular SDP. The <b>no</b> form of the command removes the debug filter.                                                                                                |
| <b>Parameters</b>  | <i>sdp-id</i> — Specifies the SDP ID for which to display information.<br><b>Default</b> All SDPs.<br><b>Values</b> 1 — 17407<br><i>vc-id</i> — Displays information about the virtual circuit identifier.<br><b>Values</b> 1 — 4294967295 |

# Ethernet Virtual Private Networks (EVPNs)

---

## In This Chapter

This chapter provides information about Ethernet Virtual Private Networks (EVPNs), process overview, and implementation notes.

Topics in this chapter include:

- [Overview on page 1031](#)
- [EVPN for VXLAN Tunnels in a Layer-2 DC GW \(EVPN-VXLAN\) on page 1032](#)
- [EVPN for VXLAN Tunnels in a Layer-2 DC with Integrated Routing Bridging Connectivity on the DC GW on page 1034](#)
- [EVPN for VXLAN Tunnels in a Layer 3 DC with Integrated Routing Bridging Connectivity among VPRNs on page 1035](#)
- [EVPN for VXLAN Tunnels in a Layer 3 DC with EVPN-Tunnel Connectivity among VPRNs on page 1037](#)
- [EVPN for MPLS Tunnels in ELAN Services on page 1039](#)
- [EVPN for PBB over MPLS Tunnels \(PBB-EVPN\) on page 1041](#)
- [VXLAN on page 1042](#)
- [BGP-EVPN Control Plane for VXLAN Overlay Tunnels on page 1053](#)
- [EVPN for VXLAN in VPLS Services on page 1057](#)
- [EVPN for VXLAN in R-VPLS Services on page 1062](#)
- [EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes on page 1064](#)
- [EVPN for VXLAN in EVPN Tunnel R-VPLS Services on page 1068](#)
- [DC GW integration with the Nuage Virtual Services Directory \(VSD\) on page 1073](#)
- [Fully-Dynamic VSD Integration Model on page 1086](#)
- [BGP-EVPN Control Plane for MPLS Tunnels on page 1095](#)

## In This Chapter

- [EVPN for MPLS Tunnels in VPLS Services \(EVPN-MPLS\) on page 1100](#)
- [BGP-EVPN Control Plane for PBB-EVPN on page 1134](#)
- [PBB-EVPN for I-VPLS and PBB Epipe Services on page 1136](#)
- [PBB-EVPN Multi-Homing in I-VPLS and PBB Epipe Services on page 1142](#)
- [ARP/ND Snooping and Proxy Support on page 1155](#)
- [BGP-EVPN MAC-Mobility on page 1161](#)
- [BGP-EVPN MAC-Duplication on page 1162](#)
- [Conditional Static MAC and Protection on page 1164](#)
- [CFM Interaction with EVPN Services on page 1165](#)
- [DC GW Policy Based Forwarding/Routing to an EVPN ESI \(Ethernet Segment Identifier\) on page 1167](#)
- [BGP and EVPN Route Selection for EVPN Routes on page 1175](#)
- [Interaction of EVPN-VXLAN and EVPN-MPLS with Existing VPLS Features on page 1177](#)
- [Interaction of PBB-EVPN with Existing VPLS Features on page 1179](#)
- [Interaction of EVPN-VXLAN with Existing VPRN Features on page 1180](#)
- [Routing Policies for BGP EVPN IP Prefixes on page 1181](#)

## Overview

EVPN is an IETF technology per RFC7432, *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 ELAN services in a similar way to RFC4364 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. Alcatel-Lucent's 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 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 RFC7432.

- **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 ELAN and ELINE services in the WAN.

The PBB-EVPN functionality is standardized in draft-ietf-l2vpn-pbb-evpn.

The 7x50 SR/ESS/XRS EVPN VXLAN implementation is integrated in the Nuage Data Center architecture, where the 7x50 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 7x50 implementation.

## EVPN for VXLAN Tunnels in a Layer-2 DC GW (EVPN-VXLAN)

Figure 110 shows the use of EVPN for VXLAN overlay tunnels on the 7x50 SR/ESS/XRS when it is used as a Layer-2 DC GW.

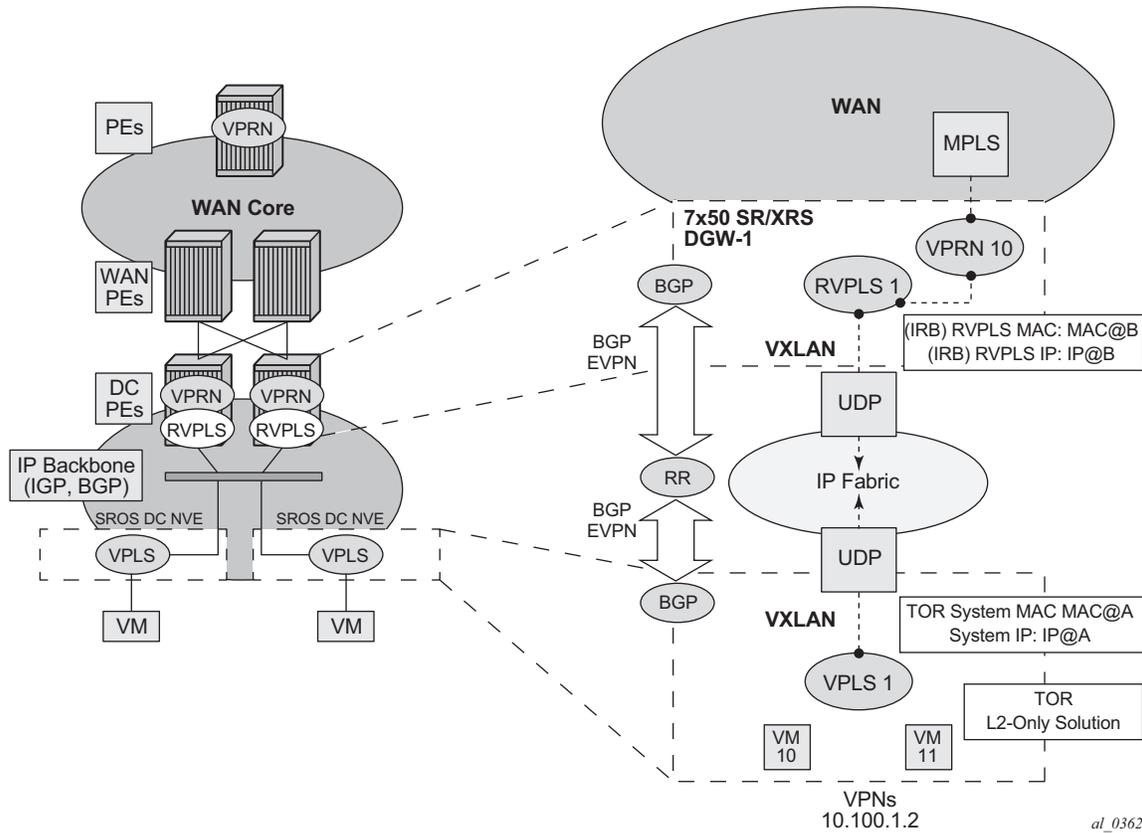


Figure 110: Layer-2 DC PE with VPLS to the WAN

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 Alcatel-Lucent 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 7x50 SROS.

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 110 shows the Layer-2 DC GW function on the 7x50, 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 for VXLAN Tunnels in a Layer-2 DC with Integrated Routing Bridging Connectivity on the DC GW

Figure 111 shows the use of EVPN for VXLAN overlay tunnels on the 7x50 SR/ESS/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.

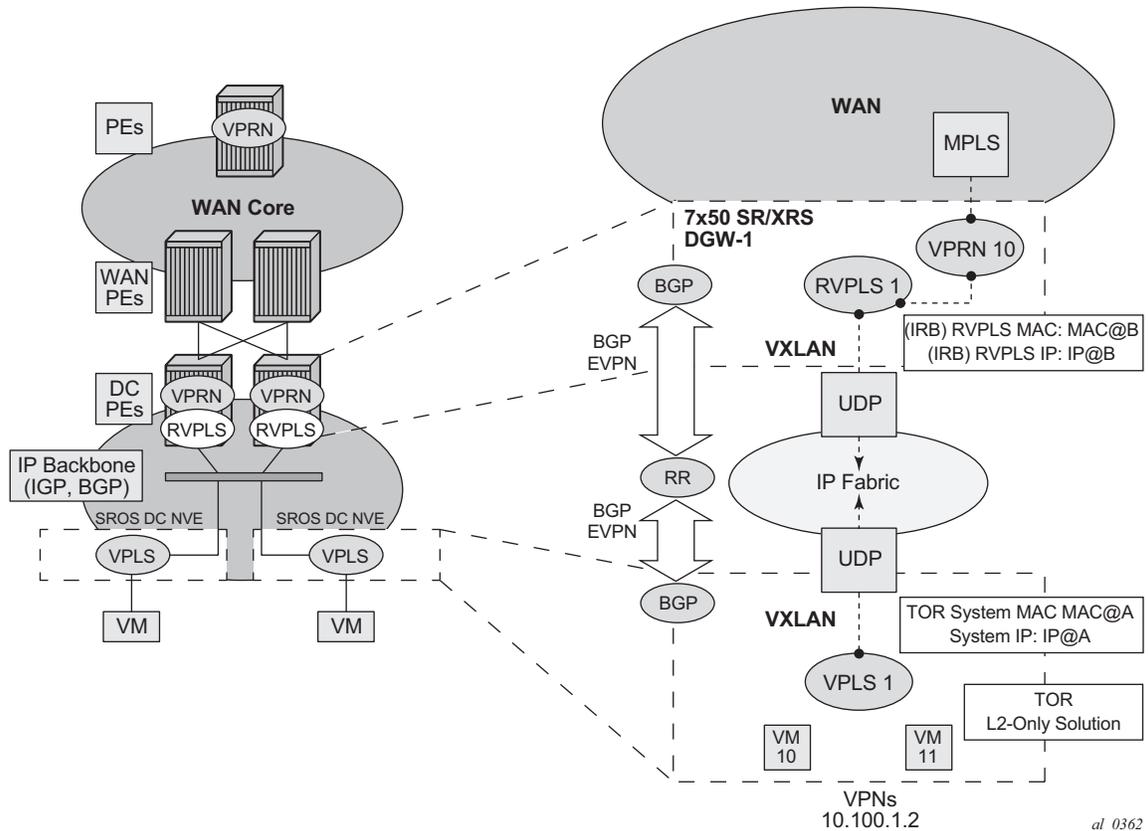


Figure 111: GW IRB on the DC PE for an L2 EVPN/VXLAN DC

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 7x50 supports IPv4 and IPv6 interfaces as default gateways in this scenario.

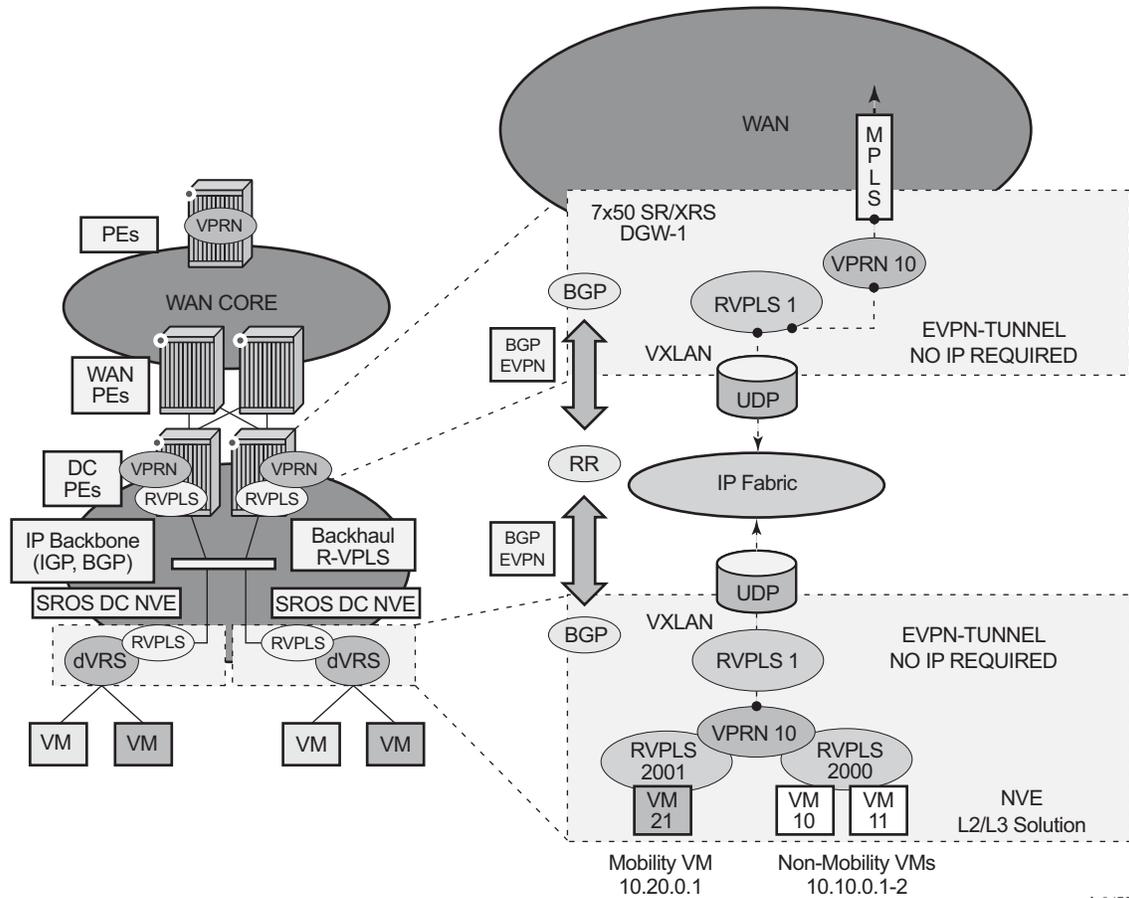


WAN. That VPRN instance will be connected to the VPRNs in the NVEs by means of an IRB (Integrated Routing and Bridging) backhaul R-VPLS. This IRB backhaul R-VPLS provides a scalable solution because it allows L3 connectivity to the WAN without the need for defining all of the subnets in the DC GW.

The 7x50 DC GW supports this 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 EVPN and IP-VPN for IP prefixes is also fully supported.

# EVPN for VXLAN Tunnels in a Layer 3 DC with EVPN-Tunnel Connectivity among VPRNs

Figure 113 shows the use of EVPN for VXLAN tunnels on the 7x50 SR/ESS/XRS, when the DC provides distributed layer-3 connectivity to the DC tenants and the VPRN instances are connected through EVPN tunnels.



al\_0473

**Figure 113: 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 on page 1035](#) 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 7x50.

## EVPN for MPLS Tunnels in ELAN Services

Figure 114 shows the use of EVPN for MPLS tunnels on the 7x50 SR/ESS/XRS. In this case, EVPN is used as the control plane for ELAN services in the WAN.

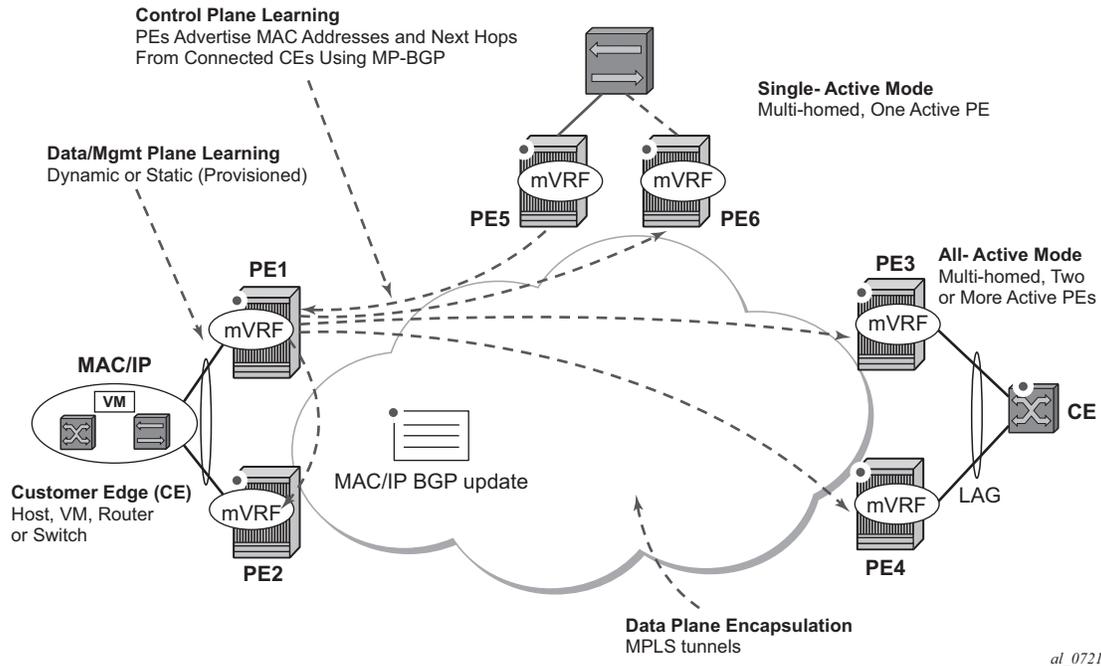


Figure 114: EVPN for MPLS in VPLS Services

EVPN-MPLS is standardized in RFC7432 as an L2VPN technology that can fill the gaps in VPLS for ELAN services. Besides the optimizations introduced by EVPN, a significant number of service providers offering ELAN services today are requesting EVPN for their multi-homing capabilities. EVPN supports all-active multi-homing (per-flow load-balancing multi-homing) in addition to 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 also perceived as a superior technology due to its mass-withdrawal capabilities to speed up convergence in scaled environments.

As well as the superior multi-homing capabilities in EVPN, the technology also provides a number of significant benefits, such as:

- An IP-VPN-like operation and control for ELAN services.

## EVPN for MPLS Tunnels in ELAN 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 7x50 SROS EVPN-MPLS implementation is compliant with RFC7432.

## EVPN for PBB over MPLS Tunnels (PBB-EVPN)

Figure 115 shows the use of EVPN for MPLS tunnels on the 7x50 SR/ESS/XRS. In this case, EVPN is used as the control plane for ELAN services in the WAN.

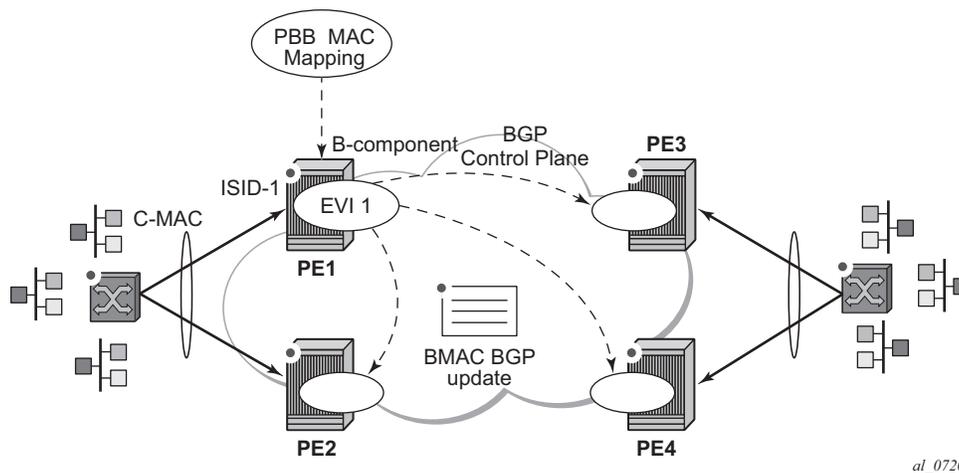


Figure 115: EVPN for PBB over MPLS

EVPN for PBB over MPLS (hereafter called PBB-EVPN) is specified in draft-ietf-l2vpn-pbb-evpn. 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 RFC7432 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 C-MACs, 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 7x50 SROS PBB-EVPN implementation supports PBB-EVPN for I-VPLS and PBB-Epipe services, including single-active and all-active multi-homing.

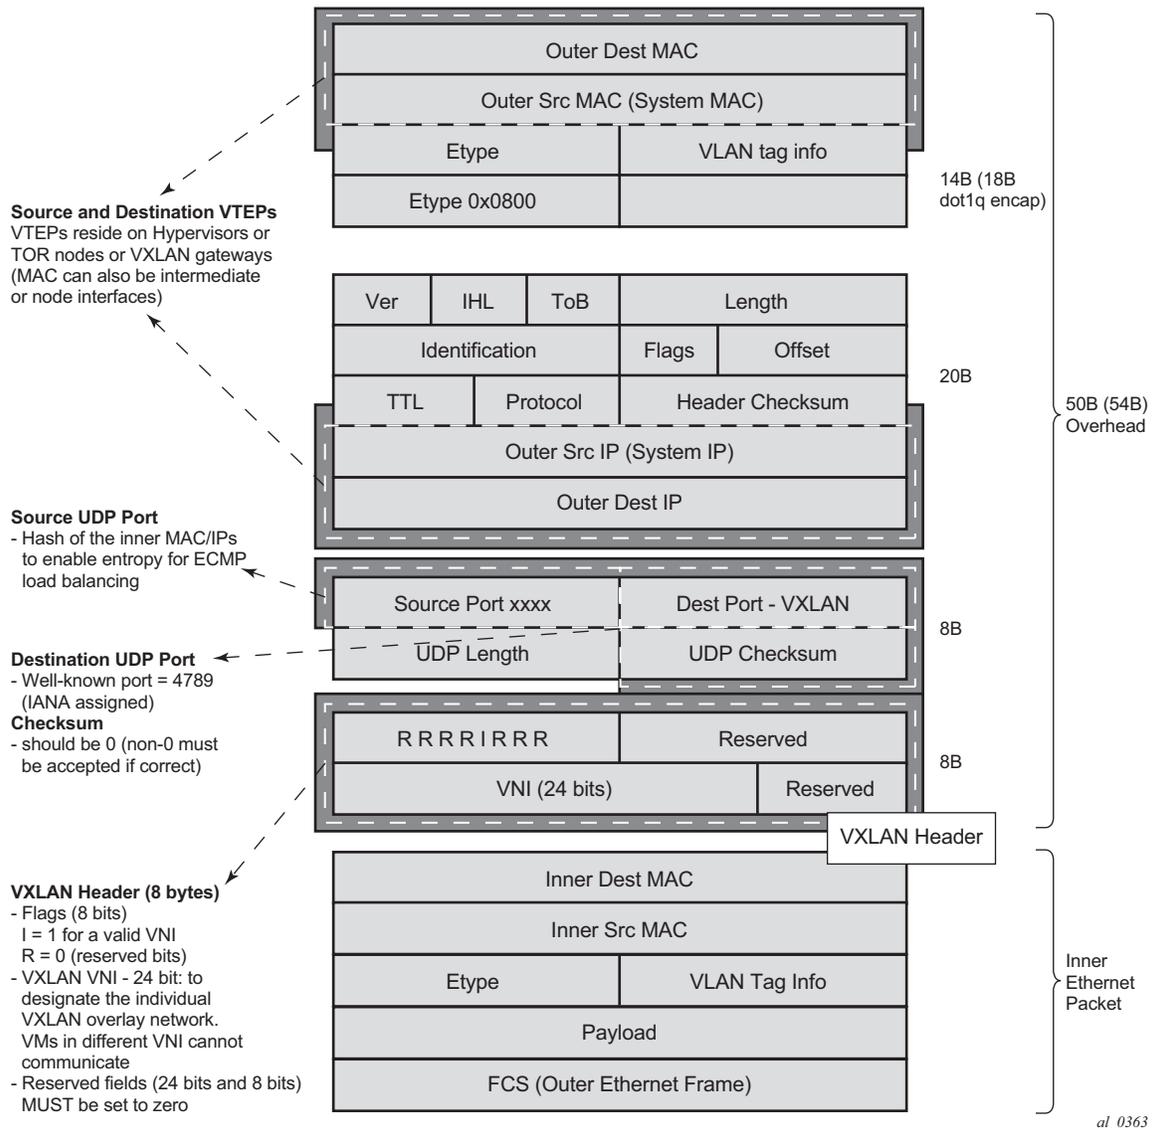
## VXLAN

The 7x50 SROS and Nuage solution for DC supports VXLAN (Virtual eXtensible Local Area Network) overlay tunnels as per RFC7348.

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 116](#) shows an example of the VXLAN encapsulation supported by the Alcatel-Lucent's implementation.



**Figure 116: VXLAN Frame Format**

As shown in Figure 116, 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 7x50 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 7x50 are:

- VXLAN is only supported on network or hybrid ports with null or dot1q encapsulation.
- VXLAN is supported on Ethernet/LAG and POS/APS.
- Only IPv4 unicast addresses are supported as VTEPs.
- Only System IP addresses are supported, as VTEPs, for originating and terminating VXLAN tunnels.

## 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 is not supported on R-VPLS services.
  - In both cases where ECMP is not supported, 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.
  - LAG spraying based on the packet hash is supported in all the cases (VPLS unicast, VPLS BUM, and R-VPLS).
- 

## VXLAN VPLS Tag Handling

The following describes the behavior on the 7x50 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 7x50 will not add any inner VLAN tag on top of the VXLAN encapsulation, and at the ingress network port, the 7x50 will ignore any VLAN tag received and will consider it as part of the payload.
- 

## 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 7x50 will never fragment or reassemble VXLAN packets. In addition, the 7x50 always sets the DF (Do not Fragment) flag in the VXLAN outer IP header.

## VXLAN QoS

VXLAN is a network port encapsulation; therefore, the QoS settings for VXLAN are controlled from the network QoS policies:

- The ingress network QoS policy is used to classify the VXLAN packets based on the outer dot1p (if present, then the DSCP, to yield a FC/profile.
- QoS control of BUM traffic received on VXLAN bindings 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.
- On egress, because the VXLAN adds a new IPv4 header, and the DSCP will be always marked based on the egress network qos policy, there is no need to specify “remarking” in the policy to mark the DSCP.

## 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 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 choose 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. In these examples, the service identifier for the VTEP source is 600; the IP Address of the terminating VTEP is 1.1.1.31; the destination VNI on the terminating VTEP is 31.

```
oam vxlan-ping test-id 1 service 600 dest-vni 31 outer-ip-destination 1.1.1.31 interval
0.1 send-count 10
vxlan-ping destination vxlan-id 31 ip-address 1.1.1.31 reply-mode udp interval 0.1s count
```

## VXLAN Ping

10

```
!!!!!!!!!!!!
---- vxlan-id 31 ip-address 1.1.1.31 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 = 0.912ms, avg = 1.355ms, max = 2.332ms, stddev = 0.425ms
round-trip-delay min = 0.679ms, avg = 0.949ms, max = 1.587ms, stddev = 0.264ms
```

```
oam vxlan-ping test-id 2 service 600 dest-vni 31 outer-ip-destination 1.1.1.31 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 pad-
ding 2000 reflect-pad fc nc profile out
```

```
vxlan-ping destination vxlan-id 31 ip-address 1.1.1.31 reply-mode overlay interval 1s
count 20
```

```
=====
=====
rc=1 Malformed Echo Request Received, rc=2 Overlay Segment Not Present, rc=3 Overlay Seg-
ment Not Operational, rc=4 Ok
=====
=====
```

```
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=1 ttl=255 rtt-time=0.722ms fwd-
time=0.000ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=2 ttl=255 rtt-time=0.750ms fwd-
time=1.508ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=3 ttl=255 rtt-time=0.974ms fwd-
time=0.588ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=4 ttl=255 rtt-time=1.714ms fwd-
time=0.819ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=5 ttl=255 rtt-time=0.799ms fwd-
time=1.776ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=6 ttl=255 rtt-time=0.892ms fwd-
time=0.000ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=7 ttl=255 rtt-time=0.843ms fwd-
time=1.560ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=8 ttl=255 rtt-time=0.825ms fwd-
time=1.253ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=9 ttl=255 rtt-time=0.958ms fwd-
time=0.000ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=10 ttl=255 rtt-time=0.963ms fwd-
time=1.673ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=11 ttl=255 rtt-time=0.929ms fwd-
time=1.697ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=12 ttl=255 rtt-time=0.973ms fwd-
time=1.362ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=13 ttl=255 rtt-time=0.813ms fwd-
time=0.000ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=14 ttl=255 rtt-time=0.887ms fwd-
time=1.676ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=15 ttl=255 rtt-time=1.119ms fwd-
time=0.000ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=16 ttl=255 rtt-time=1.017ms fwd-
time=1.887ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=17 ttl=255 rtt-time=0.873ms fwd-
```

## Ethernet Virtual Private Networks (EVPNs)

```
time=1.746ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=18 ttl=255 rtt-time=1.105ms fwd-
time=0.000ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=19 ttl=255 rtt-time=0.909ms fwd-
time=1.484ms. rc=4
2132 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=20 ttl=255 rtt-time=0.906ms fwd-
time=1.849ms. rc=4

---- vxlan-id 31 ip-address 1.1.1.31 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 = 0.000ms, avg = 0.951ms, max = 1.887ms, stddev = 0.887ms
round-trip-delay min = 0.722ms, avg = 0.948ms, max = 1.714ms, stddev = 0.202ms

oam vxlan-ping test-id 1 service 600 dest-vni 31 outer-ip-destination 1.1.1.31 send-count
10 end-system 00:00:00:00:00:01 interval 0.1

vxlan-ping destination vxlan-id 31 ip-address 1.1.1.31 reply-mode udp end-system
00:00:00:00:00:01 interval 0.1s count 10
1 1 1 1 1 1 1 1 1
---- vxlan-id 31 ip-address 1.1.1.31 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
 10 end-system present(1), 0 end-system not present(2)
forward-delay min = 0.000ms, avg = 0.000ms, max = 0.316ms, stddev = 0.520ms
round-trip-delay min = 0.704ms, avg = 0.855ms, max = 1.151ms, stddev = 0.121ms

oam vxlan-ping test-id 1 service 600 dest-vni 31 outer-ip-destination 1.1.1.31 send-count
10 end-system 00:00:00:00:00:01

vxlan-ping destination vxlan-id 31 ip-address 1.1.1.31 reply-mode udp end-system
00:00:00:00:00:01 interval 1s count 10
=====
=====
rc=1 Malformed Echo Request Received, rc=2 Overlay Segment Not Present, rc=3 Overlay Seg-
ment Not Operational, rc=4 Ok
mac=1 End System Present, mac=2 End System Not Present
=====
=====

92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=1 ttl=255 rtt-time=0.753ms fwd-time=1.240ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=2 ttl=255 rtt-time=0.785ms fwd-time=0.000ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=3 ttl=255 rtt-time=1.425ms fwd-time=2.759ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=4 ttl=255 rtt-time=1.657ms fwd-time=1.659ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=5 ttl=255 rtt-time=0.650ms fwd-time=0.982ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=6 ttl=255 rtt-time=0.894ms fwd-time=0.464ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=7 ttl=255 rtt-time=0.839ms fwd-time=0.581ms.
```

## IGMP-Snooping on VXLAN

```
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=8 ttl=255 rtt-time=0.714ms fwd-time=0.995ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=9 ttl=255 rtt-time=0.798ms fwd-time=0.881ms.
rc=4 mac=1
92 bytes from vxlan-id 31 1.1.1.31: vxlan_seq=10 ttl=255 rtt-time=0.839ms fwd-
time=1.068ms. rc=4 mac=1

---- vxlan-id 31 ip-address 1.1.1.31 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
 10 end-system present, 0 end-system not present
forward-delay min = 0.000ms, avg = 0.978ms, max = 2.759ms, stddev = 0.865ms
round-trip-delay min = 0.650ms, avg = 0.935ms, max = 1.657ms, stddev = 0.314ms
```

---

## 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 mrouter 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/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
```

## Ethernet Virtual Private Networks (EVPNs)

```

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
=====

```

## IGMP-Snooping on VXLAN

```
Multicast FIB, Service 1
=====
Source Address Group Address Sap/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
=====
```

# BGP-EVPN Control Plane for VXLAN Overlay Tunnels

The draft-ietf-bess-evpn-overlay describes EVPN as the control plane for overlay-based networks. The 7x50 supports a subset of the routes and features described in RFC7432 that are required for the DC GW function. In particular, EVPN-specific multi-homing capabilities are not supported. However, multi-homing can be supported by using regular BGP multi-homing based on the L2VPN BGP address family.

Figure 117 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

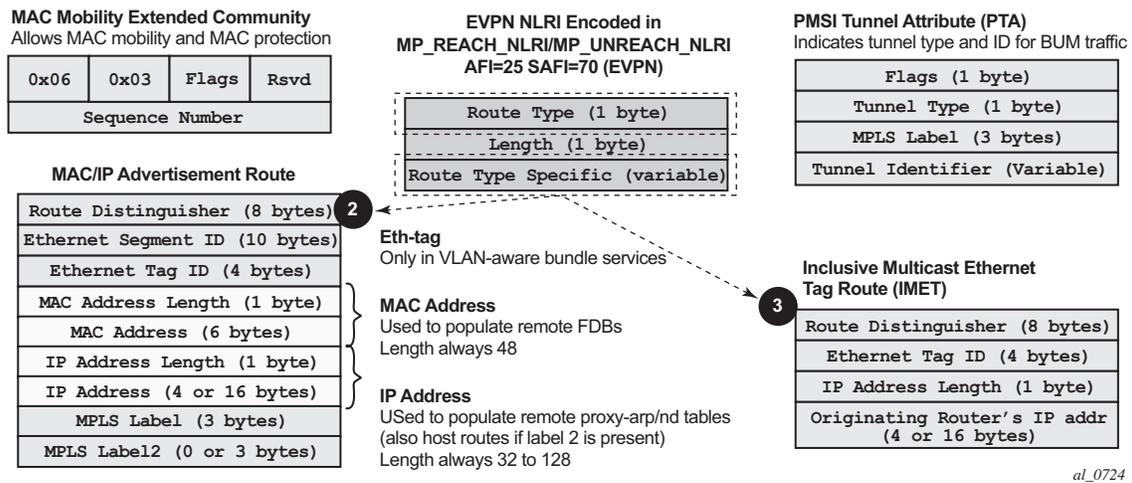


Figure 117: EVPN-VXLAN Required Routes and Communities

## EVPN Route Type 3 – Inclusive Multicast Ethernet Tag Route

Route type 3 is used for setting up the flooding tree (BUM flooding) for a specified VPLS service within the data center. The received inclusive multicast routes will add entries to the VPLS flood list in the 7x50. Only ingress replication is supported over VXLAN.

A route type 3 is generated from the 7x50 per VPLS service as soon as the service is operationally UP and 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 Tag ID: 0.
- 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 not required.
  - MPLS label = Carries the VNI configured in the VPLS service. Only one VNI can be configured per VPLS service.
  - Tunnel end-point = Equal to the originating IP address.

### **EVPN Route Type 2 – MAC/IP Advertisement Route**

The 7x50 will generate this route type for advertising MAC addresses. The 7x50 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 7x50 uses the following fields and values:

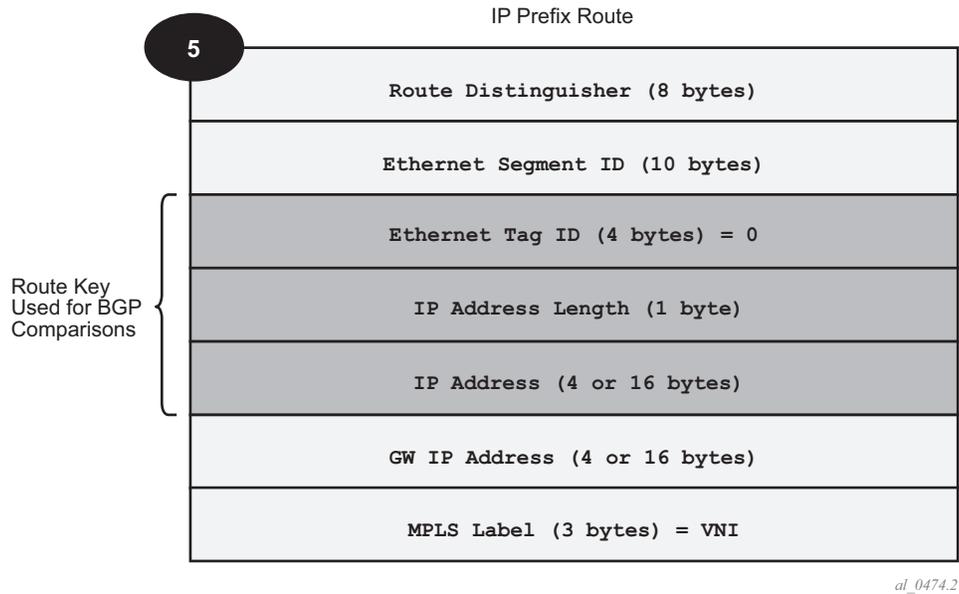
- Route Distinguisher: Taken from the RD of the VPLS service within the BGP context.  
**Note** — The RD can be configured or derived from the **bgp-evpn** evi value.
- Ethernet Segment Identifier (ESI): Value = 0:0:0:0:0:0:0:0.
- Ethernet Tag ID: 0.
- MAC address length: Always 48.
- MAC Address:
  - It will be 00:00:00:00:00:00 for the Unknown MAC route address.
  - It will be different from 00:...:00 for the rest of the advertised MACs.
- 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).
  - 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 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 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 118 shows the IP prefix route or route-type 5.



**Figure 118: EVPN Route-Type 5**

The 7x50 will generate this route type for advertising IP prefixes in EVPN. The 7x50 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 7x50 uses the following fields and values:

## BGP-EVPN Control Plane for VXLAN Overlay Tunnels

- 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.
- 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 will carry 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 will be sent along 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 along with the RFC5512 tunnel encapsulation extended community, with the tunnel type value set to VXLAN.

## EVPN for VXLAN in VPLS Services

The EVPN-VXLAN service is designed around the current VPLS objects and the additional VXLAN construct.

Figure 110 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 7x50 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 7x50s will be configured with a VPLS per tenant that will provide the VXLAN connectivity to the Nuage VPLS instances. On the 7x50, 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 7x50 will advertise new learned MACs (on the SAPs or sdp-bindings) or new conditional static MACs.

- **unknown-mac-route** and **no mac-advertisement** — The 7x50 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 7x50 is the DF (if BGP-MH is configured in the service).
- **unknown-mac-route** and **mac-advertisement** — The 7x50 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 7x50 will use in the EVPN routes generated for the VPLS service.

After the VPLS is configured and operationally UP, the 7x50 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:DGW# show service id 1 vxlan
=====
VPLS VXLAN, Ingress VXLAN Network Id: 1
=====
Egress VTEP, VNI
=====
VTEP Address Egress VNI Num. MACs Mcast Oper State L2 PBR

192.0.2.71 1 1 Yes Up No
192.0.0.72 1 1 Yes 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 7x50s and DC NVEs start advertising MAC addresses, and the 7x50s 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 7x50 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

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
=====
```

---

## 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 is provided by BGP-Multi-homing.

---

## 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.

## Use of bgp-evpn, bgp-ad, and Sites in the Same VPLS Service

- 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 7x50 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.

## Use of the unknown-mac-route

This section describes the behavior of the EVPN-VXLAN service in the 7x50 when the unknown-mac-route and BGP-MH are configured at the same time.

The use of E-VPN, as the control plane of NVO networks in the DC, provides a significant number of benefits as described in 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 E-VPN 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 7x50 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 7x50 can be configured to generate and advertise the unknown-mac-route, the 7x50 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/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.

## EVPN for VXLAN in R-VPLS Services

Figure 111 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 7x50 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 7x50s 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 7x50:

- 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":

```
A:PE73# show router 2 arp
=====
ARP Table (Service: 2)
```

```

=====
IP Address MAC Address Expiry Type Interface

10.10.10.70 d8:46:ff:ff:ff:3e 00h00m00s Evp[I] local
10.10.10.71 d8:47:ff:ff:ff:3e 00h00m00s Evp[I] local
10.10.10.73 d8:49:ff:ff:ff:3e 00h00m00s Oth[I] local

No. of ARP Entries: 3
=====

```

- When a VPLS containing proxy-ARP/ND entries is bound to a VPRN (allow-ip-int-bind) all the proxy-ARP/ND entries are moved to the VPRN ARP/ND table. ARP/ND entries will be also moved to proxy-ARP/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.

## EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes

Figure 112 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 is disabled by default and must be explicitly enabled. The command is tied to the **allow-ip-int-bind** command required for R-VPLS.

**Note** — 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
exit

```

When the above commands are enabled, the 7x50 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 7x50, 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 7x50 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.

## EVPN for VXLAN in EVPN Tunnel R-VPLS Services

Figure 113 shows an L3 connectivity model that optimizes the solution described in [EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes on page 1064](#). 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>interface>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 detail** 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
Nextthop : 10.20.1.2
From : 10.20.1.2
Res. Nextthop : 192.168.19.1
Local Pref. : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
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
Nextthop : 10.20.1.2
From : 10.20.1.2
Res. Nextthop : 192.168.19.1
Local Pref. : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
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
Gateway Address: 00:00:01:00:01:02
Prefix : 3.0.1.6/32
MPLS Label : 262140
Route Tag : 0xb
Neighbor-AS : 111
Orig Validation: N/A
Source Class : 0
Tag : 1
Route Dist. : 10.20.1.2:1
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
=====

```

## EVPN for VXLAN in EVPN Tunnel R-VPLS Services

### 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
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric : None
Connector : None
Community : target:64500:301 bgp-tunnel-encap:VXLAN
Cluster : No Cluster Members
Originator Id : None
Flags : Used Valid Best IGP
Route Source : Internal
AS-Path : No As-Path
EVPN type : IP-PREFIX
ESI : N/A
Gateway Address: fe80::da45:ffff:fe00:*
Prefix : 500::1/128
MPLS Label : 0
Route Tag : 0
Neighbor-AS : N/A
Orig Validation: N/A
Source Class : 0
Add Paths Send : Default
Last Modified : 00h41m17s
Interface Name : int-71-69
Aggregator : None
MED : 0
Peer Router Id : 192.0.2.69
Tag : 301
Route Dist. : 192.0.2.69:301
Dest Class : 0

```

### RIB Out Entries

```

Routes : 1
=====
```

## 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 7x50 DC GW), VSD uses an XMPP (eXtensible Messaging and Presence Protocol) communication channel. The 7x50 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 7x50 using the CLI, SNMP, or other supported interfaces. The VSD will only push a limited number of parameters into the configuration. This 7x50 – 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 7x50 – VSD integration comprises the following building blocks:

- An XMPP interface to the DC XMPP server, through which the 7x50 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 7x50 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.

---

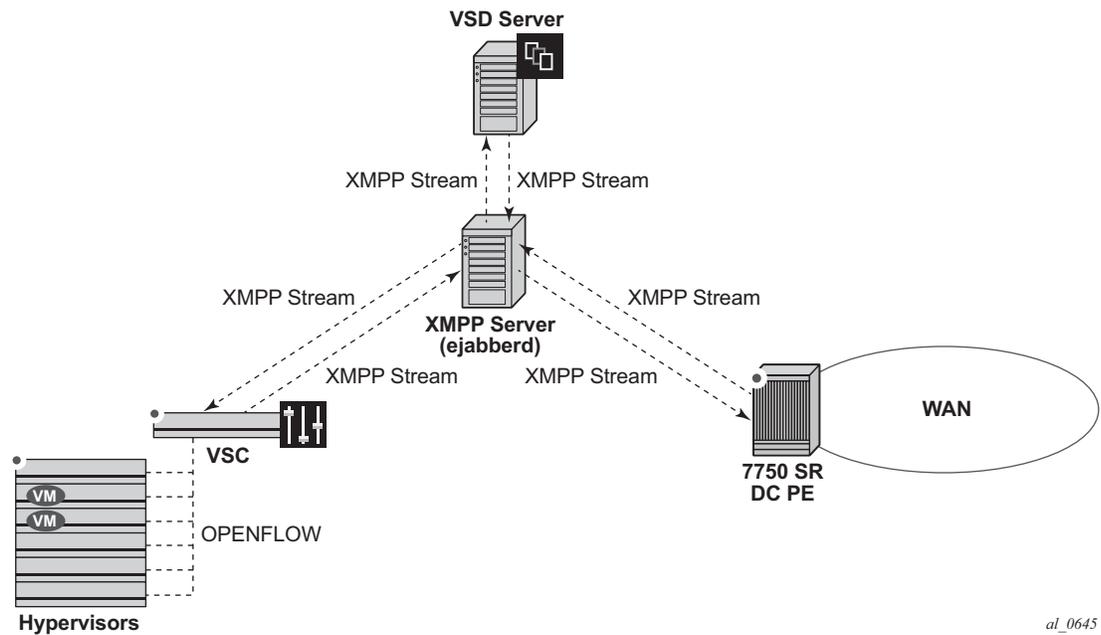
## XMPP Interface on the DC GW

The Extensible Messaging and Presence Protocol 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 7x50 DC GW.

[Figure 119](#) 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.



al\_0645

**Figure 119: Basic XMPP Architecture**

In the Nuage solution, each XMPP client, including the 7x50 SR, 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 7x50, the following command must be added to indicate to which XMPP server address the DC GW has to register, as well as the 7x50's JID:

```
A:Dut-C# configure system xmpp server
- no server <xmpp-server-name>
- server <xmpp-server-name> [domain-name <fqdn>] [username <user-name>]
 [password <password>] [create]
<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.
[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 7x50 acting as an XMPP client. Plain/MD5/anonymous authentication is supported.

## XMPP Interface on the DC GW

- 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.
- 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 7x50 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 7x50 can generate or receive XMPP stanza elements, such as presence and IQ (Information/Query) messages. IQ messages are used between the VSD and the 7x50 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 7x50 must be configured with a VSD **system-id** that uniquely identifies the 7x50 in the VSD:

```
*B:Dut>config>system>vsd# info

system-id "SR12U-46-PE"

```

After the above configuration is complete, the 7x50 will subscribe to a VSD XMPP PubSub node to discover the available VSD servers. Then, the 7x50 will be discovered in the VSD UIs. On the 7x50, 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
```

## 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 120 shows the workflow for the attachment of the WAN services defined on the DC GW to the DC domains.

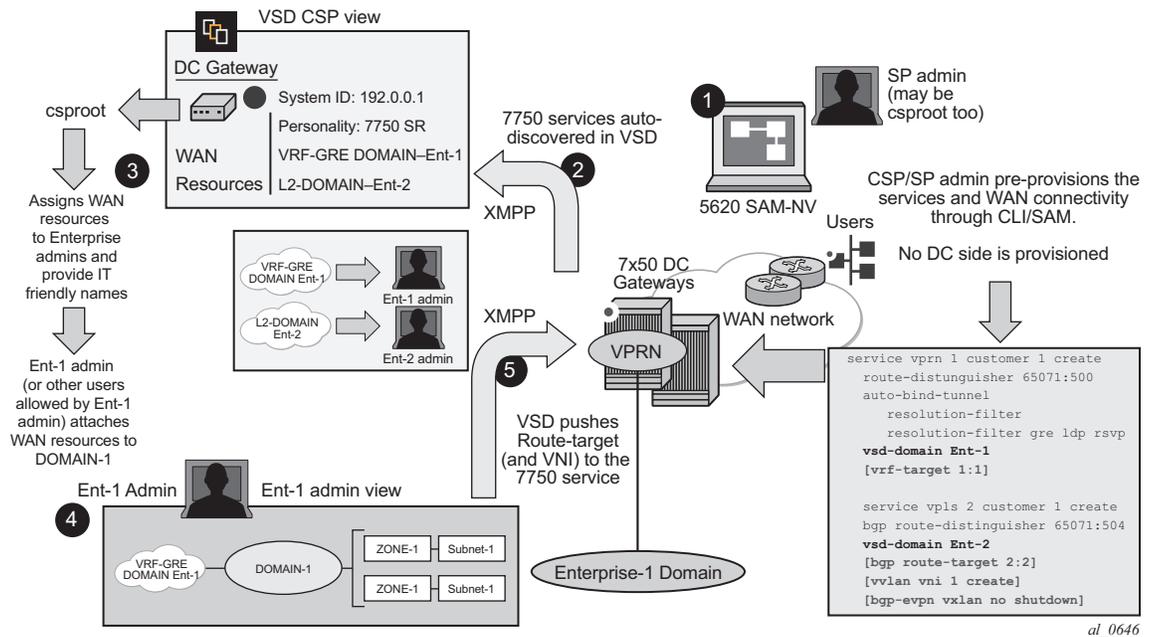


Figure 120: WAN Services Attachment Workflow

The Static-Dynamic VSD integration model can be summarized in the steps shown in Figure 120 and described as follows:

### Step 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.

---

### Step 2

The 7x50 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 7x50 will show up as DC GW with its System ID, personality (for example, 7x50) and the available WAN resources, that is, **vsd-domains** Ent-1 and Ent-2.

---

### Step 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.

---

### Step 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.

---

### Step 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.

---

## 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 7x50 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 7x50:

- L2-DOMAIN – it will be associated with a VPLS service in the 7x50 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 7x50 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 7x50 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 7x50 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
```

## 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 7x50 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 7x50 will add to the VPLS service. For the **route-target**, 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"
```

## VSD-Domains and Association to Static-Dynamic Services

```
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
```

## 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 7x50 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.

---

## 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 7x50 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
```

---

## 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 7x50 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 7x50 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 **vs-import/export** policies, as described in the section for L2-DOMAINS.

The following commands help show the association between the 7x50 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
=====

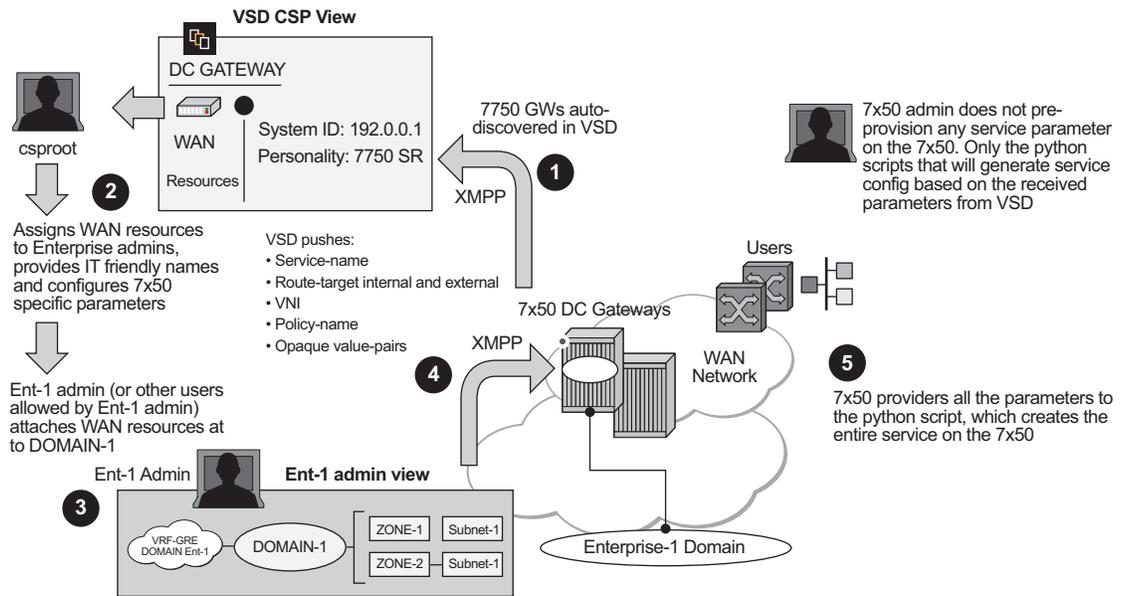
```

## Fully-Dynamic VSD Integration Model

In the Static-Dynamic VSD integration model, the VPLS/VPN 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/VPN service to the L2/L3 domain in the Data Center. In the Fully-Dynamic VSD integration model, the entire VPLS/VPN service configuration will be dynamically driven from VSD and no static configuration is required. Through the existing XMPP interface, the VSD will provide the 7x50 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 7x50, translating the VSD data model into a 7x50 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 VPN 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 7x50 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 121](#).



al\_0726

Figure 121: Fully-Dynamic VSD Integration Model Workflow

The Fully-Dynamic VSD integration model can be summarized in the steps shown in [Figure 121](#) and described as follows:

## Step 1

The 7x50 administrator only needs to provision parameters required for connectivity to the VSD, a **service-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 7x50 boots up or the gateway configuration is changed, the 7x50 sends a message to the VSD indicating its capabilities:

- System-ID
- Name and gateway type

The VSD uses this information to register the 7x50 on its list of 7x50 GWs.

Once registered, the VSD and 7x50 exchange messages where the VSD communicates its list of service-names and their domain-type to the 7x50. Based on this list, the 7x50 sends an XMPP IQ message to request the configuration of a specified service.

The 7x50 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 7x50 for a long time.

The 7x50 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}**.

---

### Step 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.

---

### Step 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.

---

### Step 4

When the Enterprise administrator attaches the service requested by the 7x50 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 7x50 (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 7x50.  
**Note** — The keys or values do not need to follow any specific format. The python script interprets and translates them into the 7x50 data model.
- Python-policy— Used by the 7x50 to find the Python script that will translate the VSD parameters into configuration.

## Step 5

When the 7x50 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*.

## Python Script Implementation Details

A python-script provides an intermediate abstraction layer between VSD and the 7750, translating the VSD data model into the 7x50 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 and it is not displayed by the info commands when executed on the service contexts. Also 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. 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', 'ser-
vicetype': '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 current 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-vsdsd-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-vsdsd-config** mode.
- Unless the CLI user enters the **enable-vsdsd-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

## Python Script Implementation Details

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' : vsdPar-
ams['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', 'ser-
vicetype': '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
```

## Python Script Implementation Details

```
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"
```

## BGP-EVPN Control Plane for MPLS Tunnels

[Table 18](#) lists all the EVPN routes supported in 7x50 SROS and their usage in EVPN-VXLAN, EVPN-MPLS, and PBB-EVPN.

**Note** — Route type 1 is not required in PBB-EVPN as per draft-ietf-l2vpn-pbb-evpn.

**Table 18: 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          | -         | -        |

RFC7432 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 117](#). See “BGP-EVPN Control Plane for VXLAN Overlay Tunnels” on page 1053. 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 within the data center. The received inclusive multicast routes will add entries to the VPLS flood list in the 7x50. Only ingress replication is supported over MPLS tunnels. The following route values are used for EVPN-MPLS services:

- 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).
  - PMSI attribute:
    - Tunnel type = Ingress replication (6).
    - Flags = Leaf not 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 MAC/IP routes for the same service unless `bgp-evpn mpls ingress-replication-bum-label` is configured in the service.
    - Tunnel end-point = Equal to the originating IP address.
- 

### EVPN Route Type 2 - MAC/IP Advertisement Route

The 7x50 generates this route type for advertising MAC addresses (and IP addresses if proxy-arp/nd is enabled). The 7x50 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 7x50 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 122](#) shows the fields in routes type 1 and 4 and their associated extended communities.

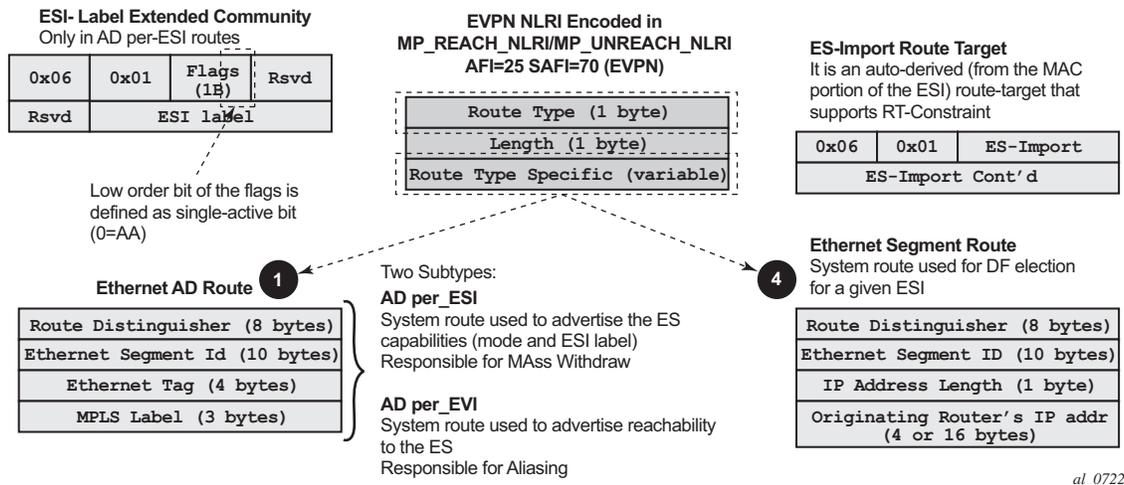


Figure 122: EVPN Routes Type 1 and 4

### **EVPN Route Type 1 - Ethernet Auto-discovery Route (AD route)**

The 7x50 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 7x50 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 Ethernet AD per-EVI route generated by a 7x50 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 7x50 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).

---

### **RFC5512 - BGP Tunnel Encapsulation Extended Community**

The following routes are sent with the RFC5512 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 7X50 processes the following BGP Tunnel Encapsulation tunnel values registered by IANA for RFC5512:

- 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 no RFC5512 encapsulation extended community is present in a received route, BGP will treat the route as MPLS or VXLAN-based configuration of the `config>router>bgp>neighbor# def-recv-evpn-encap [mpls|vxlan]` command.

## 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, LDP, RFC3107 BGP, Segment Routing IS-IS, or Segment Routing OSPF. 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 policies 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.

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 RFC7432 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.

**Note** — 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**

## EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS)

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

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
```

```

1 00:ca:fe:ca:fe:72 192.0.2.70:262140 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
=====

```

## EVPN and VPLS Integration

The 7x50 SROS 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 7x50 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 (in which all the EVPN endpoints are) is still used, but it will not be possible to refer to it on saps/spoke-sdps.

**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/ND is enabled and an IP->MAC pair is learned on a sap/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 123](#) shows an example of EVPN-VPLS integration.



## EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS)

```
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.

## 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 (not **bgp-vpls** in 13.0.R4). Note that 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**.
  - 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.

## 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.

Each **ethernet-segment** has a unique identifier called **esi** (Ethernet Segment Identifier) that is 10 bytes long and is manually configured in the 7x50.

**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.

## EVPN All-Active Multi-Homing

As described in RFC7432, 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 7x50 SROS to provide all-active multi-homing for a specified ethernet-segment:

- DF (Designated Forwarder) election
- Split-horizon
- Aliasing

[Figure 124](#) shows the need for DF election in all-active multi-homing.

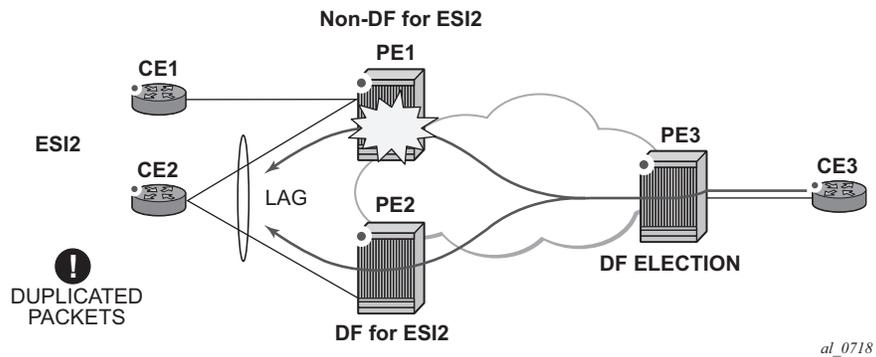


Figure 124: DF Election

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 125 shows the EVPN split-horizon concept for all-active multi-homing.

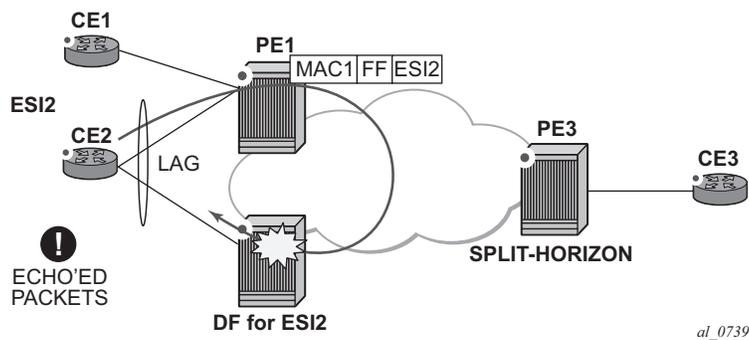


Figure 125: Split-Horizon

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 126 shows the EVPN aliasing concept for all-active multi-homing.

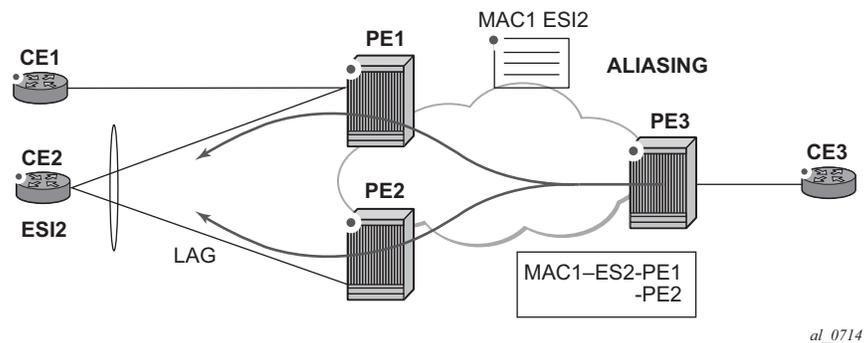


Figure 126: Aliasing

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 described in Figure 126.

```
*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
```

## EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS)

```
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 127](#).

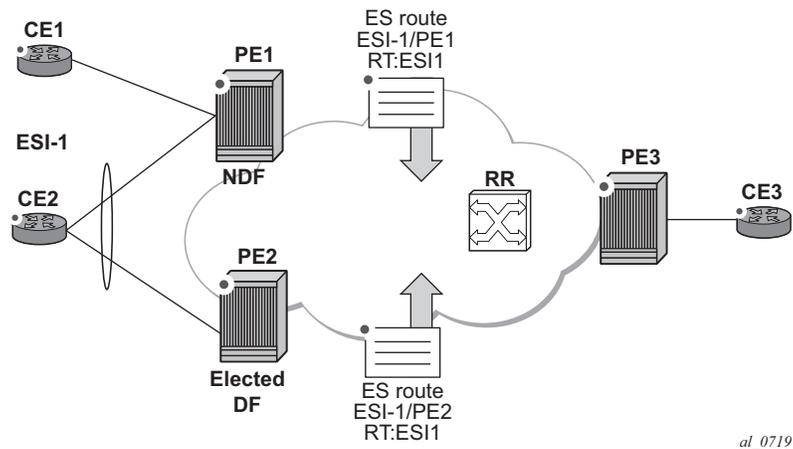


Figure 127: ES Discovery and DF Election

## 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 procedures.

## Step 2 - DF Election

Once 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 7x50 SROS is **service-carving** (as per RFC7432). 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, in addition to 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  $[V(\text{evi}) \bmod N(\text{peers}) = i(\text{ordinal})]$  to identify the DF for a specified service and ESI, as described in the following example:
  - As shown in [Figure 127](#), PE1 and PE2 are configured with ESI-1. Given that  $V(10) \bmod N(2) = 0$ , 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> primary**.
  - The system will be the PE forwarding/multicasting traffic for the **evi** identifiers included as primary. The PE will be secondary (non-DF) for the non-specified **evs**.
  - 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 irrespective of the primary status.
  - For example, as shown in [Figure 127](#): if PE1 is configured with service-carving manual evi 1 to 100 primary and PE2 with service-carving manual evi 101 to 200 primary, 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 BUM traffic. 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.

## EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS)

```
*A:PE1# show service id 1 fdb detail
```

```
Forwarding Database, Service 1
```

| ServId | MAC               | Source-Identifier           | Type  | Last Change       |
|--------|-------------------|-----------------------------|-------|-------------------|
| 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:<br>192.0.2.70:262140 | EvpnS | 06/11/15 00:09:06 |
| 1      | 00:ca:fe:ca:fe:72 | eMpls:<br>192.0.2.72:262141 | EvpnS | 06/11/15 00:09:39 |

```
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       |
|--------|-------------------|-----------------------------|-------|-------------------|
| 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:<br>192.0.2.69:262141 | EvpnS | 06/11/15 00:09:40 |
| 1      | 00:ca:fe:ca:fe:70 | eMpls:<br>192.0.2.70:262140 | EvpnS | 06/11/15 00:09:40 |

```
No. of MAC Entries: 3
```

```
Legend: L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
```

**Note** — 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 being used 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 127](#), 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
----- -
1 00:ca:ca:ba:ce:03 eES:01:00:00:00:00:71:00:00:00:01 Evpn 06/11/15 00:14:47
1 00:ca:fe:ca:fe:69 eMpls:192.0.2.69:262141 EvpnS 06/11/15 00:09:18
1 00:ca:fe:ca:fe:70 eMpls:192.0.2.70:262140 EvpnS 06/11/15 00:09:18
1 00:ca:fe:ca:fe:72 eMpls:192.0.2.72:262141 EvpnS 06/11/15 00:09:39
----- -
No. of MAC Entries: 4

Legend: L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
=====

```

## EVPN for MPLS Tunnels in VPLS Services (EVPN-MPLS)

The following command shows all the EVPN-MPLS destination bindings on PE3, including the ES destination bindings.

**Note** — 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 128 shows the behavior on the remote PEs (PE3) when there is an **ethernet-segment** failure.

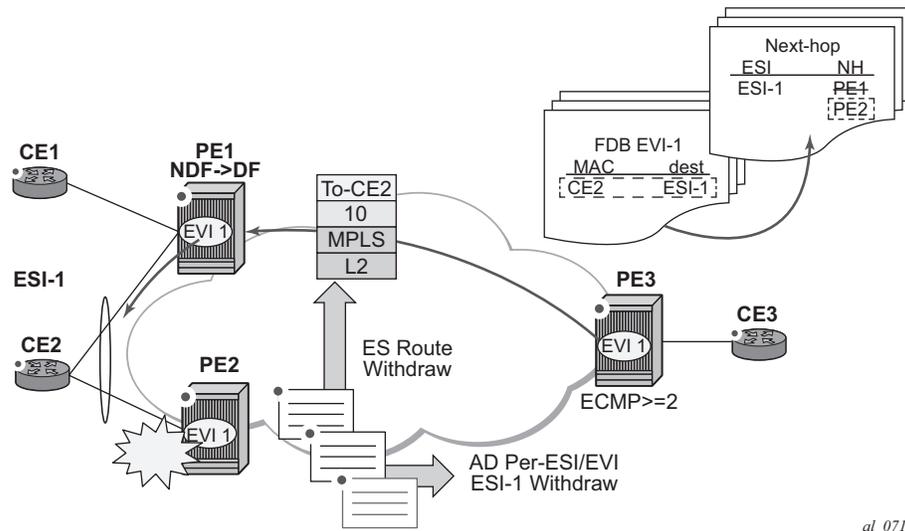


Figure 128: All-Active Multi-Homing ES Failure

The unicast traffic behavior on PE3 is as follows:

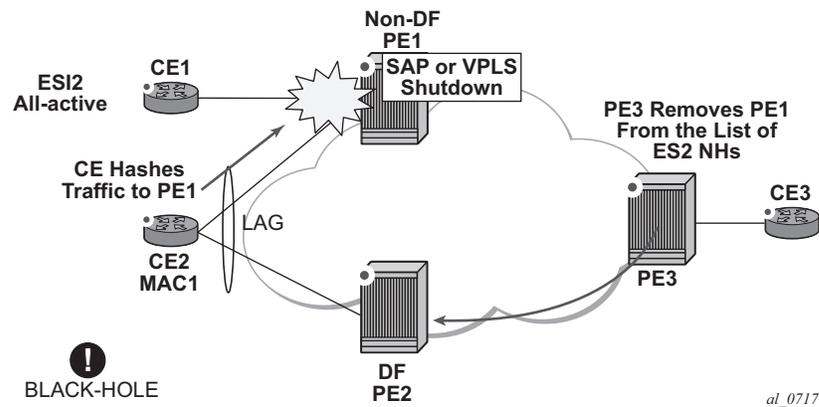
1. PE3 can only forward MAC DA = CE2 to both PE1 and PE2 when the MAC advertisement route from PE1 (or PE2) and the set of Ethernet AD per-ES routes and Ethernet AD per-EVI routes from PE1 and PE2 are active at PE3.
2. If there was a failure between CE2 and PE2, PE2 would withdraw its set of Ethernet 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 129](#):



**Figure 129: 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 *oper-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.

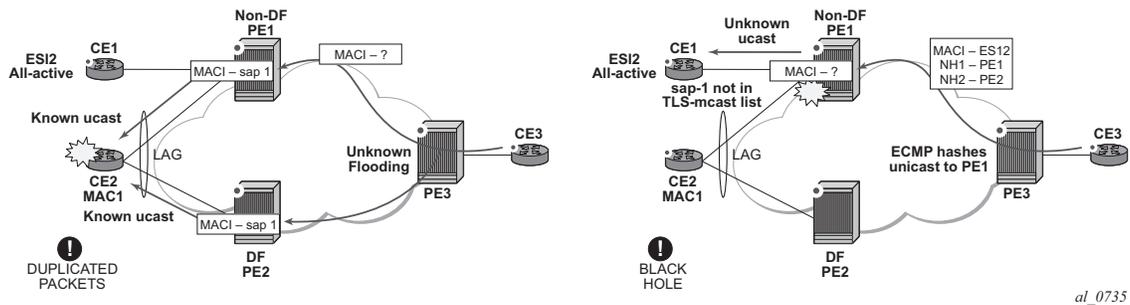
The same result occurs if the ES SAP is administratively **shutdown** instead of the service.

**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 130](#) and explained below.



**Figure 130: 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 130](#). 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 130](#). 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.

As soon as PE1 learns MAC1, the black-hole will be resolved.

---

## EVPN Single-Active Multi-Homing

The 7x50 SROS 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 7x50 SROS 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 RFC7432 procedure is known as 'Backup PE' and is shown in [Figure 131](#) for PE3.

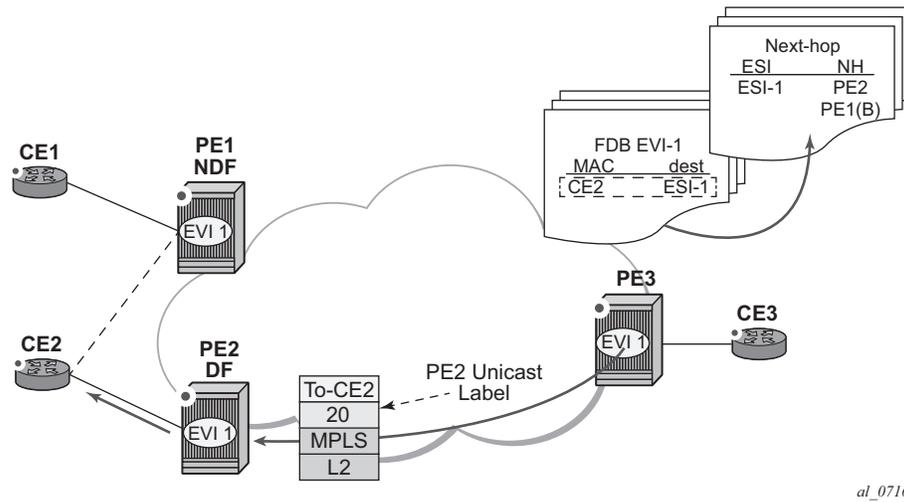


Figure 131: Backup PE

## 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 131](#).

```
*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 01:12:12:12:12:12:12:12:12:12
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-active** 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 131](#), 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 **StandByForMHPProtocol** 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 oper-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 oper-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/SDP-binding operationally down. Refer to the [ES Discovery and DF Election Procedures on page 1112](#) for more information.

The following **show** commands display the status of the single-active ESI-7413 in the non-DF.

**Note** — 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
```

## Ethernet Virtual Private Networks (EVPNs)

```
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 131](#)) 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

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

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
```

## Ethernet Virtual Private Networks (EVPNs)

```
192.0.2.254 ldp
 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 Transport	Last Change
01:74:13:00:74:13:00:00:74:13	192.0.2.73	262140 ldp	06/11/2015 19:59:03

```

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 132 shows the remote PE (PE3) behavior when there is an ethernet-segment failure.

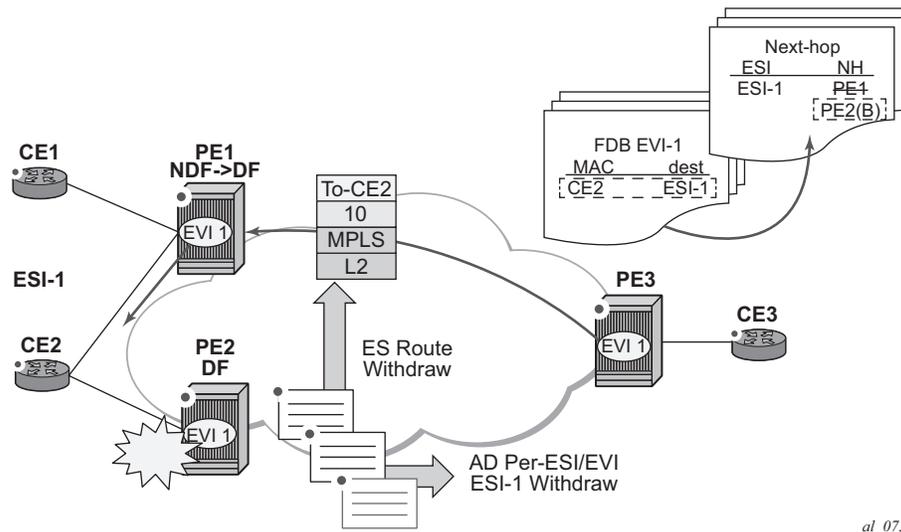


Figure 132: Single-Active Multi-Homing ES Failure

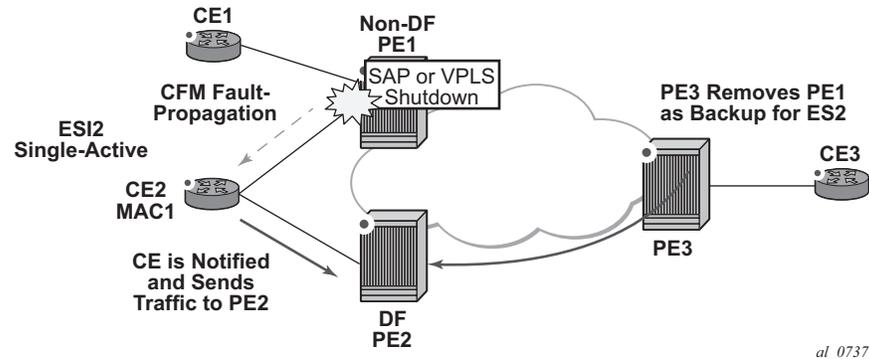
The PE3 behavior for unicast traffic is as follows:

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 to 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 once the **esi-activation-timer** expiration occurs (the timer kicks in when there is a transition from non-DF to DF).

## 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 133](#):



**Figure 133: Single-Active Multi-Homing SAP/SDP/Service Shutdown**

Common effects to consider are:

- If an individual VPLS service is administrative **shutdown** in PE1, the corresponding SAP/SDP-binding will go oper-down. This event will trigger the withdrawal of the AD per-EVI route for that particular SAP/SDP-binding. PE3 will apply its backup mechanisms and PE2 will take over as DF (if it was not the DF already). To signal the fault to CE2:
  - Eth-cfm must be used between the PEs and the multi-homed CE - if the connectivity is based on SAPs. A down MEP on the SAP going down will propagate the fault to a MEP on the CE. This network event will not produce any black-holing (in the all-active multi-homing case, this would partially black-hole the traffic).
  - PW status bits must be used between the PEs and the multi-homed CE - if the connectivity is based on pseudowires.
- The same result occurs if the ES SAP/SDP-binding is administrative **shutdown** instead of the service.

## BGP-EVPN Control Plane for PBB-EVPN

PBB-EVPN uses a reduced subset of the routes and procedures described in RFC7432. The supported routes are:

- ES routes
- MAC/IP routes
- Inclusive Multicast Ethernet Tag routes.

### 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.  
**Note** —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 end-point = equal to the originating IP address.

### EVPN Route Type 2 - MAC/IP Advertisement Route (or BMAC Routes)

The 7x50 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 7x50 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):
  - 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.
- 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.

### 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 on page 1095](#).

**Note** — The EVPN route type 1—Ethernet Auto Discovery route is not used in PBB-EVPN.

## PBB-EVPN for I-VPLS and PBB Epipe Services

The 7x50 SROS 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 134 shows the PBB-EVPN model in 7x50 SROS.

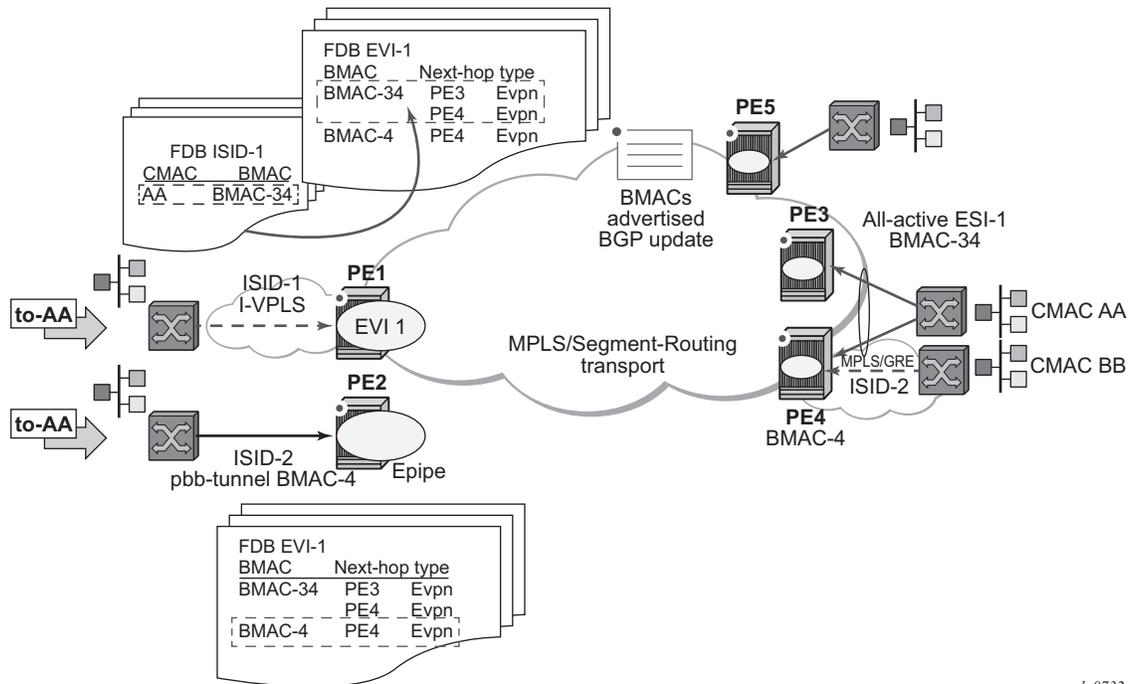


Figure 134: PBB-EVPN for I-VPLS and PFF Epipe Services

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 134, 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, in addition to 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 \(EVPN-MPLS\) on page 1100](#).

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 on page 1142](#) for more information)
  - `vpls>isid-policies`
  - `vpls>static-mac`
  - `vpls>sap/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`

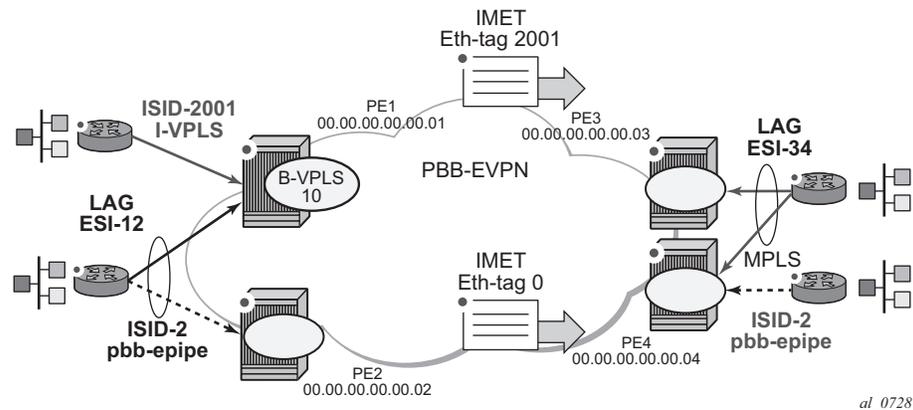
## Flood Containment for I-VPLS Services

In general, PBB technologies in 7x50 SROS 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/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 135](#) shows how the MFIBs are populated in PBB-EVPN.



**Figure 135: 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 (7x50 SROS 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.

## PBB-EVPN and PBB-VPLS Integration

The 7x50 SROS 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 [PBB-EVPN and PBB-VPLS Integration on page 1141](#) are also supported in B-VPLS services so that B-VPLS SAP/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.

## PBB-EVPN Multi-Homing in I-VPLS and PBB Epipe Services

The 7x50 SROS 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 on page 1108](#)'. 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

Draft-ietf-l2vpn-pbb-evpn 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 7x50 SROS 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 136 shows the use of all-active multi-homing in the 7x50 SROS PBB-EVPN implementation.

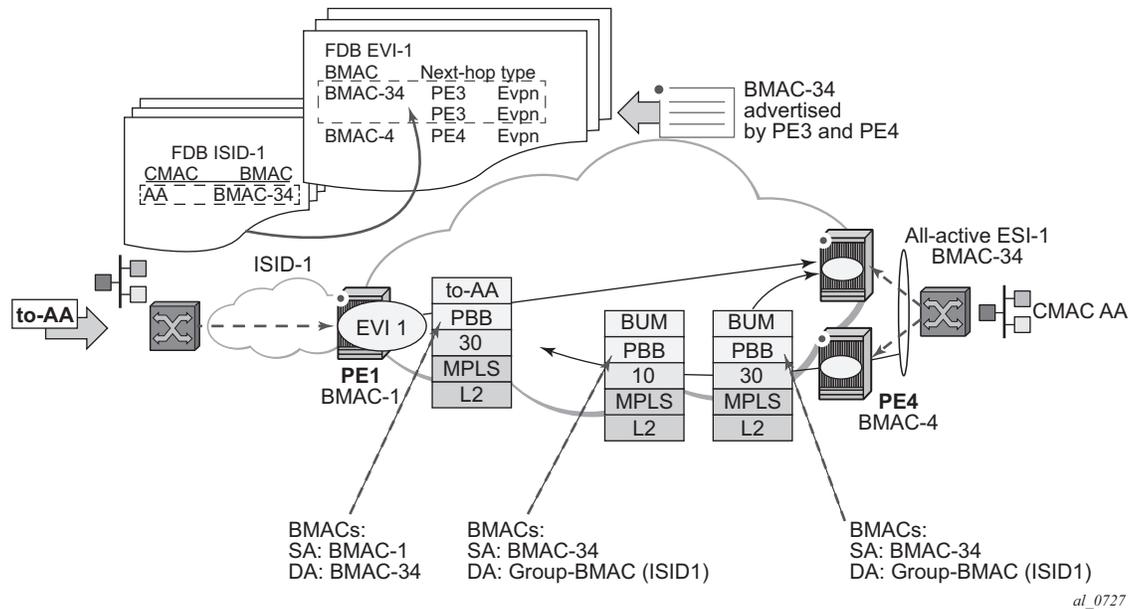


Figure 136: PBB-EVPN All-Active Multi-Homing

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 136 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
```

## PBB-EVPN all-active multi-homing service model

```
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 **pbbsource-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 136](#) 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 136](#) 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, Alcatel-Lucent 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 136](#):

```
*A:PE1# show service id 1 fdb detail
=====
Forwarding Database, Service 1
=====
ServId MAC Source-Identifier Type Last Change

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
```

## PBB-EVPN all-active multi-homing service model

```
=====
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 oper-down status.

- Change of **pbb>source-bmac**.
- Change of **es-bmac** (or removal of **pbb use-es-bmac**).
- Ethernet-segment transition to oper-down status.

**Note** — Individual saps going oper-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 on page 1143](#) 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 137](#).

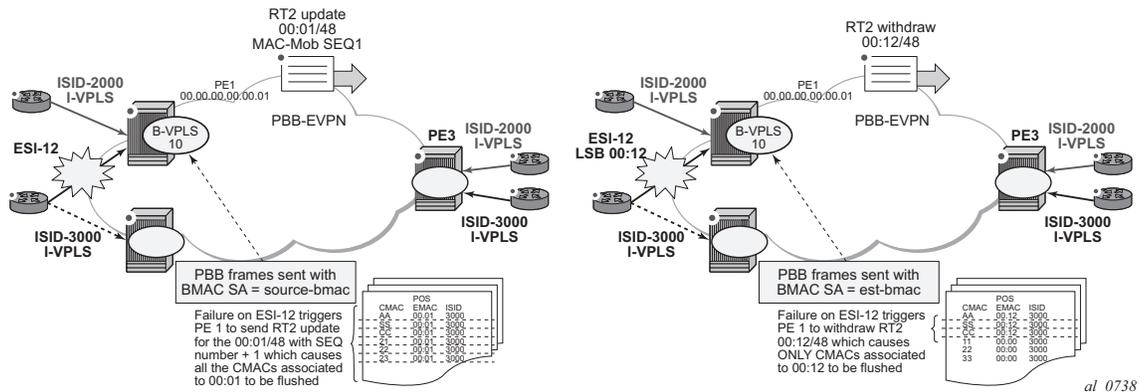


Figure 137: Source-Bmac Versus Es-Bmac CMAC Flushing

- If **es-bmacs** are used as shown in the representation on the right in Figure 137, 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 137, 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 oper-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 oper-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. Alcatel-Lucent 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 readvertise its **source-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 oper-down status. This will trigger the withdrawal of the associated BMACs, irrespective 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 oper-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 **no 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 oper-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-Epibes and EVPN Multi-Homing

EVPN multi-homing is supported with PBB-EVPN Epibes, but only in a limited number of scenarios. In general, the following applies to PBB-EVPN Epibes:

- PBB-EVPN Epibes don't support spoke-sdps that are associated with EVPN Ethernet Segments (ES).
- PBB-EVPN Epibes support all-active EVPN multi-homing as long as no local-switching is required in the Epibe instance where the ES is defined.
- PBB-EVPN Epibes support single-active EVPN multi-homing only in a two-node case scenario.

Figure 138 shows the EVPN MH support in a three-node scenario.

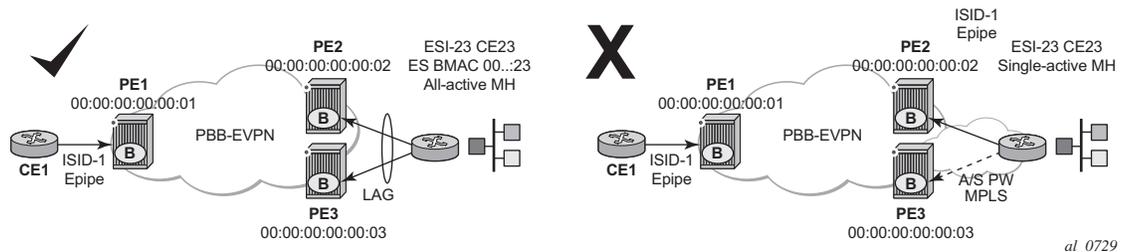


Figure 138: PBB-EVPN MH in a Three-Node Scenario

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 138). CE1 might also be multi-homed to other PEs, as long as those PEs are not PE2 or PE3. In this case, PE1 Epibe'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 138). Since PE1's Epibe **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 irrespective of the ES association to port/LAG/sdps.
- Irrespective of the EVPN multi-homing type, the CLI prevents the user from adding a spoke-sdp to an Epibe, if the corresponding SDP is part of an ES.

Figure 139 shows the EVPN MH support in a two-node scenario.

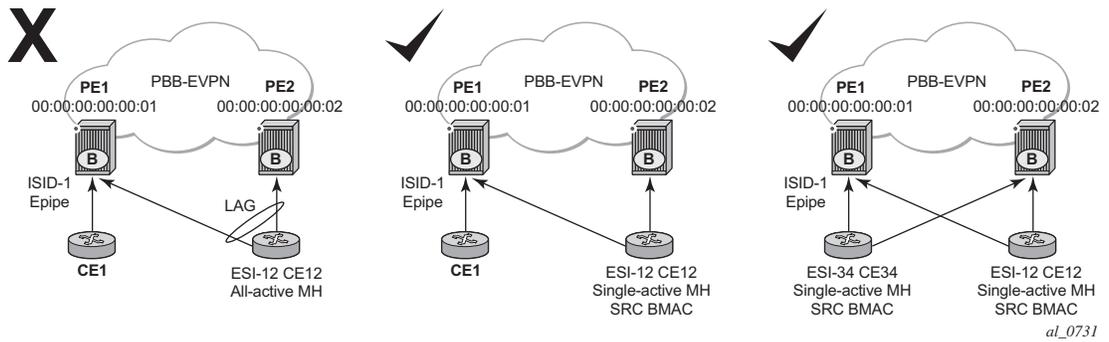


Figure 139: PBB-EVPN MH in a Two-Node Scenario

EVPN MH support in a two-node scenario has the following characteristics, as shown in Figure 139:

- 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 139.
- 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 139).

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.

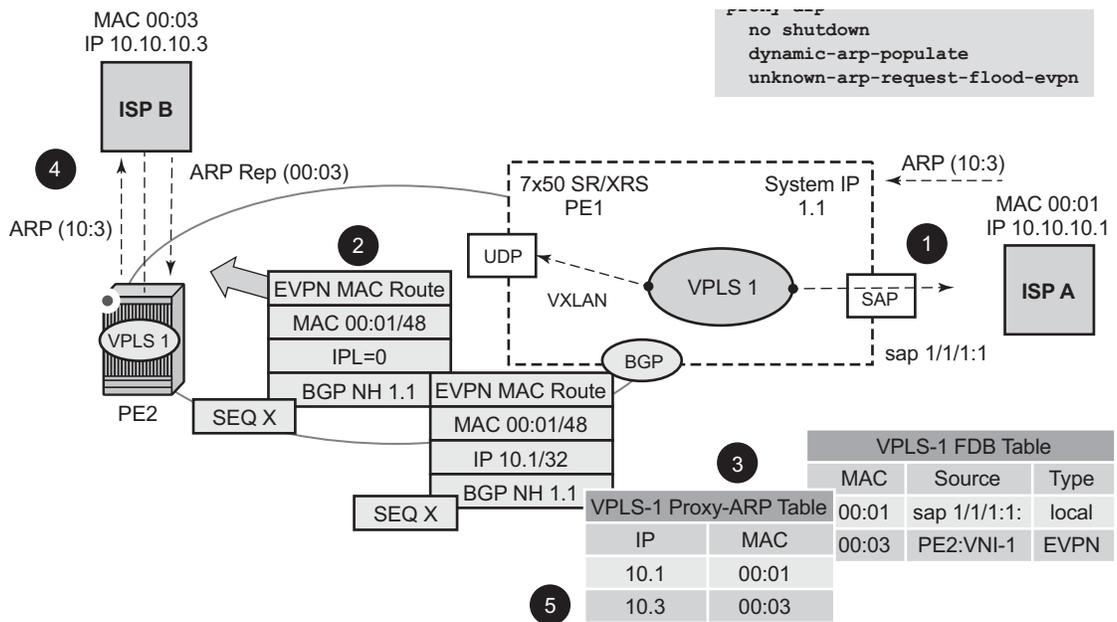
# ARP/ND Snooping and Proxy Support

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/nd no shutdown), the system will populate the corresponding proxy-ARP/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/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 140 shows an example of how proxy-ARP is used in an EVPN network. Proxy-ND would work in a similar way. Note that the MAC address notation in the diagram is shortened for readability.



ai\_0626

Figure 140: 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 140 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.
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/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/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->MACentry 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/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/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/ND offending entry's MAC is replaced by 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 has to be configured in order to drop traffic to the **anti-spoof-mac**.

Table 19 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, irrespective of the FDB entries

**NOTE**—A static entry is active in the FDB even when the service is down.

**Table 19: Proxy-arp Entry combinations**

Proxy-arp Entry Type	FDB Entry Type (for the same MAC)
Dynamic	learned
Static	learned
Dynamic	CStatic/Static

**Table 19: Proxy-arp Entry combinations**

Proxy-arp Entry Type	FDB Entry Type (for the same MAC)
Static	CStatic/Static
EVPN	EVPN
Duplicate	—

When proxy-ARP/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 19](#).

## Proxy-ARP/ND Periodic Refresh, Unsolicited Refresh and Confirm-Messages

When proxy-ARP/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.

## Proxy-ND and the Router Flag in Neighbor Advertisement messages

RFC4861 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 7x50 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.

---

### 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**.

## BGP-EVPN MAC-Mobility

EVPN defines a mechanism to allow the smooth mobility of MAC addresses from an NVE to another NVE. The 7x50 supports this procedure as well as the MAC-mobility extended community in MAC advertisement routes as follows:

- The 7x50 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.

## 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 7x50 supports 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 7x50 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 7x50 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: SVCNMR #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 7x50 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/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 7x50).
- b. The retry <in-minutes> timer expires, which will flush the MAC and restart the process.

**Note**—The other 7x50s in the VPLS instance will forward the traffic for the duplicate MAC address to the 7x50 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
```



## CFM Interaction with EVPN Services

Ethernet Connectivity & 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 & 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 dependant (ie. 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, there is no ETH-CFM functionality installed on the virtual connections. The ETH-CFM functions will only be installed on the supported SAP, spoke, and mesh 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: 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 UP MEPs on ES SAPs/SDP-bindings with EVPN multi-homing:
  - All-active multi-homing: 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: UP MEPs should be able to send or receive CFM packets normally.
  - 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.

## 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 needs 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.

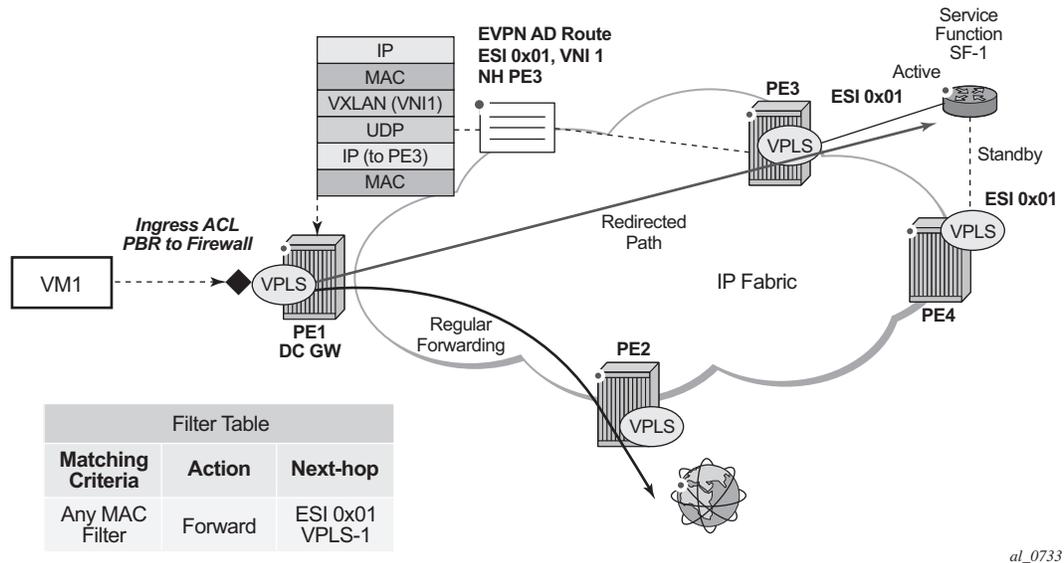
Refer to the Nuage VSP documentation for more information about the Nuage Service Chaining solution.

The 7x50 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.

---

## Policy Based Forwarding in VPLS Services for Nuage Service Chaining Integration in L2-Domains

[Figure 141](#) shows the 7x50 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 141: 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 141, the PBF filter will redirect the matching packets to ESI 0x01 in VPLS-1

**Note** — Figure 141 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:00:00:00:00:00:00:00:01
=====
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:00:00:00:00:00:00:00:01 192.0.2.72
 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

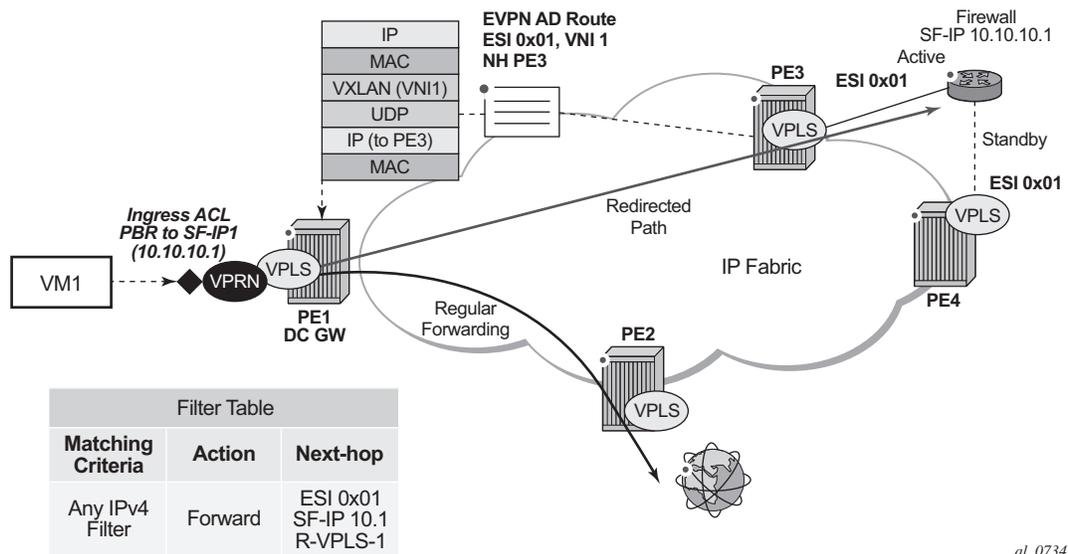
192.0.2.69 301 1 Yes Up No
192.0.2.72 301 1 Yes Up No
192.0.2.72 7272 0 No Up Yes

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.

## Policy Based Routing in VPRN Services for Nuage Service Chaining Integration in L2-DOMAIN-IRB Domains

Figure 142 shows the 7x50 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.



al\_0734

Figure 142: 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/NETCONF management interfaces. The PBR target identifies first service function in the chain (ESI 0x01 in Figure 142, 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

```

**Note** — 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 than the remote MAC/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

```

## Policy Based Routing in VPRN Services for Nuage Service Chaining Integration in L2-DOMAIN-IRB

```
=====
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

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
=====
```

As for the Layer-2 case, 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.

## 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 for the routes is 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. EVPN ES PBR MACs (see ES PBR MAC routes below)
  3. EVPN static MACs (remote protected MACs)
  4. Data plane learned MACs (regular learning on saps/sdp-bindings)
  5. EVPN MACs with higher SEQ number
  6. Lowest IP (next-hop IP of the EVPN NLRI)
  7. Lowest eth-tag (that will be zero for MPLS and might be different from zero for VXLAN)
  8. 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 given (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.

## 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 7x50-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/ND confirm messages and Eth-CFM for EVPN-MPLS.
- **eth-cfm** (meps, vmeps, mips): This command can be configured and used in EVPN VPLS service objects (service, saps and sdp-bindings). Although **vmeps** can be configured and used to local saps and sdp-bindings, **eth-cfm** tests will not work through EVPN destination bindings.
- 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-protect**—**mac-protect** cannot be used in conjunction with EVPN.
 

**Note**—EVPN provides its own protection mechanism for static mac addresses.
- **provider-tunnel**—p2mp RSVP/mLDP LSPs are not supported in the **bgp-evpn** service. The configuration of the provider-tunnel is blocked.
- MAC OAM—any MAC OAM tool is 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 cannot be enabled in the same VPLS service. BGP-MH can still be used for multi-homing as long as no ethernet-segments are configured in the service SAPs/SDP-bindings. There is no limitation in terms of 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.

- **Note**—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 **StandByForMHP** protocol flag.
- IGMP-snooping is not supported in VPLS (or I-VPLS) services when **bgp-evpn mpls** is enabled (in the service or the associated B-VPLS).
- 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/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/PIM-snooping and attributes
  - BPDU translation
  - L2PT termination
  - MAC-pinning
- Other commands not supported in conjunction with **bgp-evpn mpls** are:
  - VSD-domains
  - VXLAN cannot be not shutdown under bgp-evpn. Both bgp-evpn vxlan and bgp-evpn mpls are mutually exclusive
  - SPB configuration and attributes
  - allow-ip-int-bind (R-VPLS) and bgp-evpn ip-route-advertisement

## 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 on page 1177](#), 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 Epipe).
- PBB-EVPN multi-homing and BGP-MH cannot be enabled in the same **i-vpls** service.
- **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/SDP-bindings at the same time.
- When PBB-epipes are used with PBB-EVPN multi-homing, spoke-SDPs are not supported on **ethernet-segments**.

## Interaction of EVPN-VXLAN with Existing VPRN Features

When trying to enable existing VPRN features on interfaces linked to EVPN-VXLAN IRB backhaul or EVPN tunnel R-VPLS interfaces, consider the following:

- The following commands are not supported:
  - arp-populate
  - authentication-policy
- Dynamic routing protocols such as IS-IS, RIP, and OSPF are not supported.
- BFD is not supported on EVPN tunnel interfaces.

## 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 143 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.

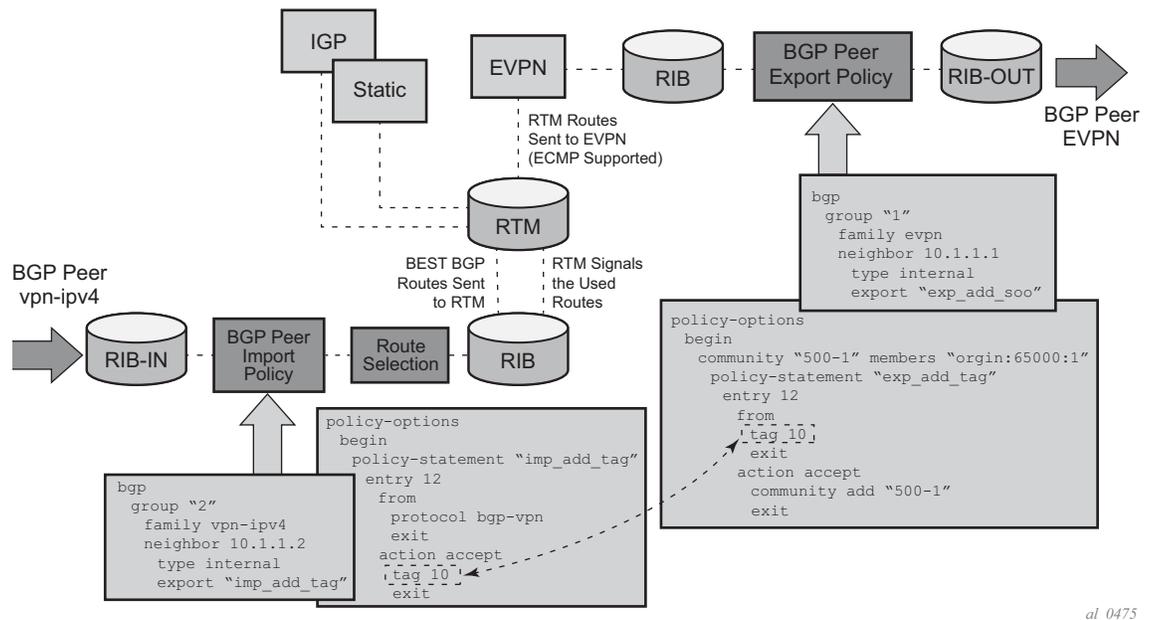


Figure 143: IP-VPN Import and EVPN Export BGP Workflow

**Note** — 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:

```

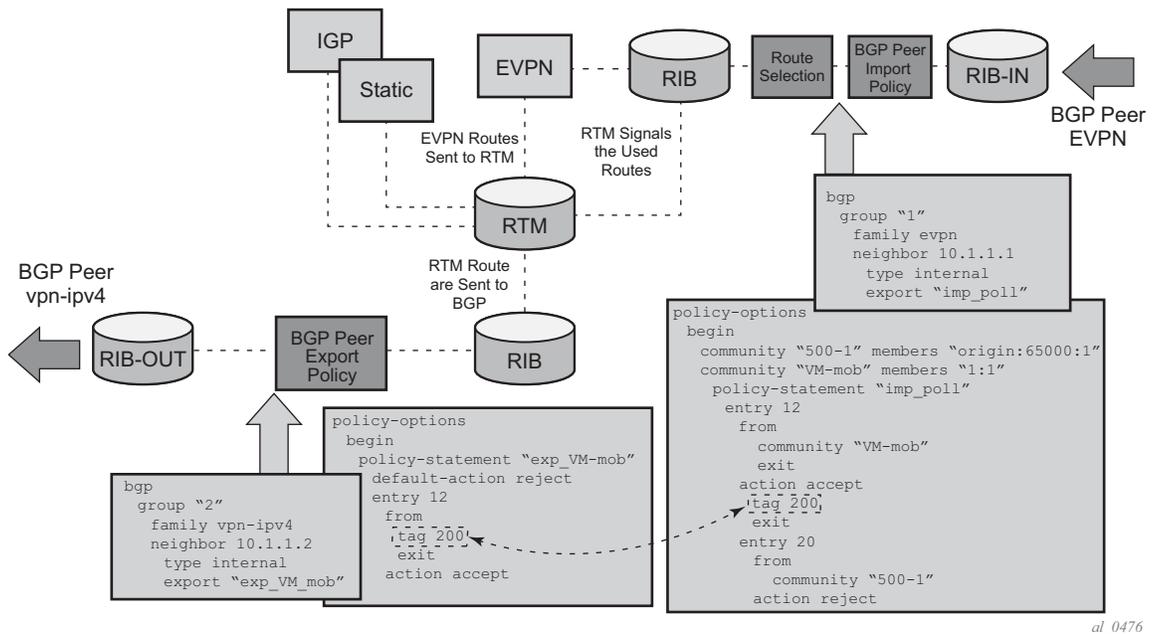
<tag> : accepts in decimal or hex
 [0x1..0xFFFFFFFF]H (for OSPF and IS-IS)
 [0x1..0xFFFF]H (for RIP)

```

## Routing Policies for BGP EVPN IP Prefixes

[0x1..0xFF]H (for BGP)

Figure 144 shows an example of the reverse workflow: routes imported from EVPN and exported from RTM to BGP VPN-IPv4.



al\_0476

Figure 144: IEVPN Import and I-VPN Export BGP Workflow

**Note** — 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.

## Configuring an EVPN Service with CLI

This section provides information to configure VPLS using the command line interface.

Topics in this section include:

- [EVPN-VXLAN Configuration Examples on page 1184](#)
  - [EVPN for VXLAN in R-VPLS Services Example on page 1186](#)
  - [EVPN for VXLAN in EVPN Tunnel R-VPLS Services Example on page 1188](#)
  - [EVPN for VXLAN in R-VPLS Services with IPv6 interfaces and prefixes Example on page 1189](#)
- [EVPN-MPLS Configuration Examples on page 1190](#)
  - [EVPN All-active Multi-homing Example on page 1190](#)
  - [EVPN Single-active Multi-homing Example on page 1193](#)
- [PBB-EVPN Configuration Examples on page 1195](#)
  - [PBB-EVPN All-active Multi-homing Example on page 1195](#)
  - [PBB-EVPN Single-active Multi-homing Example on page 1198](#)

## EVPN-VXLAN Configuration Examples

---

### Layer 2 PE Example

This section shows a configuration example for three 7x50 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

```
service vpls 2000 customer 1 create
 vxlan vni 2000 create
 bgp
 route-target 65000:2000
 route-distinguisher 65010:2000
 bgp-evpn
 no shutdown
 vxlan
 no shutdown
```

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
```

## EVPN for VXLAN in R-VPLS Services Example

This section shows a configuration example for three 7x50 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

```

## EVPN for VXLAN in EVPN Tunnel R-VPLS Services Example

The example in [EVPN for VXLAN in R-VPLS Services Example on page 1186](#) 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
```

## 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 NLRI. 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

```

## EVPN-MPLS Configuration Examples

---

### EVPN All-active Multi-homing Example

This section shows a configuration example for three 7x50 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 to 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

```

## EVPN-MPLS Configuration Examples

```
no shutdown
```

-----

The configuration on the remote PE (i.e. PE-3), which supports aliasing to PE-1 and PE-2 is shown below . Note that 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

```

## 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
```

## EVPN-MPLS Configuration Examples

```
es-activation-timer 10
service-carving
 mode auto
exit
multi-homing single-active
lag 1
no shutdown
exit
```

-----

## PBB-EVPN Configuration Examples

---

### PBB-EVPN All-active Multi-homing Example

As in the [EVPN All-active Multi-homing Example on page 1190](#), this section also shows a configuration example for three 7x50 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 to 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
```

## PBB-EVPN Configuration Examples

```
 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#

```

Note that the combination of the pbb **source-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".

## 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.

Note that 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 on page 1195](#).

```
*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
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

```

## PBB-EVPN Configuration Examples

```
 shutdown
 exit
 no shutdown

A:PE73>config>service>vpls#
```

---

# EVPN Command Reference

---

## Command Hierarchies

```

config
 — service
 — vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls] [create]
 — no vpls 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
 — 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
 — [no] cfm-mac-advertisement
 — [no] evi value
 — [no] ip-route-advertisement [incl-host]
 — [no] mac-advertisement
 — mac-duplication
 — detect num-moves num-moves window minutes
 — [no] retry minutes
 — mpls
 — auto-bind-tunnel
 — resolution {disabled|any|filter}
 — resolution-filter
 — [no] bgp
 — [no] ldp
 — [no] rsvp
 — [no] sr-isis
 — [no] sr-ospf
 — [no] control-word
 — ecmp max-ecmp-routes
 — [no] force-vlan-vc-forwarding
 — [no] ingress-replication-bum-label
 — [no] shutdown
 — [no] split-horizon-group
 — [no] unknown-mac-route
 — vxlan
 — [no] shutdown
 — pbb
 — [no] use-es-bmac
 — [no] proxy-arp
 — [no] age-time seconds

```

```

— dup-detect [anti-spoof-mac mac-address] window minutes num-moves
 count hold-down minutes|max
— [no] dynamic-arp-populate
— [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] window minutes num-moves
 count hold-down minutes|max
 — [no] dynamic-nd-populate
 — [no] evpn-nd-advertise
 — [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
— 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
 — no mac ieee-address
— vsd-domain name
— no vsd-domain vni
— vxlan vni vni-id create
— no vxlan vni
— vprn
 — interface
 — vpls
 — [no] evpn-tunnel
— vsd
 — domain name [type {l2-domain|vrf-gre|vrf-vxlan|l2-domain-irb}] [create]
 — [no] domain name
 — description discription -string
 — [no] description
 — shutdown
 — [no] shutdown
 — service-range svc-id to sve-id

config
 — service
 — system
 — [no] bgp-auto-rd-range ip-address comm-val [1..65535] to [1..65535]
 — [no] bgp-evpn
 — route-distinguisher rd
 — no ethernet-segment name [ceate]
 — [no] es-activation-timer seconds
 — [no] esi esi
 — [no] lag lag-id

```

- [no] **multi-homing** {single-active [no-esi-label
- [no] **port** *port-id*
- [no] **sdp** *sdp-id*
- **service-carving**
  - **mode** {manual | auto | off}
  - **manual**
    - [no] **evi** *start* [to *to*] primary
    - [no] **isid** *start* [to *to*] primary
- [no] **shutdown**
- [no] **source-bmac-lsb** *MAC Lsb* [ex-bmac-table-size *size*]

- 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]
      - [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*

**enable-vsd-config** *name*

## Show Commands

```
show
 — service
 — id service-id
 — bgp-evpn
 — proxy-arp
 — vxlan
 — es-pbr
 — evpn-mpls
 — esi esi
 — es-bmac ieee-address
 — 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]
 — vsd [vsd] [origin creation-origin]
 — domain [domain-name] [association]
 — root-objects
 — script
 — snippets snippet-name [instance snippet-instance] [detail]
 — statistics
 — summary
 — vxlan [VTEP ip-address]
 — evpn-mpls [TEP ip-address]
```

```
show
 — system
 — vsd
 — domain
 — xmpp
 — server
 — vsd
```

```
show
 — redundancy
 — bgp-evpn-multi-homing
```

## Clear Commands

```

clear
 — service
 — statistics
 — vsd
 — domain name
 — scripts name

clear
 — system
 — statistics
 — xmpp
 — server xmpp-server-name

```

## 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

```

## Tools Commands

```

tools
 — dump
 — service
 — domain-to-vsds-mapping
 — domain name name
 — evpn
 — usage
 — id service-id
 — evpn-mpls [clear] [default-multicast-list]

```

## Command Hierarchies

- 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
      - vsd
        - domain
          - name [*name*] refresh-config
        - fd-domain-sync {full | diff}
        - evaluate-script domain-name [*64 chars max*] 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

## EVPN Commands

### vpls

<b>Syntax</b>	<b>vpls</b> <i>service-id</i> <b>customer</b> <i>customer-id</i> <b>vpn</b> <i>vpn-id</i> [ <b>m-vpls</b> ] [ <b>bvpls</b>   <b>i-vpls</b> ] [ <b>create</b> ] <b>no vpls</b> <i>service-id</i>
<b>Context</b>	config>service
<b>Description</b>	<p>This command creates or edits a Virtual Private LAN Services (VPLS) instance. The <b>vpls</b> command is used to create or maintain a VPLS service. If the <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> exists, the context for editing the service is entered.</p> <p>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.</p> <p>When a service is created, the <b>create</b> keyword must be specified if the <b>create</b> command is enabled in the <b>environment</b> context. When a service is created, the <b>customer</b> keyword and <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> must already exist having been created using the <b>customer</b> 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.</p> <p>Once a service is created, the use of the <b>customer</b> <i>customer-id</i> is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect <i>customer-id</i> specified will result in an error.</p> <p>More than one VPLS service may be created for a single customer ID.</p> <p>By default, no VPLS instances exist until they are explicitly created.</p> <p>The <b>no</b> form of this command deletes the VPLS service instance with the specified <i>service-id</i>. 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.</p>
<b>Parameters</b>	<p><i>service-id</i> — 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 <i>service-id</i> must be the same number used for every router on which this service is defined.</p> <p><b>Values</b>     <i>service-id:</i>     1 — 2147483648                   <i>svc-name:</i>        64 characters maximum</p> <p><b>customer</b> <i>customer-id</i> — 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.</p> <p><b>Values</b>     1 — 2147483647</p> <p><b>vpn</b> <i>vpn-id</i> — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.</p> <p><b>Values</b>     1 — 2147483647</p> <p><b>Default</b>    <b>null</b> (0)</p>

## EVPN Commands

**m-vpls** — Specifies a management VPLS.

**b-vpls** | **i-vpls** — Creates a backbone-vpls or ISID-vpls.

### bgp

<b>Syntax</b>	<b>bgp</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command enables the context to configure the BGP related parameters for BGP AD, BGP VPLS and EVPN.

### route-target

<b>Syntax</b>	<b>route-target</b> { <i>ext-community</i>   {[ <b>export</b> <i>ext-community</i> ] [ <b>import</b> <i>ext-community</i> ]}}
	<b>no route-target</b>
<b>Context</b>	config>service>vpls>bgp-ad config>service>vpls>bgp
<b>Description</b>	<p>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.</p> <p>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.</p>
<b>Parameters</b>	<p><b>export</b> <i>ext-community</i> — Specify communities allowed to be sent to remote PE neighbors.</p> <p><b>import</b> <i>ext-community</i> — Specify communities allowed to be accepted from remote PE neighbors.</p>

### vsi-export

<b>Syntax</b>	<b>vsi-export</b> <i>policy-name</i> [ <i>policy-name</i> ...(up to 5 max)]
	<b>no vsi-export</b>
<b>Context</b>	config>service>vpls>bgp-ad config>service>vpls>bgp
<b>Description</b>	<p>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.</p> <p>The policy name list is handled by the SNMP agent as a single entity.</p>

## vsi-import

<b>Syntax</b>	<b>vsi-import</b> <i>policy-name</i> [ <i>policy-name...</i> (up to 5 max)] <b>no vsi-import</b>
<b>Context</b>	config>service>vpls>bgp-ad>vsi-id config>service>vpls>bgp
<b>Description</b>	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.

## route-distinguisher

<b>Syntax</b>	<b>route-distinguisher</b> [ <i>ip-addr:comm-val</i>   <i>as-number:ext-comm-val</i> ] <b>route-distinguisher auto-rd</b> <b>no route-distinguisher</b>												
<b>Context</b>	config>service>vpls>bgp												
<b>Description</b>	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: <ul style="list-style-type: none"> <li>• if BGP AD VPLS-id is configured &amp; no RD is configured under BGP node - RD = VPLS-ID</li> <li>• 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)</li> <li>• if BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails</li> </ul> Values and format (6 bytes, other 2 bytes of type will be automatically generated)  Alternatively, the <b>auto-rd</b> option allows the system to automatically generate an RD based on the <b>bgp-auto-rd-range</b> command configured at service level.												
<b>Parameters</b>	<i>ip-addr:comm-val</i> — Specifies the IP address. <table> <tr> <td><b>Values</b></td> <td>ip-addr</td> <td>a.b.c.d</td> </tr> <tr> <td></td> <td>comm-val</td> <td>0 — 65535</td> </tr> </table> <i>as-number:ext-comm-val</i> — Specifies the AS number. <table> <tr> <td><b>Values</b></td> <td>as-number</td> <td>1 — 65535</td> </tr> <tr> <td></td> <td>ext-comm-val</td> <td>0 — 4294967295</td> </tr> </table> <b>auto-rd</b> — the system will generate an RD for the service according to the IP address and range configured in the <b>bgp-auto-rd-range</b> command.	<b>Values</b>	ip-addr	a.b.c.d		comm-val	0 — 65535	<b>Values</b>	as-number	1 — 65535		ext-comm-val	0 — 4294967295
<b>Values</b>	ip-addr	a.b.c.d											
	comm-val	0 — 65535											
<b>Values</b>	as-number	1 — 65535											
	ext-comm-val	0 — 4294967295											

## vsi-import

<b>Syntax</b>	<b>vsi-import</b> <i>policy-name</i> [ <i>policy-name...</i> (up to 5 max)] <b>no vsi-import</b>
<b>Context</b>	config>service>vpls>bgp-ad>vsi-id config>service>vpls>bgp
<b>Description</b>	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-auto-rd-range

<b>Syntax</b>	<b>bgp-auto-rd-range</b> <i>ip-address comm-val comm-val to comm-val</i> <b>no bgp-auto-rd-range</b>
<b>Context</b>	config>service>system
<b>Description</b>	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 <b>bgp-enabled</b> services using <b>auto-rd</b> .  <b>Interactions:</b>  This command is used along with the <b>route-distinguisher auto-rd</b> command supported in VPLS, VPRN and Epipe services. The system forces the user to create a <b>bgp-auto-range</b> before the <b>auto-rd</b> option can be used in the services.  Note that the system will keep allocating values for services configured with <b>route-distinguisher auto-rd</b> as long as there are available community values within the configured range. Once the command is added, the following changes are allowed: <ul style="list-style-type: none"> <li>• The <i>ip-address</i> can be changed without modifying the <i>comm-val</i> range, even if services using <b>auto-rd</b> are present. The affected routes will be withdrawn and re-advertised with the new route-distinguishers.</li> <li>• The <i>comm-val</i> 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.</li> </ul>
<b>Parameters</b>	<i>ip-address</i> — Specifies the IPv4 address used in the first 4 octets of all the type-1 auto route-distinguishers selected by the system.  <i>comm-val</i> — Specifies the community value of the type-1 auto route-distinguisher.  <b>Values</b> 0 — 65535

## bgp-evpn

<b>Syntax</b>	<b>[no] bgp-evpn</b>
<b>Context</b>	config>service>vpls config>service>system
<b>Description</b>	This command enables the context to configure the BGP EVPN parameters in the base instance.

## route-distinguisher

<b>Syntax</b>	<b>route-distinguisher</b> [ <i>ip-addr:comm-val</i>   <i>as-number:ext-comm-val</i> ] <b>no route-distinguisher</b>												
<b>Context</b>	config>service>system>bgp-evpn												
<b>Description</b>	This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for 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												
<b>Default</b>	<b>no route-distinguisher</b>												
<b>Parameters</b>	<i>ip-addr:comm-val</i> — Specifies the IP address. <table> <tr> <td><b>Values</b></td> <td>ip-addr</td> <td>a.b.c.d</td> </tr> <tr> <td></td> <td>comm-val</td> <td>0 — 65535</td> </tr> </table> <i>as-number:ext-comm-val</i> — Specifies the AS number. <table> <tr> <td><b>Values</b></td> <td>as-number</td> <td>1 — 65535</td> </tr> <tr> <td></td> <td>ext-comm-val</td> <td>0 — 4294967295</td> </tr> </table>	<b>Values</b>	ip-addr	a.b.c.d		comm-val	0 — 65535	<b>Values</b>	as-number	1 — 65535		ext-comm-val	0 — 4294967295
<b>Values</b>	ip-addr	a.b.c.d											
	comm-val	0 — 65535											
<b>Values</b>	as-number	1 — 65535											
	ext-comm-val	0 — 4294967295											

## ethernet-segment

<b>Syntax</b>	<b>ethernet-segment</b> <i>name</i> <b>create</b> <b>no ethernet-segment</b>
<b>Context</b>	config>service>system>bgp-evpn
<b>Description</b>	This command configures an ethernet-segment instance its corresponding name.
<b>Parameters</b>	<i>name</i> — Specifies the 28-character ethernet-segment name.

## es-activation-timer

<b>Syntax</b>	<b>es-activation-timer</b> <i>seconds</i> <b>no es-activation-timer</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment

## EVPN Commands

<b>Description</b>	<p>This command configures the ethernet-segment activation timer for a given ethernet-segment. The <b>es-activation-timer</b> delays the activation of a given ethernet-segment on a given PE that has been elected as DF (Designated Forwarder). Only when the <b>es-activation-timer</b> 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).</p> <p>If <b>no es-activation-timer</b> is configured, the system will use the value configured at <b>config&gt;redundancy&gt;bgp-evpn-multi-homing&gt;es-activation-timer</b> if configured. Otherwise the system will use a default value of 3 seconds.</p>
<b>Default</b>	<b>no es-activation-timer</b>
<b>Parameters</b>	<i>seconds</i> — Specifies the number of seconds for the <b>es-activation-timer</b> .
<b>Values</b>	0 — 100 seconds

## esi

<b>Syntax</b>	<b>esi</b> <i>value</i> <b>no esi</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment
<b>Description</b>	This command configures the 10-byte ethernet-segment identifier 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.
<b>Parameters</b>	<i>value</i> — Specifies the 10-byte esi.
<b>Values</b>	00-11-22-33-44-55-66-77-88-99 with any of these separators ('-',':')

## lag

<b>Syntax</b>	<b>lag</b> <i>lag-id</i> <b>no lag</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment
<b>Description</b>	This command configures a lag-id associated to the ethernet-segment. When the ethernet-segment is configured as <b>all-active</b> , only a lag can be associated to the ethernet-segment. When the ethernet-segment is configured as <b>single-active</b> , 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 given lag can be part of only one ethernet-segment.
<b>Default</b>	<b>no lag</b>
<b>Parameters</b>	<i>lag-id</i> — Specifies the lag-id associated with the ethernet-segment.
<b>Values</b>	1 — 800

## multi-homing

<b>Syntax</b>	<b>multi-homing single-active</b> [no-esi-label] <b>multi-homing all-active</b> <b>no multi-homing</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment
<b>Description</b>	<p>This command configures the multi-homing mode for the ethernet-segment as <b>single-active</b> or all-active multi-homing, as defined in RFC7432.</p> <p>By default, the use of <b>esi-label</b> is enabled for <b>all-active</b> and <b>single-active</b> as defined in RFC7432 (for <b>single-active multi-homing</b>, the esi-label is used to avoid transient loops).</p> <p>When <b>single-active no-esi-label</b> 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.</p>
<b>Default</b>	<b>no multi-homing.</b>
<b>Parameters</b>	<p><i>single-active</i> — configures single-active mode for the ethernet-segment</p> <p><i>all-active</i> — configures the system to not send an esi-label for <b>single-active</b> mode</p> <p><i>no-esi-label</i> — configures single-active mode for the ethernet-segment</p>

## port

<b>Syntax</b>	<b>port port-id</b> <b>no port</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment
<b>Description</b>	<p>This command configures a port-id associated with the ethernet-segment. If the ethernet-segment is configured as <b>all-active</b> only a lag can be associated to the ethernet-segment. If the ethernet-segment is configured as <b>single-active</b>, 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 given port can be part of only one ethernet-segment. Only ethernet ports can be added to an ethernet-segment.</p>
<b>Default</b>	<b>no port</b>
<b>Parameters</b>	<i>port-id</i> — Specifies the slot/mda/port associated to the ethernet-segment.

## sdp

<b>Syntax</b>	<b>sdp sdp-id</b> <b>no sdp</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment
<b>Description</b>	<p>This command configures an sdp-id associated to the ethernet-segment. If the ethernet-segment is configured as <b>all-active</b> only a lag can be associated to the ethernet-segment. If the ethernet-segment</p>

## EVPN Commands

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 given sdp can be part of only one ethernet-segment. Only user-configured sdps can be added to an ethernet-segment.

<b>Default</b>	<b>no sdp</b>
<b>Parameters</b>	<i>sdp-id</i> — Specifies the IP address.
<b>Values</b>	1 — 17407

## service-carving

<b>Syntax</b>	<b>service-carving</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment
<b>Description</b>	The service-carving algorithm determines the PE that is the Designated Forwarder (DF) in a given ethernet-segment and for a given service. This command enables the context to configure service-carving in the ethernet-segment.

## mode

<b>Syntax</b>	<b>mode {manual   auto   off}</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment>service-carving
<b>Description</b>	This command configures the service-carving mode. This determines how the DF is elected for a given ethernet-segment and service.
<b>Default</b>	<b>mode auto</b>
<b>Parameters</b>	<b>auto</b> — This mode is the service-carving algorithm defined in RFC7432. 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. <b>manual</b> — In this mode the DF is elected based on the manual configuration added in the <b>service-carving&gt;manual</b> context. <b>off</b> — 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

<b>Syntax</b>	<b>manual</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment>service-carving
<b>Description</b>	This command enables the context to configure service-carving in a manual way, that is, configuring the evis or isids for which the PE is DF

## evi

<b>Syntax</b>	<b>evi start [to to] primary</b> <b>no evi</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment>service-carving>manual
<b>Description</b>	This command configures the evi ranges for which the PE is DF. Note that multiple individual evi values and multiple evi ranges are allowed. The PE will be non-DF for the evi values not defined as <b>primary</b> .
<b>Parameters</b>	<i>start</i> — This specifies the initial evi value of the range for which the PE is DF. <b>Values</b> 1 — 65535 <i>to</i> — This specifies the end evi value of the range for which the PE is DF. If not configured, only the individual start value will be considered. <b>Values</b> 1 — 65535 <b>primary</b> — Specifies that the PE is DF for the configured evi range.

## isid

<b>Syntax</b>	<b>isid start [to to] primary</b> <b>no sid</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment>service-carving>manual
<b>Description</b>	This command configures the <b>isid</b> ranges for which the PE is DF. Note that multiple individual <b>isid</b> values and multiple isid ranges are allowed. The PE will be non-DF for <b>isid</b> values not defined as <b>primary</b> .
<b>Parameters</b>	<i>start</i> — This specifies the initial <b>isid</b> value of the range for which the PE is DF. <b>Values</b> 1 — 16777215 <i>to</i> — This specifies the end <b>isid</b> value of the range for which the PE is DF. If not configured, only the individual start value will be considered. <b>Values</b> 1 — 16777215 <b>primary</b> — Specifies that the PE is DF for the configured <b>evi</b> range.

## shutdown

<b>Syntax</b>	<b>shutdown</b> <b>no shutdown</b>
<b>Context</b>	config>service>system>bgp-evpn>ethernet-segment
<b>Description</b>	This command changes the administrative status of the ethernet-segment.

## EVPN Commands

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 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. Note that the es-bmac will consume a separate BMAC per B-VPLS that is linked to an ethernet-segment

**Parameters**    *MACLsb* — This specifies the two least significant bytes of the es-bmac.

**Values**    1 — 65535 or xx-xx or xx:xx

*size* — This specifies the reserved space in the FDB for a given es-bmac. By default the system reserves 8 entries for a given ethernet-segment BMAC.

**Values**    1 — 511999

**Default**    8

### 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 **es-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. Once 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— 600 seconds

## 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 given ethernet-segment on a given 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**

## EVPN Commands

**Parameters** *seconds* — Specifies the number of seconds for the **es-activation-timer**.

**Values** 0— 100 seconds

### cfm-mac-advertisement

**Syntax** **cfm-mac-advertisement**  
**[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 & 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.

Note that you should 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

**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

Note that if vsi-import/export policies 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 shown in the **show service id bgp** command.

**Parameters** *value* — Specifies the *evi*.

**Values** 1 — 65535

## ip-route-advertisement

<b>Syntax</b>	<b>ip-route-advertisement</b> [ <b>incl-host</b> ] <b>no ip-route-advertisement</b>
<b>Context</b>	config>service>vpls>bgp-evpn
<b>Description</b>	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. Note that the interface host addresses are not advertised in EVPN unless the <b>ip-route-advertisement incl-host</b> command is enabled.
<b>Default</b>	no ip-route-advertisement
<b>Parameters</b>	<b>incl-host</b> — Specifies to advertise the interface host addresses in EVPN

## mac-advertisement

<b>Syntax</b>	<b>[no] mac-advertisement</b>
<b>Context</b>	config>service>vpls>bgp-evpn
<b>Description</b>	The mac-advertisement command enables the advertisement in BGP of the learnt macs on SAPs and SDP bindings. When the mac-advertisement is disabled, the local macs will be withdrawn in BGP.
<b>Default</b>	mac-advertisement

## mac-duplication

<b>Syntax</b>	<b>mac-duplication</b>
<b>Context</b>	config>service>vpls>bgp-evpn
<b>Description</b>	This command enables the context to configure the BGP EVPN mac duplication parameters.

## detect

<b>Syntax</b>	<b>detect num-moves</b> <i>num-moves</i> <b>window</b> <i>minutes</i>
<b>Context</b>	config>service>vpls>bgp-evpn>mac-duplication
<b>Description</b>	The <b>mac-duplication</b> featured is always enabled by default. This command modifies the default behavior. <b>mac-duplication</b> monitors the number of moves of a MAC address for a period of time (window).
<b>Default</b>	num-moves 5 window 3
<b>Parameters</b>	<b>num-moves</b> <i>num-moves</i> — Identifies the number of MAC moves in a VPLS service. The counter is incremented when a given 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.

## EVPN Commands

**Values** 3..10 minutes

**Default** 3 minutes

**window** *minutes* — Specifies the length of the window in minutes.

**Values** 1 — 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, once 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 BGP EVPN MAC duplication retry in minutes.

**Values** 2 — 60 minutes

### mpls

**Syntax** **mpls**

**Context** config>service>vpls>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

**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.

The following tunnel types are supported in a BGP-EVPN MPLS context in order of preference: RSVP, LDP, Segment Routing (SR), and BGP.

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 **bgp** value instructs BGP EVPN to search for a BGP LSP to the address of the BGP next-hop.

The user must set **resolution to filter** to activate the list of tunnel-types configured under **resolution-filter**.

## resolution

<b>Syntax</b>	<b>resolution {disabled any filter}</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel
<b>Description</b>	This command configures the resolution mode in the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.
<b>Parameters</b>	<p><b>any</b> — enables the binding to any supported tunnel type in a BGP-EVPN MPLS context following TTM preference.</p> <p><b>disabled</b> — disables the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.</p> <p><b>filter</b> — enables the binding to the subset of tunnel types configured under <b>resolution-filter</b>.</p>

## resolution-filter

<b>Syntax</b>	<b>resolution-filter</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel
<b>Description</b>	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.

## EVPN Commands

The following tunnel types are supported in a BGP-EVPN MPLS context in order of preference: RSVP, LDP, Segment Routing (SR), and BGP.

### bgp

<b>Syntax</b>	<b>[no] bgp</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
<b>Description</b>	Selects the BGP tunnel type.

### ldp

<b>Syntax</b>	<b>[no] ldp</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
<b>Description</b>	Selects the LDP tunnel type.

### rsvp

<b>Syntax</b>	<b>[no] rsvp</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
<b>Description</b>	Selects the RSVP-TE tunnel type.

### sr-isis

<b>Syntax</b>	<b>[no] sr-isis</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
<b>Description</b>	Selects the Segment Routing (SR) tunnel type programmed by an ISIS instance in TTM..

### sr-ospf

<b>Syntax</b>	<b>[no] sr-ospf</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
<b>Description</b>	Selects the Segment Routing (SR) tunnel type programmed by an OSPF instance in TTM.

## control-word

<b>Syntax</b>	<b>control-word</b> <b>no control-word</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls
<b>Description</b>	This command enables the transmission and reception of the <b>control-word</b> . 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.
<b>Default</b>	<b>no control-word</b>

## ecmp

<b>Syntax</b>	<b>ecmp</b> <i>value</i>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls
<b>Description</b>	This command controls the number of paths to reach a given 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 given MAC enables the 'aliasing' function described in RFC7432.
<b>Parameters</b>	<i>value</i> — 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-segmen.
	<b>Values</b> 0 — 32
	<b>Default</b> 0

## force-vlan-vc-forwarding

<b>Syntax</b>	<b>force-vlan-vc-forwarding</b> <b>no force-vlan-vc-forwarding</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls
<b>Description</b>	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.  Note that this command may be used in conjunction with the <b>sap ingress vlan-translation</b> 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.
<b>Default</b>	<b>no force-vlan-forwarding</b>

## ingress-replication-bum-label

<b>Syntax</b>	<b>[no] no-ingress-replication-bum-label</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls
<b>Description</b>	<p>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 given service. By default (<b>no ingress-replication-bum-label</b>), the same label is used for unicast and flooded BUM packets when forwarding traffic to remote PEs.</p> <p>When saving labels, this might cause transient traffic duplication for all-active multi-homing. By enabling <b>ingress-replication-bum-label</b>, 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.</p>
<b>Default</b>	<b>no ingress-replication-bum-label</b>

## shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls
<b>Description</b>	This command controls the administrative state of EVPN-MPLS in the service.

## split-horizon-group

<b>Syntax</b>	<b>split-horizon-group <i>name</i></b> <b>no split-horizon-group</b>
<b>Context</b>	config>service>vpls>bgp-evpn>mpls
<b>Description</b>	<p>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.</p> <p>If the <b>split-horizon-group</b> command for <b>bgp-evpn&gt;mpls&gt;</b> 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 <b>bgp-evpn&gt;mpls&gt; split-horizon-group</b> will only be allowed if <b>bgp-evpn&gt;mpls</b> is shutdown; no changes are allowed when <b>bgp-evpn&gt;mpls</b> is <b>no shutdown</b>.</p>

When the SAPs or/and 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

**Parameters** *name* — Specifies the split-horizon-group name.

**Default** **no split-horizon-group**

## 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 learnt 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. Note that, although the 7x50 can be configured to generate and advertise the unknown-mac-route, the 7x50 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**  
**no vxlan vni**

**Context** config>service>vpls

**Description** This command enables the use of vxlan in the VPLS service.

**Parameters** **vni vni-id** — Specifies the VXLAN network identifier configured in the VPLS service. All the 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 — 16777215

Note that the VPLS service will be operationally UP once 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.

## EVPN Commands

### vxlan

<b>Syntax</b>	<b>vxlan</b>
<b>Context</b>	config>service>vpls>bgp-evpn
<b>Description</b>	This command enables the context to configure the VXLAN parameters when BGP EVPN is used as the control plane.

### shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>service>vpls>bgp-evpn>vxlan
<b>Description</b>	This command enables/disables the automatic creation of VXLAN auto-bindings by BGP-EVPN.
<b>Default</b>	shutdown

### pbb

<b>Syntax</b>	<b>pbb</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command enables the context where the PBB parameters are configured.

### use-es-bmac

<b>Syntax</b>	<b>use-es-bmac</b>
<b>Context</b>	config>service>vpls>pbb
<b>Description</b>	<p>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.</p> <p>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 <b>config&gt;service&gt;system&gt;bgp-evpn&gt;ethernet-segment&gt;source-bmac-lsb</b> for all the packets coming from the local ethernet-segment.</p> <p>If <b>no use-es-bmac</b> is configured, the system will use the regular source-bmac (provided by the <b>config&gt;service&gt;vpls&gt;pbb&gt;source-bmac</b> command, or the chassis bmac if the source-bmac is not configured).</p>
<b>Default</b>	<b>no use-es-bmac</b>

## proxy-arp

<b>Syntax</b>	<b>proxy-arp</b> <b>no proxy-arp</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command enables the context to configure the proxy-ARP parameters in a VPLS service.
<b>Default</b>	no proxy-arp

## proxy-nd

<b>Syntax</b>	<b>proxy-nd</b> <b>[no] proxy-nd</b>
<b>Context</b>	config>service>vpls
<b>Description</b>	This command enables the context to configure the proxy-ND parameters in a VPLS service.
<b>Default</b>	no proxy-arp

## age-time

<b>Syntax</b>	<b>[no] age-time seconds</b>
<b>Context</b>	config>service>vpls>proxy-arp config>service>vpls>proxy-nd
<b>Description</b>	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 <b>age-time</b> should be set at <i>send-refresh</i> * 3 to ensure that no active entries are unnecessarily removed.
<b>Default</b>	no age-time
<b>Parameters</b>	<i>seconds</i> — Specifies the age-time in seconds.
<b>Values</b>	60— 86400

## dup-detect

<b>Syntax</b>	<b>dup-detect [anti-spoof-mac mac-address] window minutes num-moves count hold-down [minutes max]</b>
<b>Context</b>	config>service>vpls>proxy-arp config>service>vpls>proxy-nd

<b>Description</b>	<p>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 <b>window</b> <i>&lt;minutes&gt;</i>. When <i>&lt;count&gt;</i> is reached within that <b>window</b>, the proxy-ARP/ND entry for the suspected IP is marked as duplicate. An alarm is also triggered. This condition is cleared when <b>hold-down</b> time expires (max does not expire) or a <b>clear</b> command is issued.</p> <p>If the <b>anti-spoof-mac</b> is configured, the proxy-ARP/ND offending entry's MAC is replaced with this <i>&lt;mac-address&gt;</i> and advertised in an unsolicited GARP/NA for local SAP/SDP-bindings, and in EVPN to remote PEs. This mechanism assumes that the same <b>anti-spoof-mac</b> is configured in all the PEs for the same service and that traffic with destination <b>anti-spoof-mac</b> received on SAPs/SDP-bindings will be dropped. An ingress <b>mac-filter</b> must be configured in order to drop traffic to the <b>anti-spoof-mac</b>.</p>
<b>Default</b>	<b>dup-detect window 3 num-moves 5 hold-down 9</b>
<b>Parameters</b>	<p><i>minutes</i> — Specifies the window size in minutes.</p> <p style="margin-left: 2em;"><b>Values</b> 1— 15</p> <p style="margin-left: 2em;"><b>Default</b> 3</p> <p><i>count</i> — Specifies the number of moves required so that an entry is declared duplicate.</p> <p style="margin-left: 2em;"><b>Values</b> 3— 10</p> <p style="margin-left: 2em;"><b>Default</b> 5</p> <p><i>minutes max</i> — Specifies the hold-down time for a duplicate entry. Max means permanent hold-down.</p> <p style="margin-left: 2em;"><b>Values</b> 2— 60 max</p> <p style="margin-left: 2em;"><b>Default</b> 9</p> <p><i>mac-address</i> — Specifies the optional anti-spoof-mac to use.</p>

## dynamic-arp-populate

<b>Syntax</b>	<b>[no] dynamic-arp-populate</b>
<b>Context</b>	config>service>vpls>proxy-arp
<b>Description</b>	<p>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.</p> <p>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.</p>
<b>Default</b>	no dynamic-arp-populate

## garp-flood-evpn

<b>Syntax</b>	<b>[no] garp-flood-evpn</b>
<b>Context</b>	config>service>vpls>proxy-arp
<b>Description</b>	<p>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.</p> <p>The <b>no</b> form of the command only floods to local saps/binds but not to EVPN destinations.</p> <p>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.</p>
<b>Default</b>	garp-flood-evpn

## send-refresh

<b>Syntax</b>	<b>[no] send-refresh seconds</b>
<b>Context</b>	config>service>vpls>proxy-arp config>service>vpls>proxy-nd
<b>Description</b>	<p>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.</p>
<b>Default</b>	no send-refresh
<b>Parameters</b>	<i>seconds</i> — Specifies the send-refresh in seconds.
	<b>Values</b> 120— 86400

## static

<b>Syntax</b>	<b>static ip-address ieee-address</b> <b>[no] static ip-address</b>
<b>Context</b>	config>service>vpls>proxy-arp
<b>Description</b>	<p>This command configures static entries to be added to the table. Note that a static MAC-IP entry requires the addition of the MAC address to the FDB as either learnt or CStatic (conditional static mac) in order to become active.</p>
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the IPv4 address for the static entry.</p> <p><i>ieee-address</i> — Specifies the MAC address for the static entry.</p>

## table-size

<b>Syntax</b>	<b>table-size</b> <i>table-size</i>
<b>Context</b>	config>service>vpls>proxy-arp config>service>vpls>proxy-nd
<b>Description</b>	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 given IP can be replaced (a different MAC can be learnt and added) but no new IP entries will be added, regardless of the type (Static, evpn, dynamic). If the user attempts to change the <b>table-size</b> value to a value that cannot accommodate the number of existing entries, the attempt will fail.
<b>Default</b>	<b>table-size</b> 250
<b>Parameters</b>	<i>table-size</i> — Specifies the table-size as number of entries for the service.
<b>Values</b>	1— 16384

## unknown-arp-request-flood-evpn

<b>Syntax</b>	<b>[no] unknown-arp-request-flood-evpn</b>
<b>Context</b>	config>service>vpls>proxy-arp
<b>Description</b>	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 <b>no</b> form of the command will only flood to local SAPs/SDP-bindings and not to EVPN destinations.
<b>Default</b>	unknown-arp-request-flood-evpn

## dynamic-nd-populate

<b>Syntax</b>	<b>[no] dynamic-nd-populate</b>
<b>Context</b>	config>service>vpls>proxy-nd
<b>Description</b>	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 <b>dynamic-nd-populate</b> 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 two different functions: on the one hand it enables the advertisement of static or dynamic entries that are learnt as host or routers (only one option is possible for a given service). On the other hand, it 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

## 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.

## EVPN Commands

**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. Note that 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 the MAC address for the static entry.

**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 kpet in the table as *inactive*, regardless of their previous *Status*.

**Default** shutdown

## 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] 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.

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). It cannot be all zeros.

**create** — This keyword is mandatory while creating a static MAC.

**monitor fwd-status** — Specifies that this static mac is based on the forwarding status of the SAP or spoke SDP for multi-homed operation.

## EVPN Commands

### evpn-tunnel

<b>Syntax</b>	<b>[no] evpn-tunnel</b>
<b>Context</b>	config>service>vprn>interface>vpls
<b>Description</b>	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.
<b>Default</b>	no evpn-tunnel

### vsd-domain

<b>Syntax</b>	<b>vsd-domain <i>name</i></b> <b>no vsd-domain</b>
<b>Context</b>	config>service>vpls config>service>vprn
<b>Description</b>	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 7x50 to correlate configuration parameters to a service.
<b>Parameters</b>	<i>name</i> — Specifies the vsd-domain name.

### vsd

<b>Syntax</b>	<b>vsd</b>
<b>Context</b>	config>service config>service
<b>Description</b>	This command provides the context for the vsd configuration.

### domain

<b>Syntax</b>	<b>domain <i>name</i> [type {l2-domain vrf-gre vrf-vxlan l2-domain-irb}] [create]</b> <b>[no] domain <i>name</i></b>
<b>Context</b>	config>service>vsd
<b>Description</b>	This command configures a vsd-domain that can be associated to a VPLS or VPRN service.
<b>Parameters</b>	<b>type</b> — specifies the type of domain. Vrf-gre can only be associated to a VPRN service. The other three types of domains must be associated to a VPLS service. <b>Values</b> l2-domain   vrf-gre   vrf-vxlan   l2-domain-irb <b>create</b> — Creates the vsd-domain.

## description

<b>Syntax</b>	<b>description</b> <i>description-string</i>
<b>Context</b>	config>service>vsd>domain
<b>Description</b>	This command provides a description for a vsd-domain. This description must be added before the domain can be no shutdown.
<b>Parameters</b>	<b>description</b> — Specifies the text for the description.

## service-range

<b>Syntax</b>	<b>service-range</b> <i>svc-id to svc-id</i> <b>[no] service-range</b>
<b>Context</b>	config>service>vsd
<b>Description</b>	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
<b>Parameters</b>	<i>svc-id</i> — specifies the start and end service identifier values. <b>Values</b> 1— 2147483647

## shutdown

<b>Syntax</b>	<b>shutdown</b> <b>[no] shutdown</b>
<b>Context</b>	config>service>vsd>domain
<b>Description</b>	This command enables or disables a domain. A description must be provided before no shutdown is executed.

## system-id

<b>Syntax</b>	<b>system-id</b> <i>name</i> <b>[no] system-id</b>
<b>Context</b>	config>system>vsd
<b>Description</b>	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.
<b>Parameters</b>	<i>name</i> — Specifies the name.

### xmpp

<b>Syntax</b>	<b>xmpp</b>
<b>Context</b>	config>system
<b>Description</b>	This command provides the context for the xmpp configuration.

### server

<b>Syntax</b>	<b>server</b> <i>xmpp-server-name</i> [ <b>domain-name</b> <i>fqdn</i> ] [ <b>username</b> <i>user-name</i> ] [ <b>password</b> <i>password</i> ] [ <b>create</b> ] <b>[no] server</b> <i>xmpp-server-name</i>
<b>Context</b>	config>system>xmpp
<b>Description</b>	This command configures the XMPP server as well as the Jabber ID that the 7x50 will use for the XMPP communication with the server. Note that the system uses DNS to resolve the configured domain-name. <b>no server</b> <i>name</i> will remove all the dynamic configurations in all the services.
<b>Parameters</b>	<i>xmpp-server-name</i> — Specifies the name of the server in lower-case letters. <i>fqdn</i> — Specifies the Fully Qualified Domain Name of the server. <i>user-name</i> — Specifies the user-name part of the Jabber ID. <i>password</i> — Specifies the password part of the Jabber ID's user. <b>create</b> — keyword used to create the server instance.

### shutdown

<b>Syntax</b>	<b>shutdown</b> <b>[no] shutdown</b>
<b>Context</b>	config>system>xmpp>server
<b>Description</b>	This command enables or disables the communication with a given XMPP server. When the xmpp server is properly configured, <b>no shutdown</b> instructs the system to establish a TCP session with the XMPP server through the management interface first. If it fails to establish communication, the 7x50 uses an in-band communication and its system IP as source IP address. Shutdown does not remove the dynamic configurations.

## security

<b>Syntax</b>	<b>security</b>
<b>Context</b>	config>system
<b>Description</b>	This command enables the context for the configuration of the security parameters in the system.

## cli-script

<b>Syntax</b>	<b>cli-script</b>
<b>Context</b>	config>system>security
<b>Description</b>	This command enables the context for the configuration of the security parameters in the system.

## authorization

<b>Syntax</b>	<b>authorization</b>
<b>Context</b>	config>system>security>cli-script
<b>Description</b>	This command enables the context for the configuration of the authorization parameters for the cli-scripts in the system.

## vsd

<b>Syntax</b>	<b>vsd</b>
<b>Context</b>	config>system>security>cli-script>authorization
<b>Description</b>	This command enables the context for the configuration of the authorization parameters related to VSD in the system.

## cli-user

<b>Syntax</b>	<b>cli-user <i>user-name</i></b> <b>no cli-user</b>
<b>Context</b>	config>system>security>cli-script>authorization>vsd
<b>Description</b>	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 <b>tools dump service vsd-services command-list</b> to check the white-list of commands.

## EVPN 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* | *hash2* — Specifies the hashing sequence.

### 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

<b>Syntax</b>	<b>group</b> <i>name</i>
<b>Context</b>	config>router>bgp
<b>Description</b>	This command enables the context for the configuration of a bgp group in the base router.

## neighbor

<b>Syntax</b>	<b>neighbor</b> <i>ip-address</i>
<b>Context</b>	config>router>bgp>group
<b>Description</b>	This command enables the context for the configuration of a bgp group neighbor in the base router.

## def-recv-evpn-encap

<b>Syntax</b>	<b>def-recv-evpn-encap</b> { <b>mpls</b>   <b>vxlan</b> }
<b>Context</b>	config>router>bgp>group>neighbor
<b>Description</b>	This command defines how the BGP will treat a received EVPN route without RFC5512 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.
<b>Default</b>	no def-recv-evpn-encap
<b>Parameters</b>	<b>mpls</b> — Specifies that <b>mpls</b> is the default encapsulation value in the case where no RFC5512 extended community is received in the incoming BGP-EVPN route. <b>vxlan</b> — Specifies that <b>vxlan</b> is the default encapsulation value.

## python

<b>Syntax</b>	<b>python</b>
<b>Context</b>	config
<b>Description</b>	This command enables the context for the configuration of the python parameters in the system.

## python-policy

<b>Syntax</b>	<b>python-policy</b> <i>name</i>
<b>Context</b>	config>python

## EVPN Commands

**Description** This command enables the context for the configuration of the python policy parameters in the system.

### 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** **enable-vsd-config**  
**[no] enable-vsd-config**

**Context**

**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**.

## Show Commands

### service-using

<b>Syntax</b>	<b>service-using [vsd]</b> <b>service-using [origin vsd]</b>
<b>Context</b>	show>service
<b>Description</b>	When the <b>vsd</b> modifier is used, this command displays the VSD domain tags used and the associated service identifier. If the modifier <b>origin vsd</b> is used, the command displays the services created by the VSD fully-dynamic integration model. (Python will actually create the service after receiving the relevant parameters from VSD).

#### 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

<b>Syntax</b>	<b>system</b>
<b>Context</b>	show>service
<b>Description</b>	This command enables the context to display the system parameters.

**Sample Output**

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.

**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

**Context** show>service>system>bgp-evpn

## Show Commands

**ethernet-segment name *name* [all]**  
**ethernet-segment name *name evi* [evi]**  
**ethernet-segment name *name isid* [isid]**

**Description** This command enables the context to display the ethernet-segment parameters.

## 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.

### 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 Evt : 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

<b>Syntax</b>	<b>root-objects</b>
<b>Context</b>	show>service>vsd
<b>Description</b>	This command displays the root objects created by vsd.

### 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
=====

```

## Show Commands

### 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

#### 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.

**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.

**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

```

## Show Commands

Number of domain entries: 1

=====

## bgp-evpn

- Syntax** `bgp-evpn`
- Context** `show>service>id`
- Description** This command displays the bgp-evpn configured parameters for a given 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) has been enabled. . If the service is bgp-evpn mpls, the command will show the parameters corresponding to evpn-mpls.

### 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)
=====

```

## 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 given 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.

### Sample Output

```

*A:PE1# show service id 1 evpn-mpls

=====
BGP EVPN-MPLS Dest
=====
TEP Address Egr Label Num. MACs Mcast Last Change

```

# Show Commands

```

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 as well as the BGP-EVPN MPLS destinations associated to them.

**Sample Output**

```

Add this sample:

*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
=====

```

## Show Commands

```
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.

### 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.

**Sample Output**

```
A:PE1# show service id 301 es-pbr

=====
L2 ES PBR
=====
ESI Users Status
 VTEP:VNI

ff:00:00:00:00:00:00:00:01 1 Active
 192.0.2.72:7272

Number of entries : 1

=====
```

## proxy-arp

- Syntax** **proxy-arp** *ip-address* [**detail**]
- Context** show>service>id
- Description** This command displays the proxy-ARP entries existing for a particular service. This table is populated by the EVPN mac routes containing a MAC and an IP address , as well as static entries or dynamic entries from snooped ARP messages on access SAP/SDP-bindings. A 7x50 receiving an ARP request from a SAP or SDP-binding will perform a lookup in the proxy-arp table for the service. If the 7x50 finds a match, it will reply to the ARP and will not let the ARP be flooded in the VPLS service. If the 7x50 does not find a match, the ARP will be flooded within the service if the configuration allows it. The command allows for an specific IP addresses to be shown.

**Sample Output**

```
*A:DutA# show service id 1 proxy-arp

Proxy Arp

Admin State : enabled
Dyn Populate : enabled
Age Time : disabled Send Refresh : 120 secs
Table Size : 250 Total : 2
Static Count : 0 EVPN Count : 0
```

## Show Commands

```
Dynamic Count : 2 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

=====

*A:DutA# show service id 1 proxy-arp detail

Proxy Arp

Admin State : enabled
Dyn Populate : enabled
Age Time : disabled Send Refresh : 120 secs
Table Size : 250 Total : 2
Static Count : 0 EVPN Count : 0
Dynamic Count : 2 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

=====

VPLS Proxy Arp Entries

IP Address Mac Address Type Status Last Update

10.10.10.1 00:ca:ca:ba:ca:01 dyn active 07/15/2015 19:53:31
10.10.10.3 00:ca:ca:ba:ca:03 dyn active 07/15/2015 19:53:21

Number of entries : 2
=====
```

## vxlan

- Syntax** **vxlan**  
**vxlan [vtep ip-address]**
- Context** show>service>id  
show>service
- Description** This command displays the VXLAN bindings auto-created in a given 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.

### Sample Output

```
*A:DutAA# show service id 101 vxlan
show service id 101 vxlan
=====
VPLS VXLAN, Ingress VXLAN Network Id: 101

=====
Egress VTEP, VNI
=====
VTEP Address Egress VNI Num. MACs Mcast Oper State L2 PBR

192.0.2.71 101 1 No Up No

Number of Egress VTEP, VNI : 1

A:DutB# show service vxlan <vtep> 192.0.2.65 192.0.2.66
A:DutB# show service vxlan 192.0.2.65
=====
VXLAN Tunnel Endpoint: 192.0.2.65
=====
Egress VNI Service Id Oper State

60 60 Up

=====
```

## evpn-mpls

**Syntax** `evpn-mpls [TEP ip-address]`

**Context** `show>service`

**Description** This command shows the remote EVPN-MPLS tunnel endpoints in the system.

### 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
```

## Show Commands

```
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

=====
```

## 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.

### Sample Output

```
:sr12U-46-PE2# show system xmpp server
=====
XMPP Server Table
```

```

=====
Name User Name State
XMPP FQDN Last State chgd Admin State

vsdl-hy cspTest Functional
vsdl-hy.alu-srpm.us 0d 22:42:15 inService

No. of XMPP server's: 1
=====
B:Dut# show system xmpp server "vsdl-hy"
=====
XMPP Server Table
=====
XMPP FQDN : vsdl-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 as shown in the show system xmpp vsd command, a detailed view for that specific server is shown, including statistics.

### 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

```

## Show Commands

```
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.

### 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
```

## 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.

#### Sample Output

```
*A:PE2# show redundancy bgp-evpn-multi-homing
```

## Show Commands

```
=====
Redundancy BGP EVPN Multi-homing Information
=====
Boot-Timer : 10 secs
Boot-Timer Remaining : 0 secs
ES Activation Timer : 3 secs
=====
```

---

## Clear Commands

### domain

<b>Syntax</b>	<b>domain</b> [ <i>name</i> ]
<b>Context</b>	clear>service>statistics>vsd
<b>Description</b>	This command clears the statistics shown in the <b>show service vsd domain</b> <i>name</i> command.
<b>Parameters</b>	<i>name</i> — specifies the vsd domain name.

### scripts

<b>Syntax</b>	<b>scripts</b>
<b>Context</b>	clear>service>statistics>vsd
<b>Description</b>	This command clears the statistics shown in the <b>show service vsd script statistics</b> command.

### server

<b>Syntax</b>	<b>server</b> [xmpp-server-name]
<b>Context</b>	clear>system>statistics>xmpp
<b>Description</b>	This command clears the statistics shown in the <b>show system xmpp server</b> <i>name</i> command.
<b>Parameters</b>	<i>xmpp-server-name</i> — specifies the vsd domain name

### ver

<b>Syntax</b>	<b>server</b> [xmpp-server-name]
<b>Context</b>	clear>system>statistics>xmpp
<b>Description</b>	This command clears the statistics shown in the <b>show system xmpp server</b> <i>name</i> command.
<b>Parameters</b>	<i>xmpp-server-name</i> — specifies the vsd domain name

---

## Tools Commands

### service

<b>Syntax</b>	<b>service</b>
<b>Context</b>	tools>dump
<b>Description</b>	Use this command to configure tools to display service dump information.

### id

<b>Syntax</b>	<b>id <i>service-id</i></b>
<b>Context</b>	tools>dump
<b>Description</b>	Use this command to configure parameters to display service ID information.

### vxlan

<b>Syntax</b>	<b>vxlan [clear]</b>
<b>Context</b>	tools>dump>service>id
<b>Description</b>	<p>This command displays the number of times a service could not add a VXLAN binding or &lt;VTEP, Egress VNI&gt; due to the following limits:</p> <ul style="list-style-type: none"> <li>The per System VTEP limit has been reached</li> <li>The per System &lt;VTEP, Egress VNI&gt; limit has been reached</li> <li>The per Service &lt;VTEP, Egress VNI&gt; limit has been reached</li> <li>The per System Bind limit: Total bind limit or vxlan bind limit has been reached.</li> </ul> <p>The command adds a <b>clear</b> option to clear the above statistics.</p>

#### 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.

### 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 given R-VPLS service.

### 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 given VSD to a vsd-domain.

### Sample Output

## Tools Commands

```
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.

### fd-domain-sync

- Syntax** **fd-domain-sync {full | diff}**
- Context** tools>perform>service>vsd
- Description** This command instructs the system to audit the VSD to get the "DIFF" list of even the "FULL" list of all the do-mains in the VSD .

### evaluate-script

- Syntax** **evaluate-script domain-name** [64 chars max] **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 terardown 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 .

#### 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 tar-
get:64000:64000 metadata "rd=1:1, sap=1/1/10:3000"
```

Success

## name

- Syntax** `name [name] refresh-config`
- Context** `tools>perform>service>vsd>domain`
- Description** This command instructs the system to refresh the configuration of a given domain immediately instead of waiting for the next audit interval.

## 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 given evi or isid.

### 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

<b>Syntax</b>	<b>usage</b>
<b>Context</b>	tools>dump>service>evpn
<b>Description</b>	This command displays the consumed VXLAN EVPN resources in the system.

### 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
Vxlan Dests (TEP, Egress VNI) : 2
Total-Dest : 18/131071

Sdp Bind + Evpn Dests : 20/196607
RVPLS Egress VNI : 1/40959
ES L2/L3 PBR : 0/ 32767
```

## vsd-services

<b>Syntax</b>	<b>vsd-services</b>
<b>Context</b>	tools>dump>service
<b>Description</b>	This command enables the context for vsd-services commands..

## command-list

<b>Syntax</b>	<b>command-list</b>
<b>Context</b>	tools>dump>service>vsd-services
<b>Description</b>	<p>This command displays the list of CLI nodes allowed in the VSD fully dynamic provisioning model. Python will have access to the shown nodes.</p> <p>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.</p>

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 rout-er policy-options prefix-list x prefix 0.0.0.0/0** - since you can have multiple instances of prefixes.

---

## Debug Commands

### xmpp

<b>Syntax</b>	<b>xmpp [connection] [gateway] [message] [vsd] [iq] [all] [no] xmpp</b>
<b>Context</b>	debug>system
<b>Description</b>	This command enables the debug for XMPP messages sent or received by the 7x50.
<b>Parameters</b>	<b>connection</b> — filters only the messages related to the XMPP connection. <b>gateway</b> — Filters the messages related to the gateway. <b>message</b> — Filters only the messages. <b>vsd</b> — Filters the vsd messages. <b>iq</b> — Filters the IQ messages between the gateway and the vsd. <b>all</b> — Includes all the above.

### vsd

<b>Syntax</b>	<b>vsd</b>
<b>Context</b>	debug
<b>Description</b>	This command enables the context for the debug vsd commands.

### scripts

<b>Syntax</b>	<b>scripts</b> <b>scripts event [cli] [errors] [executed-cmd] [state-change] [warnings]</b> <b>scripts instance <i>instance</i> event [cli] [errors] [executed-cmd] [state-change] [warnings]</b>
<b>Context</b>	debug>vsd
<b>Description</b>	This command enables the debug of the VSD fully dynamic integration scripts.

### event

<b>Syntax</b>	<b>[no] event</b>
<b>Context</b>	debug>vsd>script

**Description** This command enables/disables the generation of all script debugging event output: cli, errors, executedcmd, warnings, state-change.

## instance

**Syntax** **[no] instance** *instance*

**Context** debug>vsd>script

**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>script>event  
debug>vsd>script>instance

**Description** This command enables/disables the generation of a specific script debugging event output: **cli**

## errors

**Syntax** **[no] errors**

**Context** debug>vsd>script>event  
debug>vsd>script>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>script>event  
debug>vsd>iscript>nstance

**Description** This command enables/disables the generation of a specific script debugging event output: **execute-cmd**.

## Debug Commands

### state-change

<b>Syntax</b>	<b>[no] state-change</b>
<b>Context</b>	debug>vsd>script>event debug>vsd>script>instance
<b>Description</b>	This command enables/disables the generation of a specific script debugging event output: <b>state-change</b> .

### warnings

<b>Syntax</b>	<b>[no] warnings</b>
<b>Context</b>	debug>vsd>script>event debug>vsd>script>instance
<b>Description</b>	This command enables/disables the generation of a specific script debugging event output: <b>warnings</b> .

# Common CLI Command Descriptions

---

## In This Chapter

This section provides information about common Command Line Interface (CLI) syntax and command usage.

Topics in this chapter include:

- [SAP syntax on page 1272](#)

## Common Service Commands

### sap

**Syntax** [no] sap *sap-id*

**Description** This command specifies the physical port identifier portion of the SAP definition.

**Parameters** *sap-id* — Specifies the physical port identifier portion of the SAP definition.

The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
null	[ <i>port-id</i>   <i>lag-id</i> ]	<i>port-id</i> : 1/1/3 <i>lag-id</i> : lag-3
dot1q	[ <i>port-id</i>   <i>lag-id</i> ]:qtag1	<i>port-id</i> :qtag1: 1/1/3:100 <i>lag-id</i> :qtag1:lag-3:102
qinq	[ <i>port-id</i>   <i>lag-id</i> ]:qtag1.qtag2	<i>port-id</i> :qtag1.qtag2: 1/1/3:100.10 <i>lag-id</i> :qtag1.qtag2: lag-10:

The values depend on the encapsulation type configured for the interface. The following table describes the allowed values for the port and encapsulation types.

Port Type	Encap-Type	Allowed Values	Comments
Ethernet	Null	0	The SAP is identified by the port.
Ethernet	Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the port. Note that a 0 qtag1 value also accepts untagged packets on the dot1q port.
Ethernet	QinQ	qtag1: *   0 — 4094 qtag2: *   null   0 — 4094	The SAP is identified by two 802.1Q tags on the port. <b>Note</b> — The following combinations of qtag1.qtag2 accept untagged packets: “0.*”, “*.null”, “*.*”.
SONET/SDH	IPCP	-	The SAP is identified by the channel. No BCP is deployed and all traffic is IP.
SONET/SDH	BCP-Null	0	The SAP is identified with a single service on the channel. Tags are assumed to be part of the customer packet and not a service delimiter.

SONET/SDH    BCP-Dot1q    0 — 4094

The SAP is identified by the 802.1Q tag on the channel.



# Standards and Protocol Support



**Note:** The information presented is subject to change without notice.

Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

## ANCP/L2CP

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

## 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*

## BGP

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 5575, *Dissemination of Flow Specification Rules*
- RFC 5668, *4-Octet AS Specific BGP Extended Community*
- draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*
- draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

### Circuit Emulation

- MEF-8, *Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004*
- 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, *Ethernet Linear Protection Switching*
- ITU-T G.8032, *Ethernet Ring Protection Switching*
- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

### EVPN

- RFC7432, *BGP MPLS-Based Ethernet VPN*
- draft-ietf-bess-evpn-overlay-01, *A Network Virtualization Overlay Solution using EVPN*
- draft-ietf-bess-evpn-prefix-advertisement-01, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-vpls-seamless-integ-00, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*  
 draft-ietf-l2vpn-pbb-evpn-10, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*  
 draft-snr-bess-evpn-proxy-arp-nd-00, *Proxy-ARP/ND function in EVPN networks*

## Fast Reroute

RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*  
 RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*  
 draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*  
 draft-kastran-mofir-02, *Multicast only Fast Re-Route*

## Frame Relay

ANSI T1.617 Annex D, *DSSI - 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*

## IP — General

RFC 768, *User Datagram Protocol*  
 RFC 793, *Transmission Control Protocol*  
 RFC 854, *TELNET Protocol Specifications*

RFC 951, *Bootstrap Protocol (BOOTP)*  
 RFC 1034, *Domain Names - Concepts and Facilities*  
 RFC 1035, *Domain Names - Implementation and Specification*  
 RFC 1350, *The TFTP Protocol (revision 2)*  
 RFC 1534, *Interoperation between DHCP and BOOTP*  
 RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*  
 RFC 2131, *Dynamic Host Configuration Protocol*  
 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 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 3046, *DHCP Relay Agent Information Option (Option 82)*  
 RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*  
 RFC 3596, *DNS Extensions to Support IP version 6*  
 RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*  
 RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*  
 RFC 4251, *The Secure Shell (SSH) Protocol Architecture*  
 RFC 4254, *The Secure Shell (SSH) Connection Protocol*  
 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*  
draft-grant-tacacs-02, *The TACACS+ Protocol*  
draft-ietf-vrrp-unified-spec-02, *Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6*

### IP — Multicast

RFC 1112, *Host Extensions for IP Multicasting*  
RFC 2236, *Internet Group Management Protocol, Version 2*  
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 Rendezvous 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 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 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*  
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 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*  
draft-dolganow-l3vpn-mvpn-expl-track-00, *Explicit tracking in MPLS/BGP IP VPNs*

### IP — Version 4

RFC 791, *Internet Protocol*  
RFC 792, *Internet Control Message Protocol*  
RFC 826, *An Ethernet Address Resolution Protocol*  
RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*  
RFC 1812, *Requirements for IPv4 Routers*  
RFC 2401, *Security Architecture for Internet Protocol*

RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

## 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 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*

RFC 3587, *IPv6 Global Unicast Address Format*

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 Only)*

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 5952, *A Recommendation for IPv6 Address Text Representation*

RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*

RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

## IPsec

RFC 2401, *Security Architecture for the Internet Protocol*

RFC 2406, *IP Encapsulating Security Payload (ESP)*

RFC 2409, *The Internet Key Exchange (IKE)*

RFC 2560, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

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 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*

RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*

draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*

## IS-IS

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*

## Standards and Protocol Support

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*  
draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*  
draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

## Management

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*  
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*  
RFC 1157, *A Simple Network Management Protocol (SNMP)*  
RFC 1215, *A Convention for Defining Traps for use with the SNMP*  
RFC 1724, *RIP Version 2 MIB Extension*  
RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIPv2*  
RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIPv2*  
RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*  
RFC 2206, *RSVP Management Information Base using SMIPv2*  
RFC 2213, *Integrated Services Management Information Base using SMIPv2*  
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 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 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 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 4292, *IP Forwarding Table MIB*
- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- 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 6241, *Network Configuration Protocol (NETCONF)*

- RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
- draft-ietf-snmv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- draft-ietf-idr-bgp4-mib-05, *Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4)*
- 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 SMIPv2*
- 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*

### MPLS — General

- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3032, *MPLS Label Stack Encoding*
- RFC 3443, *Time To Live (TTL) Processing in Multiprotocol 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*

### MPLS — GMPLS

- 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*
- draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

### MPLS — LDP

- 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 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
- RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
- draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*
- draft-ietf-mpls-ldp-ipv6-15, *Updates to LDP for IPv6*

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*  
 draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*  
 draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*  
 draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*

## MPLS — 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 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*

## MPLS — OAM

RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*  
 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*

## MPLS — RSVP-TE

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*

## Standards and Protocol Support

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

### NAT

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 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

### OpenFlow

ONF *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

### OSPF

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 4970, *Extensions to OSPF for Advertising Optional Router Capabilities*

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*

draft-ietf-ospf-prefix-link-attr-06, *OSPFv2 Prefix/Link Attribute Advertisement*

draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

### Policy Management and Credit Control

3GPP TS 29.212, *Policy and Charging Control (PCC) over Gx/Sd Reference Point (Release 11 and Release 12) Gx support as it applies to wireline environment (BNG)*

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

**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 2661, *Layer Two Tunneling Protocol "L2TP"*
- RFC 2686, *The Multi-Class Extension to Multi-Link PPP*
- RFC 2878, *PPP Bridging Control Protocol (BCP)*
- RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*
- RFC 5072, *IP Version 6 over PPP*

**Pseudowire**

- MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane 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*
- 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*
- draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*

### Quality of Service

- 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 3260, *New Terminology and Clarifications for Diffserv*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 3140, *Per Hop Behavior Identification Codes*

### RIP

- RFC 1058, *Routing Information Protocol*
- RFC 2080, *RIPng for IPv6*
- RFC 2082, *RIP-2 MD5 Authentication*
- RFC 2453, *RIP Version 2*

### SONET/SDH

- ITU-G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

### 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*
- 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*

### Voice and Video Performance

- 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)*

### 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)*



# Customer Documentation and Product Support



## Customer Documentation

<http://documentation.alcatel-lucent.com>



## Technical Support

<http://support.alcatel-lucent.com>



## Documentation Feedback

[documentation.feedback@alcatel-lucent.com](mailto:documentation.feedback@alcatel-lucent.com)

