



Alcatel-Lucent 7450

ETHERNET SERVICE SWITCH | RELEASE 13.0.R1

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 |
| Technical Support | 22 |
| VLL Services | |
| In This Chapter | 23 |
| ATM VLL (Apipe) Services | 24 |
| ATM VLL For End-to-End ATM Service | 24 |
| ATM Virtual Trunk Over IP/MPLS Packet-Switched Network | 25 |
| Circuit Emulation Services (Cpipe) | 27 |
| Mobile Infrastructure | 27 |
| Circuit Emulation Modes | 28 |
| Circuit Emulation Parameters | 31 |
| Circuit Emulation Modes | 31 |
| Absolute Mode Option | 31 |
| Payload Size | 31 |
| Jitter Buffer | 34 |
| CES Circuit Operation | 34 |
| Services for Transporting CES Circuits | 35 |
| Network Synchronization Considerations | 36 |
| Cpipe Payload | 37 |
| Ethernet Pipe (Epipe) Services | 38 |
| Epipe Service Overview | 39 |
| Epipe Service Pseudowire VLAN Tag Processing | 40 |
| Epipe Up Operational State Configuration Option | 44 |
| Epipe with PBB | 45 |
| Epipe over L2TPv3 | 46 |
| Traffic Management Support | 47 |
| Ingress SAP Classification and Marking | 47 |
| Egress Network EXP Marking | 47 |
| Ingress Network Classification | 47 |
| IP Interworking VLL (Ipipe) Services | 48 |
| Ipipe VLL | 48 |
| IP Interworking VLL Datapath | 50 |
| Extension to IP VLL for Discovery of Ethernet CE IP Address | 51 |
| VLL Ethernet SAP Procedures | 51 |
| IPv6 Support on IP Interworking VLL | 54 |
| IPv6 Datapath Operation | 54 |
| IPv6 Stack Capability Signaling | 56 |
| Services Configuration for MPLS-TP | 58 |
| MPLS-TP SDPs | 58 |
| VLL Spoke SDP Configuration | 60 |
| Epipe VLL Spoke-SDP Termination on IES, VPRN and VPLS | 63 |

Table of Contents

| | |
|--|-----|
| Configuring MPLS-TP Lock Instruct and Loopback | 64 |
| MPLS-TP PW Lock Instruct and Loopback Overview | 64 |
| Lock PW End-Point Model | 65 |
| PW Redundancy and Lock Instruct and Loopback | 65 |
| Configuring a Test SAP for an MPLS-TP PW | 66 |
| Configuring an Administrative Lock | 67 |
| Configuring a Loopback | 68 |
| VCCV BFD support for VLL, Spoke-SDP Termination on IES and VPRN, and VPLS Services | 69 |
| VCCV BVD Support | 69 |
| VCCV BFD Encapsulation on a Pseudowire | 70 |
| BFD Session Operation | 70 |
| Configuring VCCV BFD | 71 |
| Pseudowire Switching | 73 |
| Pseudowire Switching with Protection | 74 |
| Pseudowire Switching Behavior | 76 |
| Static-to-Dynamic Pseudowire Switching | 78 |
| Ingress VLAN Swapping | 79 |
| Ingress VLAN Translation | 80 |
| Pseudowire Redundancy | 81 |
| Dynamic Multi-Segment Pseudowire Routing | 82 |
| Overview | 82 |
| Pseudowire Routing | 87 |
| Configuring VLLs using Dynamic MS-PWs | 90 |
| Pseudowire Redundancy | 93 |
| VCCV OAM for Dynamic MS-PWs | 95 |
| VCCV-Ping on Dynamic MS-PWs | 95 |
| VCCV-Trace on Dynamic MS-PWs | 96 |
| Example Dynamic MS-PW Configuration | 97 |
| VLL Resilience with Two Destination PE Nodes | 100 |
| Master-Slave Operation | 102 |
| Epipe Using BGP-MH Site Support for Ethernet Tunnels | 110 |
| Operational Overview | 112 |
| Detailed Operation | 113 |
| BGP-MH Site Support for Ethernet Tunnels Operational-Group Model | 117 |
| BGP-MH Specifics for MH Site Support for Ethernet Tunnels | 117 |
| PW Redundancy for BGP MH Site Support for Ethernet Tunnels | 117 |
| T-LDP Status Notification Handling Rules of BGP-MH Epipes | 118 |
| Access Node Resilience Using MC-LAG and Pseudowire Redundancy | 129 |
| VLL Resilience for a Switched Pseudowire Path | 131 |
| Pseudowire Redundancy Service Models | 133 |
| Redundant VLL Service Model | 133 |
| T-LDP Status Notification Handling Rules | 135 |
| Processing Endpoint SAP Active/Standby Status Bits | 135 |
| Processing and Merging | 135 |
| High-Speed Downlink Packet Access (HSDPA) Off Load Fallback over ATM | 137 |
| Primary Spoke-SDP Fallback to Secondary SAP | 138 |
| Reversion to Primary Spoke SDP Path | 138 |
| MC-APS and MC-LAG | 139 |
| Failure Scenarios | 141 |

| | |
|---|-----|
| VLL Using G.8031 Protected Ethernet Tunnels | 142 |
| BGP Virtual Private Wire Service (VPWS) | 143 |
| Single-Homed BGP VPWS | 143 |
| Dual-Homed BGP VPWS | 144 |
| VLL Service Considerations | 154 |
| SDPs | 154 |
| SDP Statistics for VPLS and VLL Services | 155 |
| SAP Encapsulations and Pseudowire Types | 156 |
| QoS Policies | 157 |
| Filter Policies | 157 |
| MAC Resources | 157 |
| Configuring a VLL Service with CLI | 159 |
| Basic Configurations | 160 |
| Common Configuration Tasks | 160 |
| Configuring VLL Components | 161 |
| Creating an Epipe Service | 162 |
| Creating an Ipipe Service | 173 |
| Using Spoke SDP Control Words | 177 |
| Same Fate Epipe VLANs Access Protection | 178 |
| Pseudowire Configuration Notes | 180 |
| Configuring Two VLL Paths Terminating on T-PE2 | 182 |
| Configuring VLL Resilience | 185 |
| Configuring VLL Resilience for a Switched Pseudowire Path | 186 |
| Configuring BGP Virtual Private Wire Service (VPWS) | 188 |
| Single-Homed BGP VPWS | 188 |
| Dual-Homed BGP VPWS | 190 |
| Service Management Tasks | 196 |
| Modifying Epipe Service Parameters | 197 |
| Disabling an Epipe Service | 197 |
| Re-Enabling an Epipe Service | 198 |
| Deleting an Epipe Service | 198 |
| Modifying Ipipe Service Parameters | 199 |
| Disabling an Ipipe Service | 200 |
| Re-enabling an Ipipe Service | 201 |
| Deleting an Ipipe Service | 201 |
| VLL Services Command Reference | 203 |
| Command Hierarchies | 203 |
| VLL Service Configuration Commands | 223 |
| VLL Show Commands | 345 |

Virtual Private LAN Service

| | |
|---|-----|
| In This Chapter | 409 |
| VPLS Service Overview | 411 |
| VPLS Packet Walkthrough | 412 |
| VPLS Features | 415 |
| VPLS Enhancements | 415 |
| VPLS over MPLS | 416 |
| VPLS Service Pseudowire VLAN Tag Processing | 417 |
| VPLS MAC Learning and Packet Forwarding | 421 |
| MAC Learning Protection | 421 |

Table of Contents

| | |
|--|-----|
| DEI in IEEE 802.1ad | 423 |
| VPLS Using G.8031 Protected Ethernet Tunnels | 424 |
| Pseudowire Control Word | 425 |
| Table Management | 426 |
| FIB Size | 426 |
| FIB Size Alarms | 426 |
| Local and Remote Aging Timers | 427 |
| Disable MAC Aging | 427 |
| Disable MAC Learning | 427 |
| Unknown MAC Discard | 427 |
| VPLS and Rate Limiting | 428 |
| MAC Move | 428 |
| Auto-Learn MAC Protect | 429 |
| Split Horizon SAP Groups and Split Horizon Spoke SDP Groups | 433 |
| VPLS and Spanning Tree Protocol | 434 |
| Spanning Tree Operating Modes | 434 |
| Multiple Spanning Tree | 436 |
| MSTP for QinQ SAPs | 438 |
| Provider MSTP | 438 |
| Enhancements to the Spanning Tree Protocol | 440 |
| Egress Multicast Groups | 443 |
| Egress Multicast Group Provisioning | 443 |
| VPLS Redundancy | 454 |
| Spoke SDP Redundancy for Metro Interconnection | 454 |
| Spoke SDP Based Redundant Access | 456 |
| Inter-Domain VPLS Resiliency Using Multi-Chassis Endpoints | 457 |
| Support for Single Chassis Endpoint Mechanisms | 461 |
| Using B-VPLS for Increased Scalability and Reduced Convergence Times | 465 |
| MAC Flush Additions for PBB VPLS | 467 |
| VPLS Access Redundancy | 470 |
| STP-Based Redundant Access to VPLS | 470 |
| Redundant Access to VPLS Without STP | 471 |
| Object Grouping and State Monitoring | 472 |
| VPLS Applicability — Block on VPLS a Failure | 472 |
| MAC Flush Message Processing | 474 |
| Dual Homing to a VPLS Service | 476 |
| ACL Next-Hop for VPLS | 478 |
| SDP Statistics for VPLS and VLL Services | 479 |
| BGP Auto-Discovery for LDP VPLS | 480 |
| BGP AD Overview | 480 |
| Information Model | 480 |
| FEC Element for T-LDP Signaling | 482 |
| BGP-AD and Target LDP (T-LDP) Interaction | 483 |
| SDP Usage | 485 |
| Automatic Creation of SDPs | 485 |
| Manually Provisioned SDP | 485 |
| Automatic Instantiation of Pseudowires (SDP Bindings) | 486 |
| Mixing Statically Configured and Auto-Discovered Pseudowires in a VPLS | 487 |
| Resiliency Schemes | 487 |
| BGP VPLS | 488 |

| | |
|--|-----|
| Pseudowire Signaling Details | 489 |
| Supported VPLS Features | 491 |
| VCCV BFD Support for VPLS Services | 493 |
| BGP Multi-Homing for VPLS | 494 |
| Information Model and Required Extensions to L2VPN NLRI | 495 |
| Supported Services and Multi-Homing Objects | 496 |
| Blackhole Avoidance | 497 |
| BGP Multi-Homing for VPLS Inter-Domain Resiliency | 498 |
| Multicast-Aware VPLS | 499 |
| PIM Snooping for VPLS | 499 |
| Multicast Listener Discovery (MLD) Snooping and MAC-Based Multicast Forwarding | 501 |
| PIM and IGMP Snooping Interaction | 502 |
| VPLS Multicast-Aware High Availability Features | 502 |
| RSVP and LDP P2MP LSP for Forwarding VPLS/B-VPLS BUM and IP Multicast Packets | 503 |
| Routed VPLS and I-VPLS | 505 |
| IES or VPRN IP Interface Binding | 505 |
| Assigning a Service Name to a VPLS Service | 505 |
| Service Binding Requirements | 506 |
| Bound Service Name Assignment | 506 |
| Binding a Service Name to an IP Interface | 506 |
| Bound Service Deletion or Service Name Removal | 507 |
| IP Interface Attached VPLS Service Constraints | 507 |
| IP Interface and VPLS Operational State Coordination | 507 |
| IP Interface MTU and Fragmentation | 508 |
| Unicast IP Routing into a VPLS Service | 508 |
| ARP and VPLS FIB Interactions | 509 |
| Routed VPLS Specific ARP Cache Behavior | 510 |
| The allow-ip-int-binding VPLS Flag | 511 |
| Routed VPLS SAPs Only Supported on Standard Ethernet Ports | 511 |
| Routed VPLS SAPs Only Supported on FP2 (or later) Based Systems or IOM/IMM | 511 |
| Network Ports Restricted to FP2-Based Systems or IOMs | 511 |
| LAG Port Membership Constraints | 512 |
| Routed VPLS Feature Restrictions | 513 |
| Routed I-VPLS Feature Restrictions | 513 |
| IES IP Interface VPLS Binding and Chassis Mode Interaction | 514 |
| VPRN IP Interface VPLS Binding and Forwarding Plane Constraints | 514 |
| Route Leaking Between Routing Contexts | 514 |
| Ingress LAG and FP1 to Routed VPLS Discards | 515 |
| IPv4 Multicast Routing Support | 516 |
| BGP Auto Discovery (BGP-AD) for Routed VPLS Support | 517 |
| Routed VPLS Caveats | 518 |
| VPLS SAP Ingress IP Filter Override | 518 |
| IP Interface Defined Egress QoS Reclassification | 518 |
| Remarking for VPLS and Routed Packets | 519 |
| 7450 Mixed Mode Chassis | 519 |
| IPv4 Multicast Routing | 519 |
| Routed VPLS Supported Routing Related Protocols | 519 |
| Spanning Tree and Split Horizon | 520 |
| VPLS Service Considerations | 521 |
| SAP Encapsulations | 521 |

Table of Contents

| | |
|--|-----|
| VLAN Processing | 521 |
| Ingress VLAN Swapping | 522 |
| Service Auto-Discovery using Multiple VLAN Registration Protocol (MVRP) | 523 |
| Configure the MVRP Infrastructure using an M-VPLS Context | 524 |
| Instantiate Related VLAN FIBs and Trunks in MVRP Scope | 524 |
| MVRP Activation of Service Connectivity | 527 |
| MVRP Control Plane | 530 |
| STP-MVRP Interaction | 530 |
| VPLS E-Tree Services | 533 |
| VPLS E-Tree Services Overview | 533 |
| Leaf-ac and Root-ac SAPs | 534 |
| Leaf-ac and Root-ac SDP Binds | 535 |
| Root-leaf-tag SAPs | 535 |
| Root-leaf-tag SDP Binds | 536 |
| Interaction between VPLS E-Tree Services and Other Features | 537 |
| Configuring a VPLS Service with CLI | 539 |
| Basic Configuration | 540 |
| Common Configuration Tasks | 542 |
| Configuring VPLS Components | 543 |
| Configuring Egress Multicast Groups | 544 |
| Creating a VPLS Service | 545 |
| Enabling Multiple MAC Registration Protocol (MMRP) | 546 |
| Configuring GSMP Parameters | 555 |
| Configuring a VPLS SAP | 556 |
| Applying an Egress Multicast Group to a VPLS Service SAP | 567 |
| Configuring SAP Subscriber Management Parameters | 568 |
| MSTP Control over Ethernet Tunnels | 569 |
| Configuring SDP Bindings | 570 |
| Configuring Overrides on Service SAPs | 571 |
| Configuring VPLS Redundancy | 583 |
| Creating a Management VPLS for SAP Protection | 583 |
| Creating a Management VPLS for Spoke SDP Protection | 586 |
| Configuring Load Balancing with Management VPLS | 589 |
| Configuring Selective MAC Flush | 594 |
| Configuring Multi-Chassis Endpoints | 595 |
| Configuring BGP Auto-Discovery | 599 |
| Configuration Steps | 599 |
| LDP Signaling | 602 |
| Pseudowire Template | 604 |
| Configuring BGP VPLS | 606 |
| Configuring a VPLS Management Interface | 608 |
| Configuring Policy-Based Forwarding for Deep Packet Inspection (DPI) in VPLS | 609 |
| Configuring VPLS E-Tree Services | 612 |
| Service Management Tasks | 613 |
| Modifying VPLS Service Parameters | 613 |
| Modifying Management VPLS Parameters | 614 |
| Deleting a Management VPLS | 614 |
| Disabling a Management VPLS | 615 |
| Deleting a VPLS Service | 616 |
| Disabling a VPLS Service | 616 |

| | |
|---|-----|
| Re-Enabling a VPLS Service | 617 |
| VPLS Services Command Reference | 619 |
| Command Hierarchies | 619 |
| VPLS Service Configuration Commands | 659 |
| VPLS Show Commands | 881 |

IEEE 802.1ah Provider Backbone Bridging

| | |
|---|------|
| In This Chapter | 1059 |
| IEEE 802.1ah Provider Backbone Bridging (PBB) Overview | 1060 |
| PBB Features | 1061 |
| Integrated PBB-VPLS Solution | 1061 |
| PBB Technology | 1063 |
| PBB Mapping to Existing VPLS Configurations | 1064 |
| SAP and SDP Support | 1066 |
| PBB B-VPLS | 1066 |
| PBB I-VPLS | 1066 |
| PBB Packet Walkthrough | 1068 |
| PBB Control Planes | 1069 |
| Shortest Path Bridging MAC Mode (SPBM) | 1070 |
| Flooding and Learning Versus Link State | 1070 |
| SPB for B-VPLS | 1071 |
| Control B-VPLS and User B-VPLS | 1071 |
| Shortest Path and Single Tree | 1074 |
| Data Path and Forwarding | 1077 |
| SPB Ethernet OAM | 1077 |
| SPB Levels | 1078 |
| SPBM to Non-SPBM Interworking | 1079 |
| Static MACs and Static ISIDs | 1079 |
| Epipe Static Configuration | 1079 |
| SPBM ISID Policies | 1081 |
| ISID Policy Control | 1083 |
| Static ISID Advertisement | 1083 |
| I-VPLS for Unicast Service | 1083 |
| Default Behaviors | 1084 |
| Example Network Configuration | 1085 |
| Sample Configuration for Dut-A | 1086 |
| IEEE 802.1ak MMRP for Service Aggregation and Zero Touch Provisioning | 1093 |
| MMRP Support Over B-VPLS SAPs and SDPs | 1095 |
| I-VPLS Changes and Related MMRP Behavior | 1095 |
| Limiting the Number of MMRP Entries on a Per B-VPLS Basis | 1095 |
| Optimization for Improved Convergence Time | 1096 |
| Controlling MRP Scope using MRP Policies | 1096 |
| PBB and BGP-AD | 1100 |
| PBB ELINE Service | 1100 |
| Non-Redundant PBB Epipe Spoke Termination | 1100 |
| PBB Using G.8031-Protected Ethernet Tunnels | 1101 |
| Solution Overview | 1101 |
| Detailed Solution Description | 1103 |
| Detailed PBB Emulated LAG Solution Description | 1106 |
| Support Service and Solution Combinations | 1108 |

Table of Contents

| | |
|---|------|
| Periodic MAC Notification | 1109 |
| MAC Flush | 1110 |
| PBB Resiliency for B-VPLS Over Pseudowire Infrastructure | 1110 |
| Access Multi-Homing for Native PBB (B-VPLS over SAP Infrastructure) | 1115 |
| Solution Description for I-VPLS Over Native PBB Core | 1116 |
| Solution Description for PBB Epipe over G.8031 Ethernet Tunnels | 1120 |
| BGP Multi-homing for I-VPLS | 1124 |
| Access Multi-Homing over MPLS for PBB Epipes | 1125 |
| PBB and IGMP/MLD Snooping | 1128 |
| PBB QoS | 1129 |
| Transparency of Customer QoS Indication through PBB Backbone | 1130 |
| Egress B-SAP per ISID Shaping | 1136 |
| B-SAP Egress ISID Shaping Configuration | 1136 |
| Provisioning Model | 1138 |
| Egress Queue Scheduling | 1140 |
| B-SAP per-ISID Shaping Configuration Example | 1142 |
| PBB OAM | 1145 |
| Mirroring | 1146 |
| OAM Commands | 1146 |
| CFM Support | 1146 |
| Configuration Examples | 1147 |
| PBB using G.8031 Protected Ethernet Tunnels | 1147 |
| MC-LAG Multihoming for Native PBB | 1150 |
| Access Multi-Homing over MPLS for PBB Epipes | 1152 |
| PBB Command Reference | 1155 |
| Command Hierarchies | 1155 |
| PBB Service Commands | 1163 |
| PBB Show Commands | 1193 |
| PBB Clear Commands | 1210 |
| PBB Debug Commands | 1212 |

Ethernet Virtual Private Networks (EVPN)

| | |
|---|------|
| In This Chapter | 1215 |
| Overview | 1216 |
| EVPN for VXLAN Tunnels in a Layer-2 DC GW | 1217 |
| EVPN for VXLAN Tunnels in a Layer-2 DC | 1219 |
| EVPN for VXLAN Tunnels in a Layer 3 DC | 1220 |
| EVPN for VXLAN Tunnels in a Layer 3 DC | 1222 |
| VXLAN | 1224 |
| VXLAN ECMP and LAG | 1227 |
| VXLAN VPLS Tag Handling | 1227 |
| VXLAN MTU Considerations | 1227 |
| VXLAN QoS | 1228 |
| VXLAN Ping | 1228 |
| IGMP-snooping on VXLAN | 1232 |
| BGP-EVPN Control Plane for VXLAN Overlay Tunnels | 1235 |
| EVPN for VXLAN in VPLS Services | 1239 |
| Resiliency and BGP Multi-Homing | 1241 |
| Use of bgp-evpn, bgp-ad, and Sites in the Same VPLS Service | 1241 |
| Use of the unknown-mac-route | 1243 |

| | |
|---|-------------|
| ARP/ND snooping and proxy support | 1244 |
| Proxy-ARP/ND periodic refresh, unsolicited refresh and confirm-messages | 1247 |
| Proxy-ND and the Router Flag in Neighbor Advertisement messages | 1248 |
| BGP-EVPN MAC-Mobility | 1248 |
| Conditional Static MAC and Protection | 1250 |
| EVPN for VXLAN in R-VPLS Services | 1252 |
| EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes | 1254 |
| EVPN for VXLAN in EVPN Tunnel R-VPLS Services | 1258 |
| Interaction of EVPN and VXLAN with Existing VPLS Features | 1263 |
| Interaction of EVPN and VXLAN with Existing VPRN Features | 1264 |
| Routing Policies for BGP EVPN IP Prefixes | 1265 |
| DC GW integration with the Nuage Virtual Services Directory (VSD) | 1268 |
| XMPP Interface on the DC GW | 1269 |
| Overview of the Static-Dynamic VSD Integration Model | 1273 |
| VSD-domains and Association to Static-Dynamic Services | 1275 |
| VSD-domain Type L2-DOMAIN | 1275 |
| VSD-domain Type L2-DOMAIN-IRB | 1278 |
| VSD-domain Type VRF-GRE | 1278 |
| VSD-domain Type VRF-VXLAN | 1278 |
| Configuring a EVPN Service with CLI | 1281 |
| EVPN Configuration Examples | 1282 |
| Layer 2 PE Example | 1282 |
| EVPN for VXLAN in R-VPLS Services Example | 1284 |
| EVPN for VXLAN in EVPN Tunnel R-VPLS Services Example | 1286 |
| EVPN for VXLAN in R-VPLS Services with IPv6 interfaces and prefixes Example | 1287 |
| EVPN Command Reference | 1289 |
| Command Hierarchies | 1289 |
| Show Commands | 1309 |
| Clear Commands | 1314 |
| Tools Commands | 1315 |
| Debug Commands | 1318 |
| Common CLI Command Descriptions | |
| In This Chapter | 1319 |
| Common Service Commands | 1320 |
| Standards and Protocol Support | 1323 |

Table of Contents

List of Tables

VLL Services

| | |
|--|-----|
| Table 1: Mobile Infrastructure Definitions | 27 |
| Table 2: Unstructured Payload Defaults | 32 |
| Table 3: Structured Number of Frames Defaults | 32 |
| Table 4: Epipe Spoke SDP VLAN Tag Processing: Ingress | 41 |
| Table 5: Epipe Spoke SDP VLAN Tag Processing: Egress | 42 |
| Table 6: Mapping of Real Services to Test Service Types | 68 |
| Table 7: SAP MEP Signaling | 115 |
| Table 8: Supported SAP Types | 164 |
| Table 9: Default QinQ and TopQ SAP Dot1P Evaluation | 307 |
| Table 10: Bottom Position QinQ and TopQ SAP Dot1P Evaluation | 309 |
| Table 11: Show Service Egress Label Output Fields | 345 |

Virtual Private LAN Service

| | |
|--|-----|
| Table 12: VPLS Mesh and Spoke SDP VLAN Tag Processing: Ingress | 418 |
| Table 13: VPLS Mesh and Spoke SDP VLAN Tag Processing: Egress | 419 |
| Table 14: SAP Chain Creation | 451 |
| Table 15: Ingress Routed to VPLS Next-Hop Behavior | 510 |
| Table 16: Egress Routed VPLS Next-Hop Behavior | 510 |
| Table 17: MSTP and MVRP Interaction Table | 530 |
| Table 18: Spoke SDP BPDU Encapsulation States | 580 |
| Table 19: Default QinQ and TopQ SAP Dot1P Evaluation | 799 |
| Table 20: Bottom Position QinQ and TopQ SAP Dot1P Evaluation | 800 |

IEEE 802.1ah Provider Backbone Bridging

| | |
|--|------|
| Table 21: B-VPLS Control Planes | 1070 |
| Table 22: SPB Ethernet OAM Operation Summary | 1078 |
| Table 23: SPBM ISID Policies Table | 1082 |

Ethernet Virtual Private Networks (EVPN)

List of Tables

List of Figures

VLL Services

| | |
|---|-----|
| Figure 1: ATM VLL for End-to-End ATM Service | 25 |
| Figure 2: VT Application Example | 25 |
| Figure 3: Mobile Infrastructure | 27 |
| Figure 4: RFC 4553 (SAToP) MPLS PSN Encapsulation | 29 |
| Figure 5: CESoPSN Packet Format for an MPLS PSN | 29 |
| Figure 6: MEF8 PSN Encapsulation | 30 |
| Figure 7: CESoPSN MPLS Encapsulation | 37 |
| Figure 8: Epipe/VLL Service | 39 |
| Figure 9: L2TPv3 SDP Illustration | 46 |
| Figure 10: IP Interworking VLL (Ipipe) | 49 |
| Figure 11: IP Interworking VLL Datapath | 50 |
| Figure 12: Data Path for Ethernet CE to PPP Attached CE | 55 |
| Figure 13: Pseudowire Service Switching Node | 73 |
| Figure 14: VLL Resilience with Pseudowire Redundancy and Switching | 74 |
| Figure 15: Ingress VLAN Swapping | 79 |
| Figure 16: Ingress VLAN Translation | 80 |
| Figure 17: Dynamic MS-PW Overview | 82 |
| Figure 18: MS-PW Addressing using FEC129 All Type 2 | 83 |
| Figure 19: Advertisement of PE Addresses by PW Routing | 84 |
| Figure 20: Signaling of Dynamic MS-PWs using T-LDP | 85 |
| Figure 21: Mapping of All to SAP | 85 |
| Figure 22: VLL Using Dynamic MS-PWs, Inter-AS Scenario | 86 |
| Figure 23: Pseudowire Redundancy | 93 |
| Figure 24: Dynamic MS-PW Example | 97 |
| Figure 25: VLL Resilience | 100 |
| Figure 26: Master-Slave Pseudowire Redundancy | 103 |
| Figure 27: Example of SAP OAM Interaction with Master-Slave Pseudowire Redundancy | 105 |
| Figure 28: VLL Resilience | 107 |
| Figure 29: VLL Resilience with Pseudowire Switching | 109 |
| Figure 30: BGP-MH Site Support for Ethernet Tunnels | 110 |
| Figure 31: G.8031 for Slave Operation | 112 |
| Figure 32: Full Redundancy G.8031 Epipe & BGP-MH | 114 |
| Figure 33: Sample Topology Full Redundancy | 120 |
| Figure 34: Access Node Resilience | 129 |
| Figure 35: VLL Resilience with Pseudowire Redundancy and Switching | 131 |
| Figure 36: Redundant VLL Endpoint Objects | 133 |
| Figure 37: HSDPA Off Load Fallback over ATM | 137 |
| Figure 38: HSDPA Off Load Fallback with MC-APS | 139 |
| Figure 39: Ethernet Failure At Cell Site | 141 |
| Figure 40: Single-Homed BGP-VPWS Example | 143 |
| Figure 41: Dual-Homed BGP VPWS with Single Pseudowire | 144 |
| Figure 42: Dual-homed BGP VPWS with Active/Standby Pseudowires | 145 |
| Figure 43: BGP VPWS Update Extended Community Format | 146 |
| Figure 44: BGP VPWS NLRI | 148 |

List of Figures

| | |
|---|-----|
| Figure 45: BGP VPWS NLRI TLV Extension Format | 148 |
| Figure 46: Circuit Status Vector TLV Type | 148 |
| Figure 47: SDP Statistics for VPLS and VLL Services | 155 |
| Figure 48: SDPs — Uni-Directional Tunnels | 170 |
| Figure 49: VLL Resilience with Pseudowire Redundancy and Switching | 182 |
| Figure 50: VLL Resilience | 185 |
| Figure 51: VLL Resilience with Pseudowire Switching | 186 |
| Figure 52: Single-Homed BGP VPWS Configuration Example | 188 |
| Figure 53: Example of Dual-Homed BGP VPWS with Single Pseudowire | 190 |
| Figure 54: Example of Dual-homed BGP VPWS with Active/Standby Pseudowires | 193 |

Virtual Private LAN Service

| | |
|--|-----|
| Figure 55: VPLS Service Architecture | 412 |
| Figure 56: Access Port Ingress Packet Format and Lookup | 412 |
| Figure 57: Network Port Egress Packet Format and Flooding | 413 |
| Figure 58: Access Port Egress Packet Format and Lookup | 414 |
| Figure 59: MAC Learning Protection | 422 |
| Figure 60: DE Bit in the 802.1ad S-TAG | 423 |
| Figure 61: Access Resiliency | 437 |
| Figure 62: HVPLS with Spoke Redundancy | 455 |
| Figure 63: HVPLS Resiliency Based on AS Pseudowires | 457 |
| Figure 64: Multi-Chassis Pseudowire Endpoint for VPLS | 458 |
| Figure 65: MC-EP in Passive Mode | 461 |
| Figure 66: MAC Flush in the MC-EP Solution | 463 |
| Figure 67: MC-EP with B-VPLS | 466 |
| Figure 68: MC-EP with B-VPLS Failure Scenario | 467 |
| Figure 69: MC-EP with B-VPLS Mac Flush Solution | 468 |
| Figure 70: Dual Homed MTU-s in Two-Tier Hierarchy H-VPLS | 470 |
| Figure 71: Dual Homed CE Connection to VPLS | 476 |
| Figure 72: Application 1 Diagram | 478 |
| Figure 73: SDP Statistics for VPLS and VLL Services | 479 |
| Figure 74: BGP AD NLRI versus IP VPN NLRI | 481 |
| Figure 75: Generalized Pseudowire-ID FEC Element | 482 |
| Figure 76: BGP-AD and T-LDP Interaction | 484 |
| Figure 77: BGP VPLS Solution | 488 |
| Figure 78: BGP Multi-Homing for VPLS | 494 |
| Figure 79: BGP MH-NLRI for VPLS Multi-Homing | 495 |
| Figure 80: BGP MH Used in an HVPLS Topology | 498 |
| Figure 81: IPv4 Multicast with a Router VPLS service | 516 |
| Figure 82: Ingress VLAN Swapping | 522 |
| Figure 83: Infrastructure for MVRP Exchanges | 523 |
| Figure 84: Service Instantiation with MVRP - QinQ to PBB Example | 527 |
| Figure 85: E-Tree Service | 534 |
| Figure 86: Mapping PE Model to 7x50 VPLS Service | 535 |
| Figure 87: Leaf and Root Tagging Dot1q | 536 |
| Figure 88: Leaf and Root Tagging PW | 537 |
| Figure 89: SDPs — Uni-Directional Tunnels | 572 |
| Figure 90: Example Configuration for Protected VPLS SAP | 584 |
| Figure 91: Example Configuration for Protected VPLS Spoke SDP | 587 |

| | |
|---|-----|
| Figure 92: Example Configuration for Load Balancing Across Two Protected VPLS Spoke SDPs. | 589 |
| Figure 93: BGP AD Configuration Example. | 599 |
| Figure 94: BGP-AD CLI Command Tree | 600 |
| Figure 95: BGP AD Triggering LDP Functions | 602 |
| Figure 96: Show Router LDP Session Output | 603 |
| Figure 97: Show Router LDP Bindings FEC-Type Services | 603 |
| Figure 98: PW-Template CLI Tree. | 604 |
| Figure 99: PW-Template-Binding CLI Syntax | 605 |
| Figure 100: BGP VPLS Example. | 606 |
| Figure 101: Policy-Based Forwarding For Deep Packet Inspection | 609 |

IEEE 802.1ah Provider Backbone Bridging

| | |
|---|------|
| Figure 102: Large HVPLS Deployment | 1061 |
| Figure 103: Large PBB-VPLS Deployment | 1062 |
| Figure 104: QinQ Payload in Provider Header Example | 1063 |
| Figure 105: PBB Mapping to VPLS Constructs | 1064 |
| Figure 106: PBB Packet Walkthrough | 1068 |
| Figure 107: Control and User B-VPLS with FIDs. | 1072 |
| Figure 108: Sample Partial Mesh network. | 1074 |
| Figure 109: Unicast Paths for Low-path-id and High-path-id | 1075 |
| Figure 110: Multicast Paths for Low-path-id and High-path-id. | 1076 |
| Figure 111: Static MACs Example. | 1080 |
| Figure 112: Static ISIDs Example | 1081 |
| Figure 113: ISID Policy Example | 1084 |
| Figure 114: Sample Network | 1085 |
| Figure 115: Customer Services Transported in 1 B-VPLS (M:1 Model) | 1093 |
| Figure 116: Flood Containment Requirement in M:1 Model | 1094 |
| Figure 117: Inter-Domain Topology | 1097 |
| Figure 118: Limiting the Scope of MMRP Advertisements | 1097 |
| Figure 119: Mobile Backhaul Use Case. | 1102 |
| Figure 120: PBB-Epipe with B-VPLS over Ethernet Tunnel | 1103 |
| Figure 121: G.8031 P2P Tunnels and LAG-Like Loadsharing Co-Existence | 1104 |
| Figure 122: Ethernet Tunnel Overlay. | 1106 |
| Figure 123: TCN Triggered PBB Flush-ALI-But-Mine Procedure | 1113 |
| Figure 124: Access Dual-Homing into PBB BEBs - Topology View | 1115 |
| Figure 125: PBB Active Topology and Access Multi-Homing | 1116 |
| Figure 126: Access Multi-Homing - Link Failure | 1118 |
| Figure 127: Access Multi-Homing Solution for PBB Epipe | 1120 |
| Figure 128: Access Dual-Homing for PBB ELINE - BEB Failure. | 1121 |
| Figure 129: Solution for Access Dual-Homing with Local Switching for PBB Eline/Epipe. | 1122 |
| Figure 130: Active/Standby PW into PBB Epipes | 1125 |
| Figure 131: PCP, DE Bits Transparency in PBB | 1130 |
| Figure 132: Egress Queue Scheduling | 1140 |
| Figure 133: PBB OAM View for MPLS Infrastructure | 1145 |

Ethernet Virtual Private Networks (EVPN)

| | |
|---|------|
| Figure 134: Layer-2 DC PE with VPLS to the WAN. | 1217 |
| Figure 135: GW IRB on the DC PE for an L2 EVPN/VXLAN DC. | 1219 |
| Figure 136: GW IRB on the DC PE for an L3 EVPN/VXLAN DC. | 1220 |
| Figure 137: EVPN-Tunnel GW IRB on the DC PE for an L3 EVPN/VXLAN DC | 1222 |

List of Figures

| | |
|--|------|
| Figure 138: VXLAN Frame Format | 1225 |
| Figure 139: EVPN-VXLAN Required Routes and Communities | 1235 |
| Figure 140: EVPN Route-Type 5 | 1238 |
| Figure 141: Proxy-ARP example usage in an EVPN Network | 1244 |
| Figure 142: IP-VPN Import and EVPN Export BGP Workflow | 1265 |
| Figure 143: EVPN Import and IP-VPN Export BGP Workflow | 1266 |
| Figure 144: Basic XMPP Architecture | 1270 |
| Figure 145: WAN Services Attachment Procedure | 1273 |

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 7450 ESS 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 7450 ESS documentation set is composed of the following guides:

- **7450 ESS Basic System Configuration Guide**
This guide describes basic system configurations and operations.
- **7450 ESS System Management Guide**
This guide describes system security and access configurations as well as event logging and accounting logs.
- **7450 ESS Interface Configuration Guide**
This guide describes card, Media Dependent Adapter (MDA) and port provisioning.
- **7450 ESS 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.
- **7450 ESS Routing Protocols Guide**
This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.
- **7450 ESS MPLS Guide MPLS Guide**
This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- **7450 ESS Services Overview Guide**
This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
- **7450 ESS Layer 2 Services and EVPN Guide**
This guide describes Virtual Leased Lines (VLL), Virtual Private LAN Service (VPLS), Provider Backbone Bridging (PBB), and Ethernet VPN (EVPN).
- **7450 ESS Layer 3 Services Guide**
This guide describes Internet Enhanced Services (IES) and Virtual Private Routed Network (VPRN) services.
- **7450 ESS Versatile Service Module Guide**
This guide describes how to configure service parameters for the Versatile Service Module (VSM).
- **7450 ESS OAM and Diagnostics Guide**
This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- **7450 ESS Triple Play Guide**
This guide describes Triple Play services and support provided by the 7450 ESS and presents examples to configure and implement various protocols and services.

- **7450 ESS Quality of Service Guide**
This guide describes how to configure Quality of Service (QoS) policy management.
- **Multi-Service Integrated Service Adapter Guide**
This guide describes services provided by integrated service adapters such as Application Assurance, ad insertion (ADI) and Network Address Translation (NAT).

Technical Support

If you purchased a service agreement for your 7450 ESS 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:

<http://support.alcatel-lucent.com>

VLL Services

In This Chapter

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

Topics in this chapter include:

- [ATM VLL \(Apipe\) Services on page 24](#)
- [Circuit Emulation Services \(Cpipe\) on page 27](#)
- [Ethernet Pipe \(Epipe\) Services on page 38](#)
- [IP Interworking VLL \(Ipipe\) Services on page 48](#)
- [Services Configuration for MPLS-TP on page 58](#)
- [VCCV BFD support for VLL, Spoke-SDP Termination on IES and VPRN, and VPLS Services on page 69](#)
- [Pseudowire Switching on page 73](#)
- [Pseudowire Redundancy on page 81](#)
- [Dynamic Multi-Segment Pseudowire Routing on page 82](#)
- [Epipe Using BGP-MH Site Support for Ethernet Tunnels on page 110](#)
- [VLL Using G.8031 Protected Ethernet Tunnels on page 142](#)
- [BGP Virtual Private Wire Service \(VPWS\) on page 143](#)

ATM VLL (Apipe) Services

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

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

Topics in this section include:

- [ATM VLL For End-to-End ATM Service on page 24](#)
 - [ATM Virtual Trunk Over IP/MPLS Packet-Switched Network on page 25](#)
 - [Common Configuration Tasks on page 160](#)
 - [Configuring VLL Components on page 161](#)
 - [Service Management Tasks on page 196](#)
-

ATM VLL For End-to-End ATM Service

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

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

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

[Figure 1](#) shows an example of ATM VLL for end-to-end ATM service.

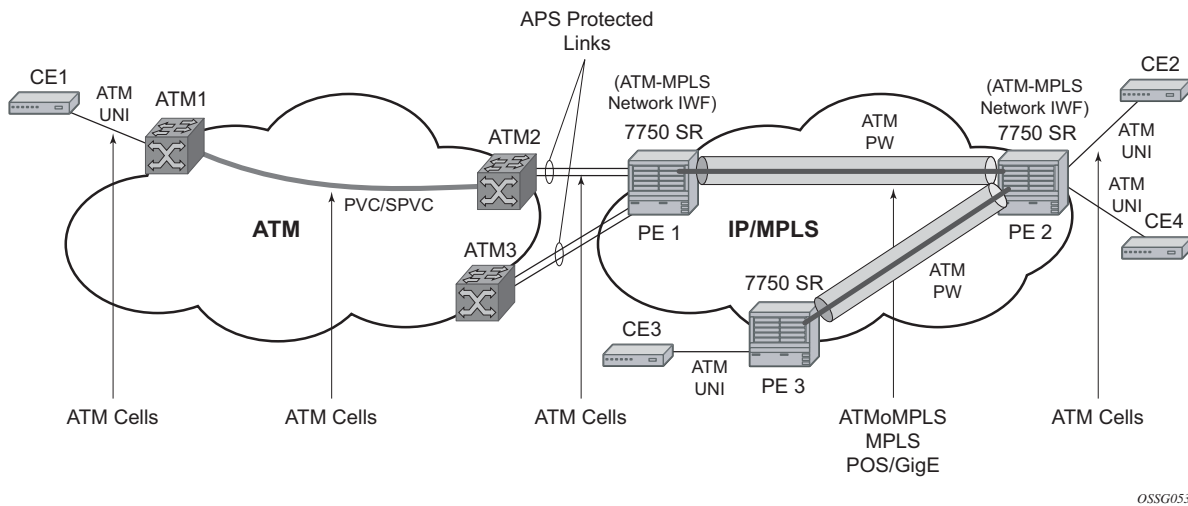


Figure 1: ATM VLL for End-to-End ATM Service

ATM Virtual Trunk Over IP/MPLS Packet-Switched Network

ATM virtual trunk (VT) implements a transparent trunking of user and control traffic between two ATM switches over an ATM pseudowire. Figure 2 depicts ATM 2 and ATM 3 switches that appear as if they are directly connected over an ATM link. Control traffic includes PNNI signaling and routing traffic.




Figure 2: VT Application Example

The virtual trunk (VT) SAP on a PE is identified by a tuple (port, VPI-range) meaning that all cells arriving on the specified port within the specified VPI range are fed into a single ATM pseudowire for transport across the IP/MPLS network. Note that a user can configure the whole ATM port as a

VT and does not need to specify a VPI range. No VPI/VCI translation is performed on ingress or egress. Cell order is maintained within a VT. Note that, as a special case, the two ATM ports could be on the same PE node.

By carrying all cells from all VPIs making up the VT in one pseudowire, a solution is provided that is both robust, for example no black holes on some VPIs but not others, as well as operationally efficient since the entire VT can be managed as a single entity from the Network Manager (single point for configuration, status, alarms, statistics, etc.).

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

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

Circuit Emulation Services (Cpipe)

Mobile Infrastructure

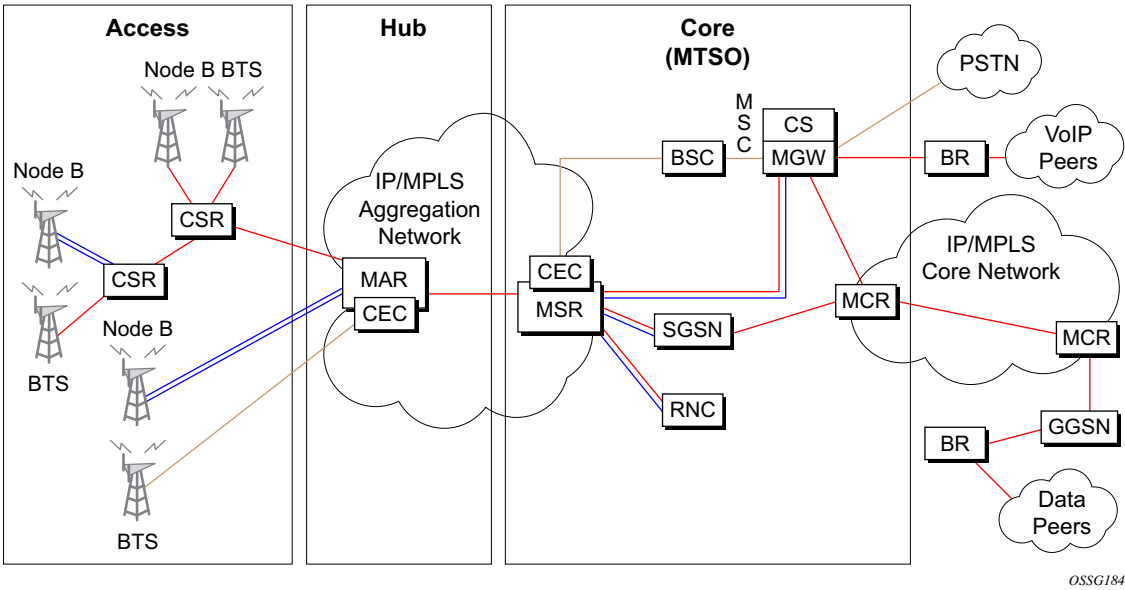


Figure 3: Mobile Infrastructure

Table 1: Mobile Infrastructure Definitions

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

Packet infrastructure is required within 2G, 2.5G and 3G mobile networks to handle SMS messaging, web browsing and emerging applications such as streaming video, gaming and video

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

Typically, low cost devices are used at the many cell sites to transport multiple DS1 or E1 using ATM/IMA and TDM over an Ethernet/MPLS infrastructure. In Alcatel-Lucent applications, the circuit emulation would currently be performed using the 7705 SAR. This could be performed by DMXplore at the cell site. However, a large number of cell sites aggregate into a single switching center. Book-ending 7705 SAR nodes would require a very large number of systems at the switching center ([Figure 3](#)). Therefore, a channelized OC3/STM1 solution is much more efficient at the switching center. With the introduction of a channelized OC3/STM1 CES CMA/MDA in the 7750 SR, Alcatel-Lucent can provide a converged, flexible solution for IP/MPLS infrastructures for 2G/2.5G/3G mobile networks supporting both the CES (by CES CMA/MDA) and ATM/IMA transported traffic (by the ASAP MDA).

Circuit Emulation Modes

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

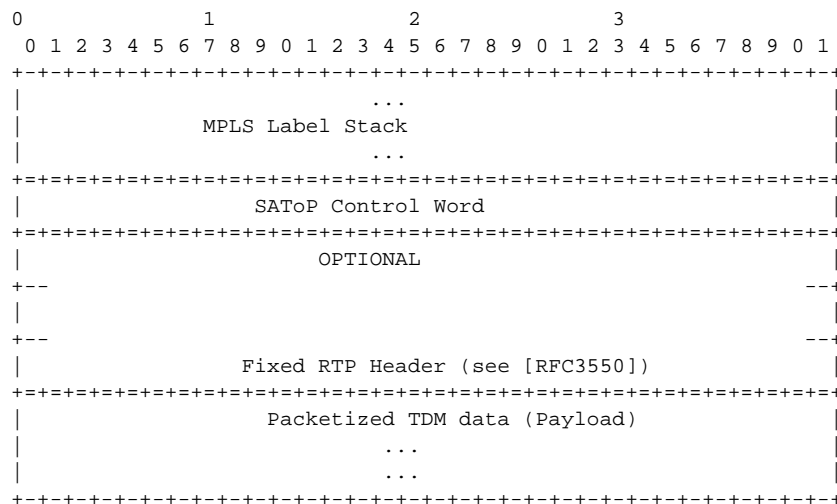


Figure 4: RFC 4553 (SAToP) MPLS PSN Encapsulation

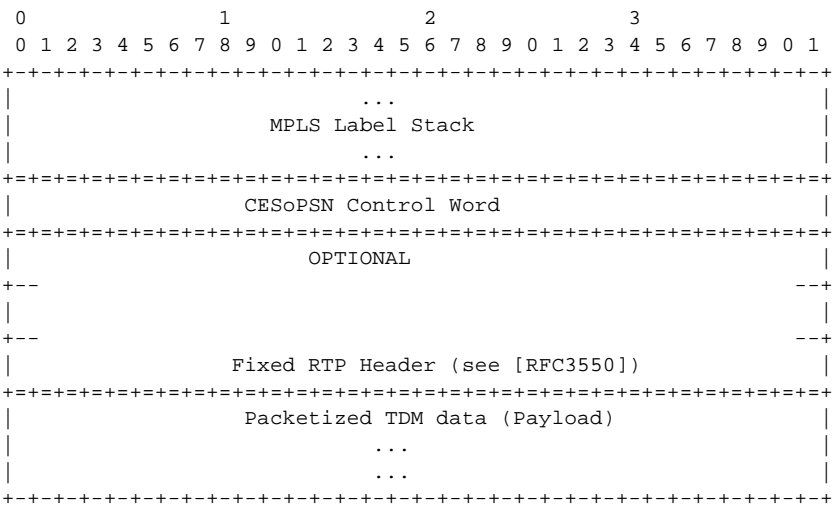


Figure 5: CESoPSN Packet Format for an MPLS PSN

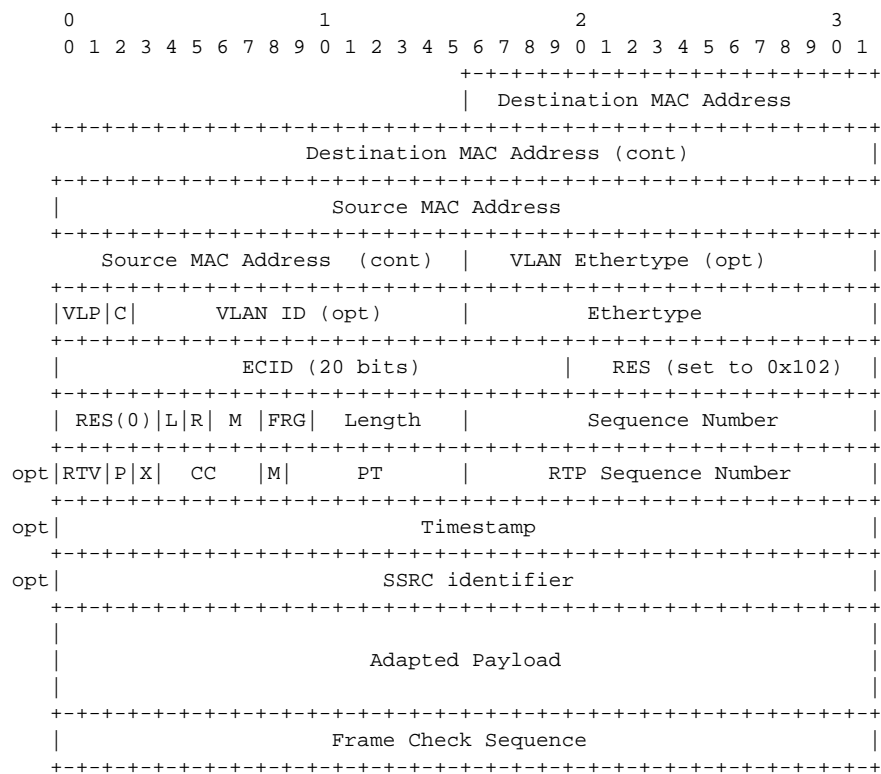


Figure 6: MEF8 PSN Encapsulation

Circuit Emulation Parameters

Circuit Emulation Modes

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

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

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

Absolute Mode Option

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

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

Payload Size

For DS3, E3, DS1 and E1 circuit emulation, the payload size can be configurable in number of octets. The default values for this parameter are shown in [Table 2](#). Unstructured payload sizes can be set to a multiple of 32 octets and minimally be 64 octets.

Table 2: Unstructured Payload Defaults

| TDM Circuit | Default Payload Size |
|--------------------|-----------------------------|
| DS1 | 192 octets |
| E1 | 256 octets |

For $n \times 64$ kbps circuits, the number of octets or DS1/E1 frames to be included in the TDM payload needs to be configurable in the range 4 to 128 DS1/E1 frames in increments of 1 or the payload size in octets. The default number of frames is shown in the table below with associated packet sizes. For the number of 64 kbps channels included (N), the following number of frames defaults apply for no CAS: $N=1$, 64 frames; $2 \leq N \leq 4$, 32 frames; $5 \leq N \leq 15$, 16 frames; $N \geq 16$, 8 frames. For CAS circuits, the number of frames can be 24 for DS1 and 16 for E1 which yields a payload size of $N \times 24$ octets for T1 and $N \times 16$ octets for E1. For CAS, the signaling portion is an additional $((N+1)/2)$ bytes where N is the number of channels. The additional signaling bytes are not included in the TDM payload size, although they are included in the actual packet size shown in [Table 3](#).

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

Table 3: Structured Number of Frames Defaults

| Num Timeslots | no CAS | | | DS1 CAS | | E1 CAS | |
|----------------------|---------------------------|------------------------|------------------------|-----------------------------|--------------------|-----------------------------|--------------------|
| | num-frames default | Default Payload | Minimum Payload | Pay-load (24 frames) | Packet Size | Pay-load (16 frames) | Packet Size |
| 1 | 64 | 64 | 40 | 24 | 25 | 16 | 17 |
| 2 | 32 | 64 | 64 | 48 | 49 | 32 | 33 |
| 3 | 32 | 96 | 96 | 72 | 74 | 48 | 50 |
| 4 | 32 | 128 | 128 | 96 | 98 | 64 | 66 |
| 5 | 16 | 80 | 80 | 120 | 123 | 80 | 83 |
| 6 | 16 | 96 | 96 | 144 | 147 | 96 | 99 |
| 7 | 16 | 112 | 112 | 168 | 172 | 112 | 116 |
| 8 | 16 | 128 | 128 | 192 | 196 | 128 | 132 |
| 9 | 16 | 144 | 144 | 216 | 221 | 144 | 149 |
| 10 | 16 | 160 | 160 | 240 | 245 | 160 | 165 |

Table 3: Structured Number of Frames Defaults (Continued)

| Num Timeslots | no CAS | | | DS1 CAS | | E1 CAS | |
|---------------|--------------------|-----------------|-----------------|----------------------|-------------|----------------------|-------------|
| | num-frames default | Default Payload | Minimum Payload | Pay-load (24 frames) | Packet Size | Pay-load (16 frames) | Packet Size |
| 11 | 16 | 176 | 176 | 264 | 270 | 176 | 182 |
| 12 | 16 | 192 | 192 | 288 | 294 | 192 | 198 |
| 13 | 16 | 208 | 208 | 312 | 319 | 208 | 215 |
| 14 | 16 | 224 | 224 | 336 | 343 | 224 | 231 |
| 15 | 16 | 240 | 240 | 360 | 368 | 240 | 248 |
| 16 | 8 | 128 | 128 | 384 | 392 | 256 | 264 |
| 17 | 8 | 136 | 136 | 408 | 417 | 272 | 281 |
| 18 | 8 | 144 | 144 | 432 | 441 | 288 | 297 |
| 19 | 8 | 152 | 152 | 456 | 466 | 304 | 314 |
| 20 | 8 | 160 | 160 | 480 | 490 | 320 | 330 |
| 21 | 8 | 168 | 168 | 504 | 515 | 336 | 347 |
| 22 | 8 | 176 | 176 | 528 | 539 | 352 | 363 |
| 23 | 8 | 184 | 184 | 552 | 564 | 368 | 380 |
| 24 | 8 | 192 | 192 | 576 | 588 | 384 | 396 |
| 25 | 8 | 200 | 200 | NA | NA | 400 | 413 |
| 26 | 8 | 208 | 208 | NA | NA | 416 | 429 |
| 27 | 8 | 216 | 216 | NA | NA | 432 | 446 |
| 28 | 8 | 224 | 224 | NA | NA | 448 | 462 |
| 29 | 8 | 232 | 232 | NA | NA | 464 | 479 |
| 30 | 8 | 240 | 240 | NA | NA | 480 | 495 |
| 31 | 8 | 248 | 248 | NA | NA | NA | NA |

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

Jitter Buffer

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

Basic CES (DS1 & E1):

N=1, 32 ms

2<=N<= 4, 16 ms

5<=N<=15, 8 ms

N>=16, 5 ms

CES Circuit Operation

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

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

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

Services for Transporting CES Circuits

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

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

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

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

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

Network Synchronization Considerations

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

Each DS-1 or E-1 channel without CAS signaling enabled can be independently configured to be loop-timed, node-timed, adaptive-timed or differential-timed. Each DS-1 or E-1 channel with CAS signaling enabled can be independently configured to be loop-timed or node-timed. Adaptive-timed and differential-timed are not supported on DS-1 or E-1 channels with CAS signaling enabled.

A CES circuit's adaptive recovered clock can be used a timing reference source for the node (ref1 or ref2). This is required to distribute network timing to network elements which only have packet connectivity to the network. One timing source on the CMA/MDA can be monitored for timing integrity. Both timing sources can be monitored if they are configured on separate CMA/MDAs while respecting the timing subsystem slot requirements. If a CES circuit is being used for adaptive clock recovery at the remote end (such that the local end is now an adaptive clock master), it is recommended to set the DS-1/E-1 to be node-timed to prevent potential jitter issues in the recovered adaptive clock at the remote device.

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

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

On the 7710 SR CES CMA, a BITS port is also provided. The BITS port can be used as one of the two timing reference sources in the system timing subsystem. The operation of BITS ports configured as ref1 or ref2 is the same as existing ports configured as ref1 and ref2 with all options supported. The operation of the 7750/7450 BITS source is unchanged and the BITS ports are not available on the CES MDAs (only SF/CPM BITS are currently available).

Cpipe Payload

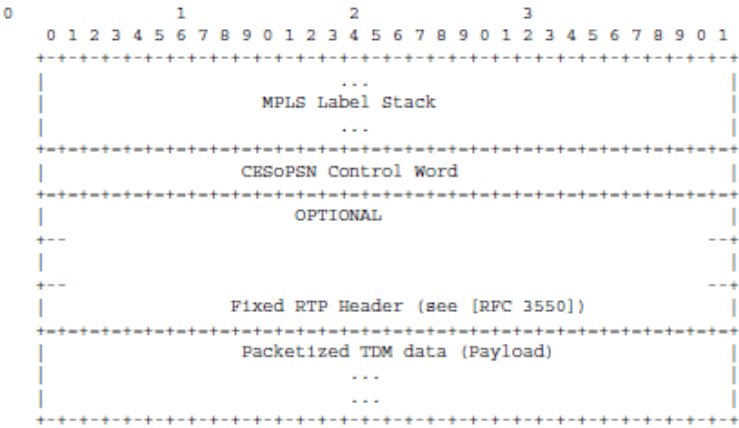


Figure 7: CESoPSN MPLS Encapsulation

Figure 7 shows the format of the CESoPSN TDM payload (with and without CAS) for packets carrying trunk-specific 64 kb/s service. In CESoPSN, the payload size is dependent on the number of timeslots used.

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 39](#)
 - [SAP Encapsulations and Pseudowire Types on page 156](#)
 - [QoS Policies on page 157](#)
 - [Filter Policies on page 157](#)
 - [MAC Resources on page 157](#)
- [Basic Configurations on page 160](#)
- [Common Configuration Tasks on page 160](#)
 - [Configuring VLL Components on page 161](#)
 - [Creating an Epipe Service on page 162](#)
- [Service Management Tasks on page 196](#)

Epipe Service Overview

An Epipe 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 Epipe 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 Epipe service is completely transparent to the subscriber's data and protocols. The 7450 ESS Epipe service does not perform any MAC learning. A local Epipe service consists of two SAPs on the same node, whereas a distributed Epipe service consists of two SAPs on different nodes. SDPs are not used in local Epipe services.

Each SAP configuration includes a specific port on which service traffic enters the 7450 ESS 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 8: Epipe/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 4](#) and [Table 5](#) 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 4 and Table 5 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 4: 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 5: 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 5: 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 7750 SR and 7450 ESS (mixed mode) and 7950 XRS platforms, in mode D with FP2+ (FP3 recommended).

ETH-CFM is supported for OAM services.

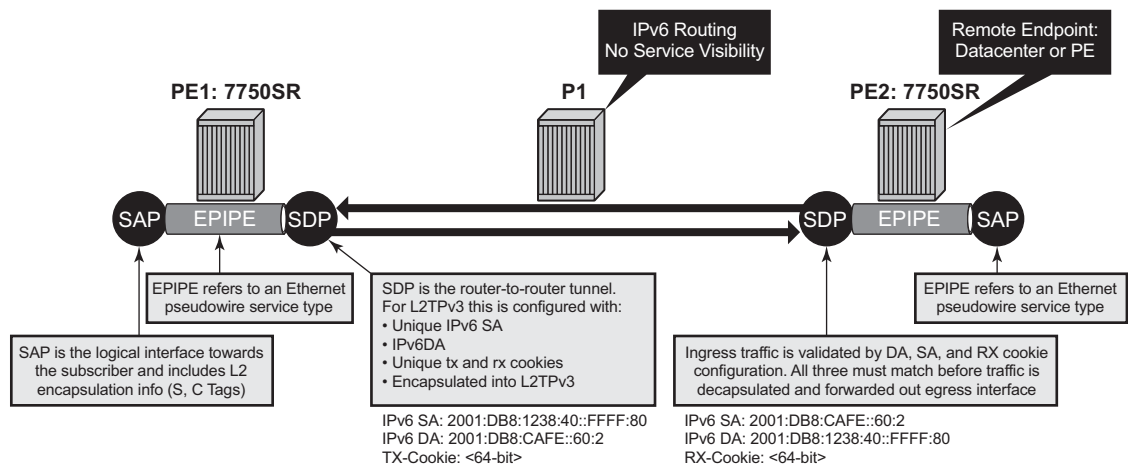


Figure 9: L2TPv3 SDP Illustration

Traffic Management Support

Ingress SAP Classification and Marking

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

Egress Network EXP Marking

FC-to-EXP mapping is as per the Network Egress QoS policy. Marking of the EXP field in both label stacks is performed.


Ingress Network Classification

Classification is based on the EXP value of the pseudowire label and EXP-to-FC mapping is as per Network Ingress QoS policy.

IP Interworking VLL (Ipipe) Services

- [IP Interworking VLL \(Ipipe\) Services on page 48](#)
 - [Ipipe VLL on page 48](#)
 - [IP Interworking VLL Datapath on page 50](#)
 - [IPv6 Support on IP Interworking VLL on page 54](#)
- [Basic Configurations on page 160](#)
- [Common Configuration Tasks on page 160](#)
 - [Configuring VLL Components on page 161](#)
 - [Creating an Ipipe Service on page 173](#)
- [Service Management Tasks on page 196](#)

Ipipe VLL






Figure 10: IP Interworking VLL (Ipipe)

Figure 10 provides an example of IP connectivity between a host attached to a point-to-point access circuit (FR, PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts appear to be on the same LAN segment. This feature enables service interworking between different link layer technologies. A typical use of this application is in a Layer 2 VPN when upgrading a hub site to Ethernet while keeping the spoke sites with their existing Frame Relay routed encapsulation.

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

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

IP Interworking VLL Datapath




Figure 11: IP Interworking VLL Datapath

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

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

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

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

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. The router does not support invATM ARP on an ATM interface.

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 7450 ESS 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.

VLL FR SAP Procedures

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

- Allow the service to come up without the CE address parameter configured at both the SAP and spoke SDP. If one or both parameters are configured, they are ignored.
- The operator cannot configure the **ce-address** parameter under SAP or spoke SDP when the **ce-address-discovery** option under the VLL service is enabled. Conversely, the operator is not allowed to enable the **ce-address-discovery** option under the Ipipe service if it has a SAP and/or spoke SDP with a user-entered **ce-address** parameter.
- If the router receives an invFR ARP request message over the FR SAP, it updates the ARP cache with the FR CE address. It also replies with the IP address of the CE attached to the remote PE if a valid address was advertised in the address list TLV by this remote PE. Otherwise, the router updates the ARP cache but does not reply to the invFR ARP.
- If the Ipipe service 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 is used.
- If the remote PE included the address list TLV in the received label mapping message, the local router updates the remote PE node with the most current IP address of the FR CE using a T-LDP status notification message containing an LDP address list. The notification message is sent each time an IP address different from the current value in the ARP cache is learned. This includes when the ARP is flushed and the CE address is reset to the value of 0.0.0.0.
- If the remote PE did not include the address list TLV in the received label mapping message, the local router does not send any notification messages containing the address list TLV during the lifetime of the IP pseudowire.

IPv6 Support on IP Interworking VLL

The 7450 ESS supports 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 12](#) illustrates the data path operation for IPv6 on an IP Interworking VLL between the Ethernet and PPP (IPv6CP) SAPs.

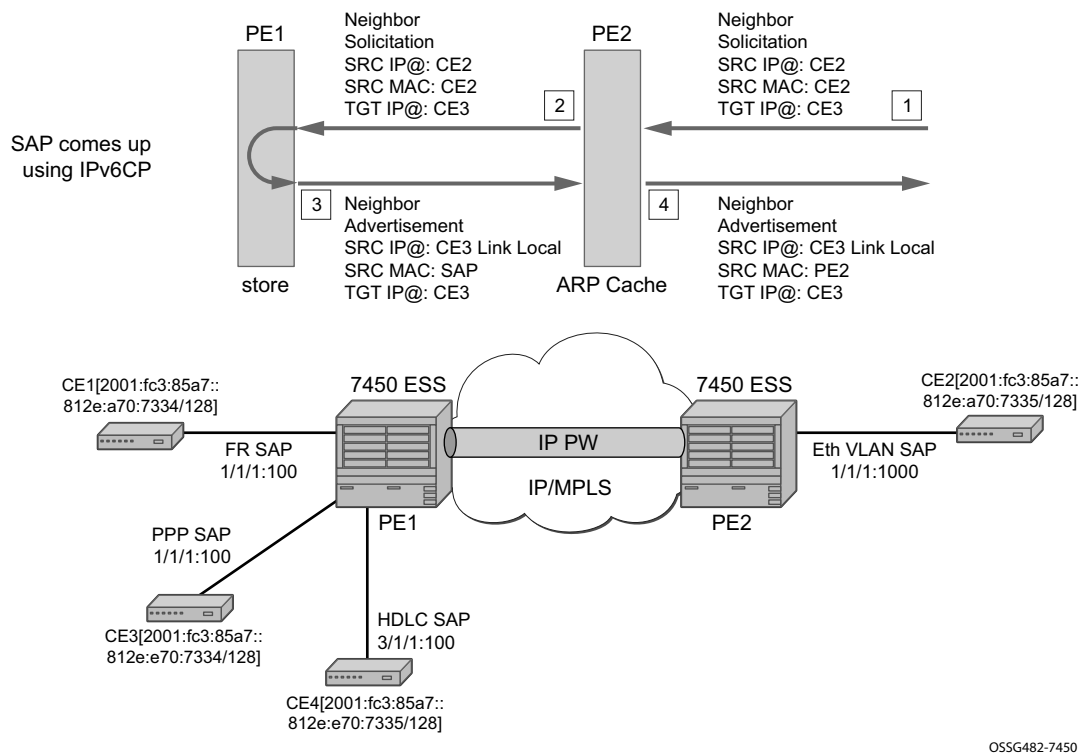


Figure 12: Data Path for Ethernet CE to PPP Attached CE

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

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

Services Configuration for MPLS-TP

MPLS-TP PWs are supported in epipe, apipe and cpipe VLLs and epipe spoke termination on IES/ VPRN and VPLS, iVPLS and B-VPLS.

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

MPLS-TP SDPs

Only MPLS SDPs are supported.

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

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

The following CLI shows the new MPLS-TP options:.

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

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

The system will accept the node-id being entered in either 4-octet IP address format <a.b.c.d> or unsigned integer format.

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

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

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

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

Signaling LDP or BGP is blocked if:

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

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

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

VLL Spoke SDP Configuration

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

vc-switching is supported.

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

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

```

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

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

```

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

The **control-channel-status** command enables static pseudowire status signaling. This is valid for any spoke-sdp where **signaling none** is configured on the SDP (for example, where T-LDP

signaling is not in use). The refresh timer is specified in seconds, from 10-65535, with a default of 0 (off). This value can only be changed if **control-channel-status** is **shutdown**. Commands that rely on PW status signaling are allowed if control-channel-status is configured for a spoke-sdp bound to an SDP with signaling off, but the system will use control channel status signaling rather than T-LDP status signaling. The ability to configure control channel status signaling on a given spoke-sdp is determined by the credit based algorithm described earlier. Control-channel-status for a particular pseudowire only counts against the credit based algorithm if it is in a **no shutdown** state and has a non-zero refresh timer and a non-zero request timer.

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

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

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

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

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

Only static pw to static pw switching is supported for MPLS-TP. Therefore, the vc-switching command is mutually exclusive with the configuration of the MPLS-TP parameters if the mate PW is not configured for an SDP with signaling off. However, vc-switching is supported if the mate SDP has signaling off.

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

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

In the vc-switching case, if configured on a mate spoke-sdp, then the TAI of the spoke-sdp must match the SAI of its mate, and SAI of spoke-sdp has to match the TAI of its mate.

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

- in network-mode D
- sdp signaling is off

- control-word is enabled (control-word by default is disabled)
- the service type is epipe, apipe, vpls, cpipe, or IES/VP RN interface
- mate sdp signaling is off (in vc-switched services)
- pw-status-signaling is enabled (see below)
- pw-path-id is configured for this spoke.

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

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

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

The CLI commands to configure control channel status requests are shown, below:

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

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

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

retry-timer <timer2> : 3-60s

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

timeout-multiplier <value> - 3-15.

- If a requesting node does not hear back after retry-timer times multiplier, then it must assume that the peer is down. This defaults to zero (0) when no retry-timer.

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

All existing commands (except for those explicitly specified below) are supported for spoke-sdp termination on IES, VPRN and VPLS (VPLS, iVPLS and bVPLS and routed VPLS) services. In addition, the MPLS-TP commands listed above are supported. The syntax and default values, and functional behavior of these commands is the same as for Epipe VLLs, as specified above.

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

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

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

Configuring MPLS-TP Lock Instruct and Loopback

MPLS-TP supports lock instruct and loopback for PWs. The topics in this section are:

- [MPLS-TP PW Lock Instruct and Loopback Overview on page 64](#)
 - [Lock PW End-Point Model on page 65](#)
 - [PW Redundancy and Lock Instruct and Loopback on page 65](#)
 - [Configuring a Test SAP for an MPLS-TP PW on page 66](#)
 - [Configuring an Administrative Lock on page 67](#)
 - [Configuring a Loopback on page 68](#)
 - [Configuring a Loopback on page 68](#)
-

MPLS-TP PW Lock Instruct and Loopback Overview

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

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

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

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

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

- An MPLS-TP loopback can be created for an epipe, cpipe or apipe VLL

- Test traffic can be inserted at an epipe, cpipe or apipe VLL endpoint or at an Epipe spoke-sdp termination on a VPLS interface

Lock PW End-Point Model

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

- Both ends of a PW or MS-PW represented by a spoke-sdp must be administratively locked.
 - Test traffic can be injected into the spoke-sdp using a SAP defined within a test service. The test service must be identified in the **tools** command at one end of the locked PW.
 - All traffic is forwarded to and from the test SAP defined in the test service, which must be of a type that is compatible with the spoke-sdp.
 - Traffic to and from a non-test-SAP is dropped. If no test SAP is defined all traffic received on the spoke-SDP is dropped, and all traffic received on the paired SAP is also dropped.
 - If a spoke-sdp is administratively locked, it is treated as operationally down. If a VLL SAP is paired with a spoke-sdp that is administratively locked, the SAP OAM treats this as if the spoke-sdp is operationally down.
 - If a VPLS interface is paired to a spoke-sdp that is administratively locked, the L2 interface is taken down locally.
 - Control-channel-status must be shutdown prior to administratively locking a spoke-sdp.
-

PW Redundancy and Lock Instruct and Loopback

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

Configuring a Test SAP for an MPLS-TP PW

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

```

config
  service
    epipe <service-id> [test][create]
      [no] sap <sap-id>
      [no] shutdown
    [no] shutdown
config
  service
    apipe <service-id> [vc-type {atm-vcc | atm-sdu | atm-vpc | atm-cell}
      [test][create]
      [no] sap <sap-id>
      [no] shutdown
    [no] shutdown
config
  service
    cpipe <service-id> [vc-type {satop-e1 | satop-t1 | satop-e3 | satop-t3 | cesopsn |
      cesopsn-cas} [test][create]
      [no] sap <sap-id>
      [no] shutdown
    [no] shutdown

```

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

- Ethernet NULL, .1q, Q-in-Q
- ATM VC, VP, VT and so on
- TDM E1, E3, DS0, DS3 and so on

The following constraints and conditions apply:

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

Configuring an Administrative Lock

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

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

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

- Can be configured either on a spoke-sdp that is bound to a SAP, another spoke-sdp or a VPLS interface.
- Is only allowed if a PW path ID is defined (for example, for static PWs with MPLS-TP identifiers).
- Cannot be configured on spoke-sdps that are an ICB or if the vc-switching keyword is present.
- The control-channel-status must be shutdown. The operator should also shutdown control-channel-status on spoke-sdps belonging to an MS-PW at an S-PE whose far ends are administratively locked at its T-PEs. This should be enforced throughout the network management if using the Service Access Manager.
- When enabled, all traffic on the spoke-sdp is sent to and from a paired SAP that has the **test** keyword present, if such a SAP exists in the X endpoint. Otherwise all traffic to and from the paired SAP is dropped.
- Can be configured at a spoke-sdp that is bound to a VLL SAP or a VPLS interface.
- The **test-svc-id** parameter refers to the test service that should be used to inject test traffic into the service. The test service must be of a compatible type to the existing spoke-sdp under test (see [Table 6](#)).
- If the **test-svc-id** parameter is not configured on an admin-locked spoke-sdp, then user traffic is simply blocked on the spoke-sdp.

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

[Table 6](#) illustrates the mapping between supported real services and their corresponding test services.

Table 6: Mapping of Real Services to Test Service Types

| Service | Test Service |
|----------|--------------|
| CPIPE | CPIPE |
| EPIPE | EPIPE |
| APIPE | APIPE |
| VPLS | EPIPE |
| PBB VPLS | EPIPE |

Configuring a Loopback

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

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

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

The following constraints and conditions apply for PW loopback configuration:

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

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

Topics include:

- [VCCV BVD Support on page 69](#)
- [VCCV BFD Encapsulation on a Pseudowire on page 70](#)
- [BFD Session Operation on page 70](#)
- [Configuring VCCV BFD on page 71](#)

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**.

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 13](#), 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.

```
| 7x50 PE1 (Epipe) | ---sdp 2:10--- | 7x50 PW SW (Epipe) | ---sdp 7:15--- | 7x50 PE2 (Epipe)
```

Figure 13: Pseudowire Service Switching Node

Configuration examples can be found in [Configuring Two VLL Paths Terminating on T-PE2 on page 182](#).

Pseudowire Switching with Protection

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

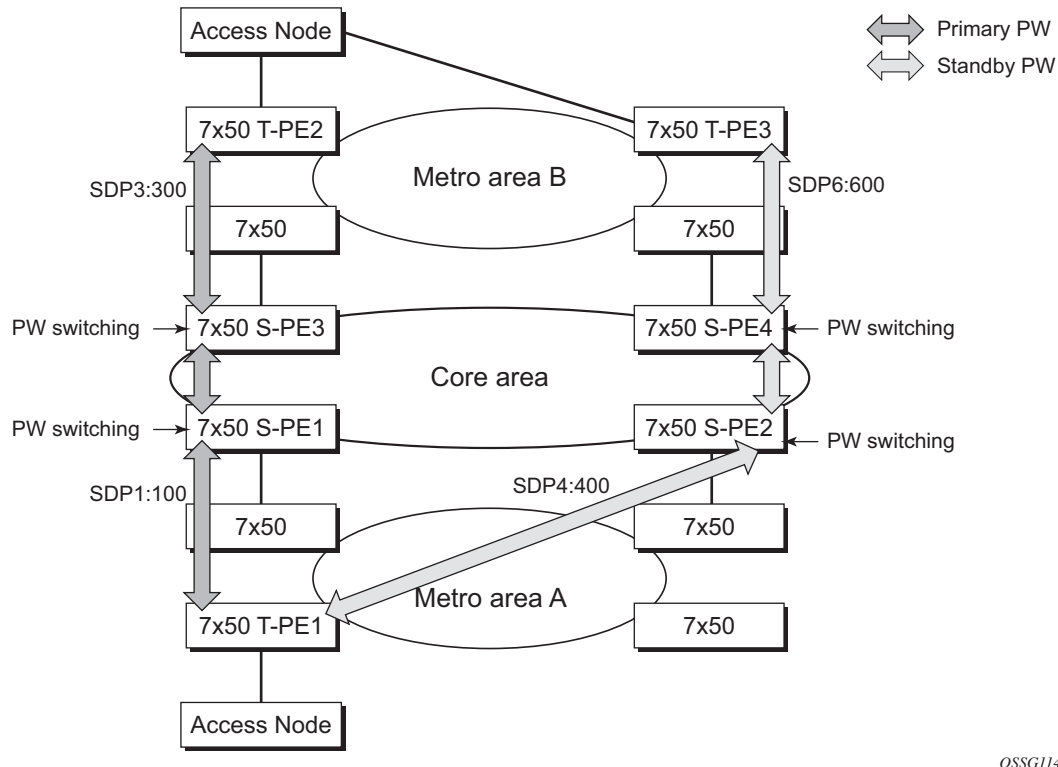


Figure 14: VLL Resilience with Pseudowire Redundancy and Switching

In the network in [Figure 14](#), PE nodes act as masters and pseudowire switching nodes act as slaves for the purpose of pseudowire signaling. A switching node will need to pass the SAP Interface Parameters of each PE to the other. 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 14](#), PE nodes act as masters and pseudowire switching nodes act as slaves for the purpose of pseudowire signaling. This is because a switching node will need to pass the SAP interface parameters of each PE to the other. T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node, for example, S-PE1. It will include the SAP interface parameters, such as MTU, in the label mapping message. S-PE1 checks the FEC against the local information and if a match exists, 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.

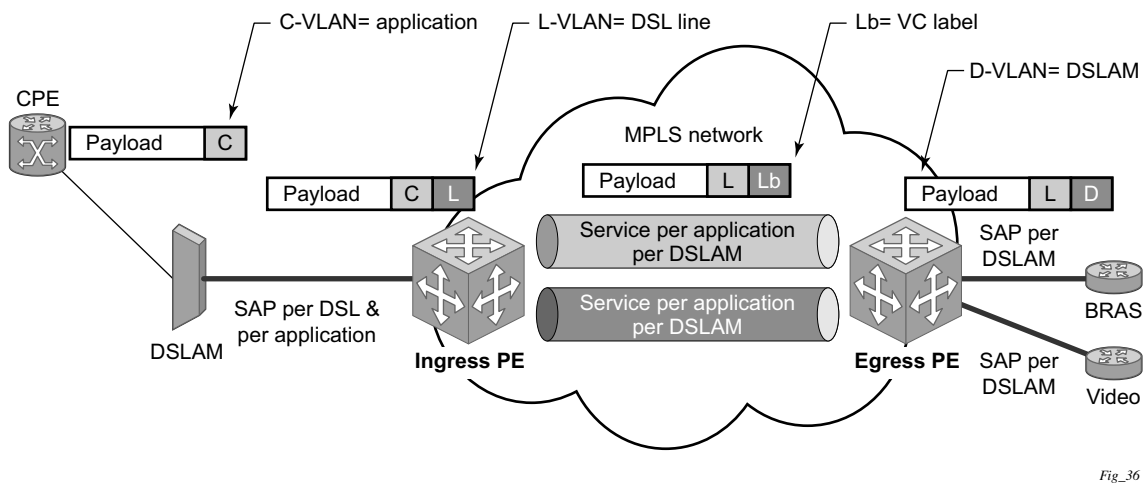
When the label mapping for the egress direction is also received from the T-LDP peer, and the information in the FEC matches that of the local configuration, the static-to-dynamic cross-connect is effected.

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 15: Ingress VLAN Swapping

The network diagram in [Figure 15](#) 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 16: Ingress VLAN Translation

The drawing in [Figure 16](#) 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 16](#)) 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 17 illustrates the operation of dynamic MS-PWs.




Figure 17: Dynamic MS-PW Overview

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

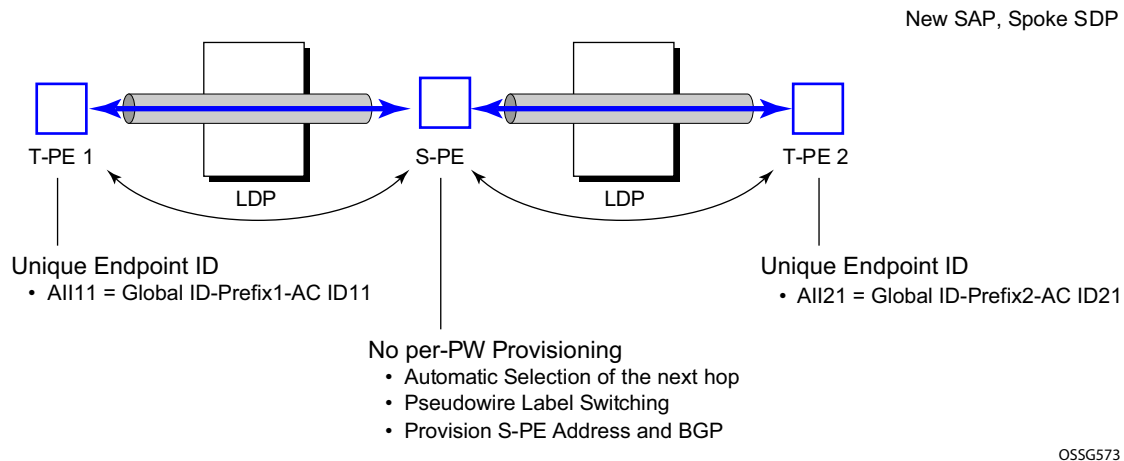


Figure 18: MS-PW Addressing using FEC129 AII Type 2

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

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 19](#):

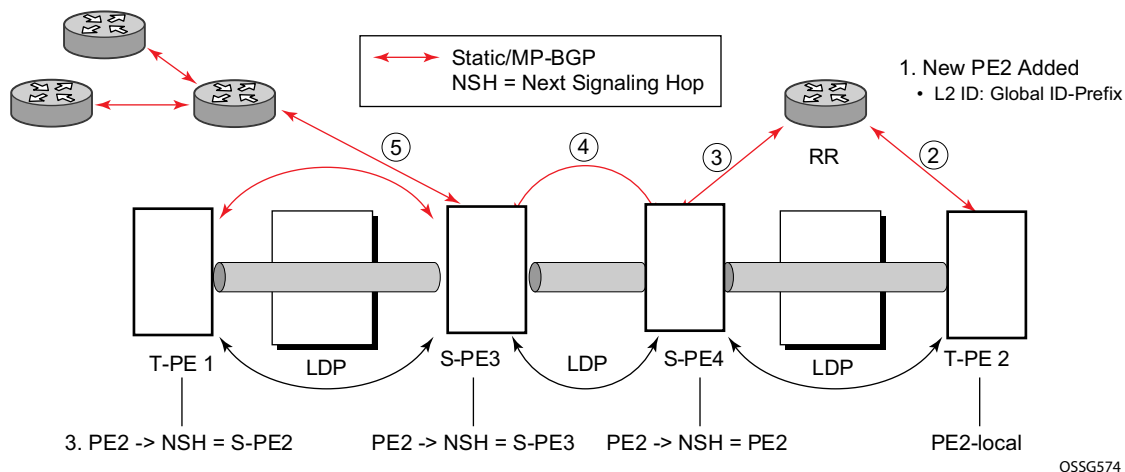


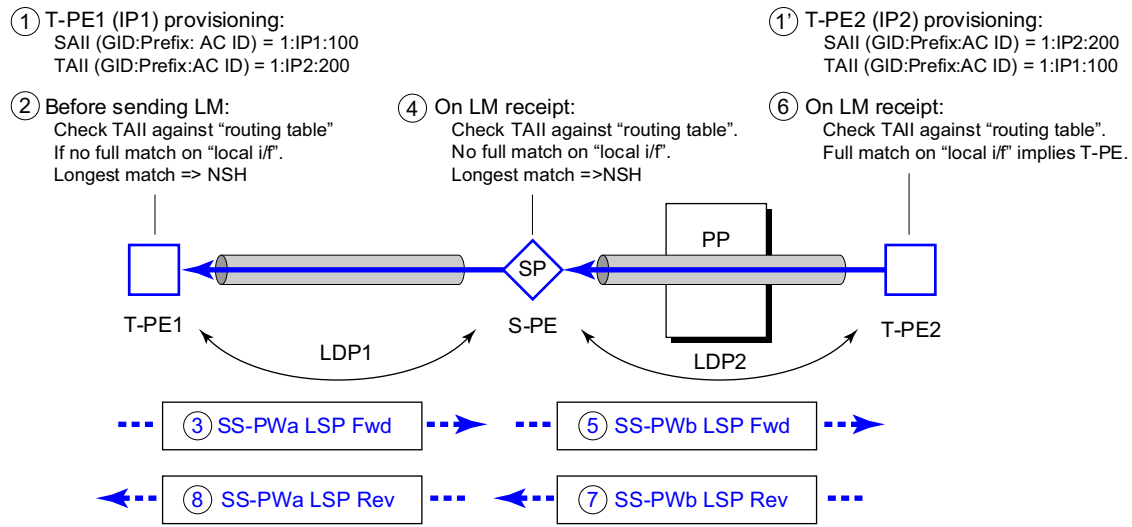
Figure 19: Advertisement of PE Addresses by PW Routing

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.

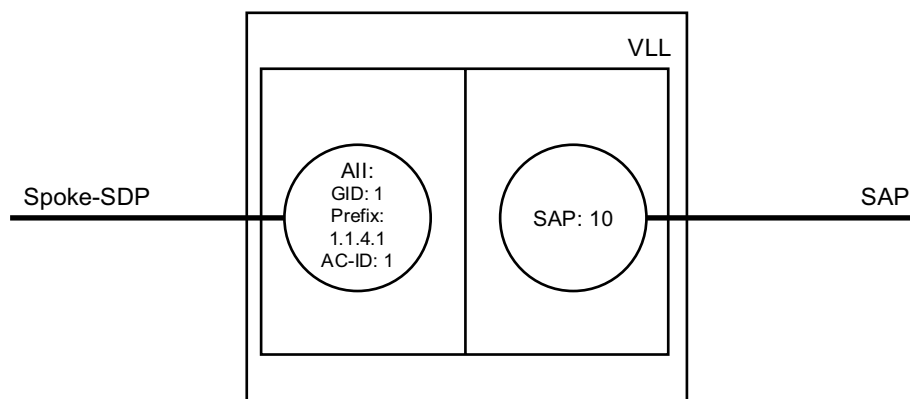
VLL services can then be established, as illustrated in [Figure 20](#).



OSSG575

Figure 20: 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 (SAI), Target AII (TAI). On the 7x50, the AIIs are locally configured for each spoke SDP, according to the model shown in Figure 21. 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 21: Mapping of AII to SAP

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

This signaling hop is then used to choose the T-LDP session to the chosen next-hop S-PE. Signaling proceeds through each subsequent S-PE using similar matching procedures to determine the next signaling hop. Otherwise, if a subsequent S-PE does not support dynamic MS-PW routing 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 22](#).




Figure 22: 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:

```
*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 TAII identifies the remote end. The SAII and TAII are each structured as follows:

- **Global-ID** — This is a 4 byte identifier that uniquely identifies an operator or the local network.
- **Prefix** — A 4-byte prefix, which should correspond to one of the local prefixes assigned under pw-routing.
- **AC-ID** — A 4-byte identifier for 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 TAII as unsigned integers. The T-PE with SAII>TAII assumes the active role. However, it is possible to override this behavior using the signaling **{master | auto}** command under 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 TAII associated with that spoke-sdp-fec. It allows a single-sided provisioning model where an incoming label mapping message with a TAII that matches the SAII of that spoke SDP 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 TAII 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 23 illustrates the use of pseudowire redundancy.

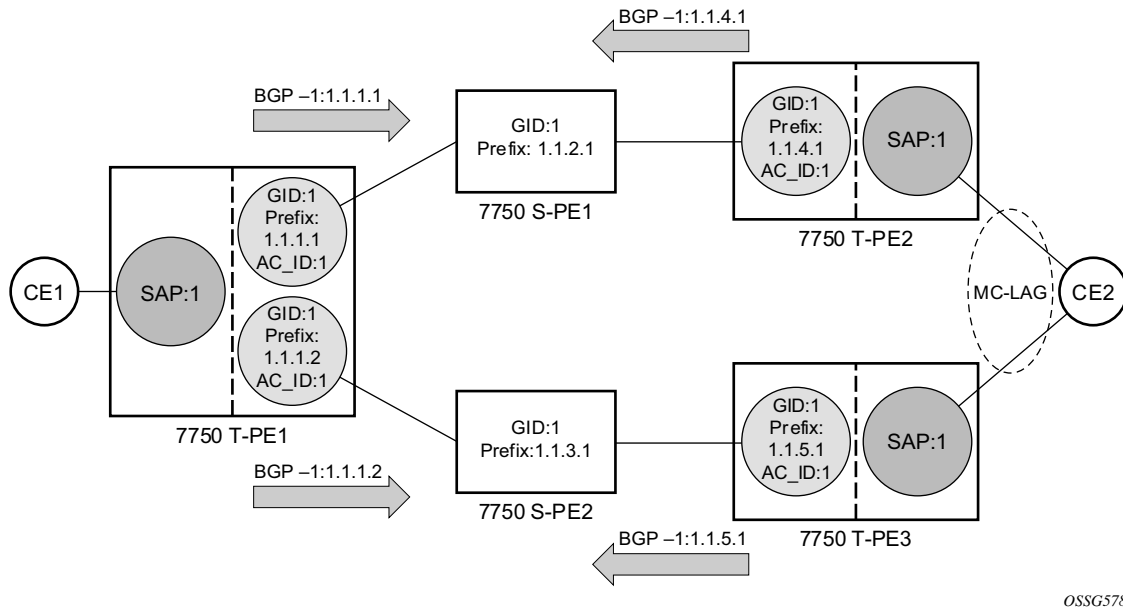


Figure 23: Pseudowire Redundancy

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

- Each MS-PW in the redundant set must have a unique 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 SAI and TAI to use.

If the SAI:TAI is entered by the user in the `vccv-ping` command, then those values are be used for the `vccv-ping` echo request, but their order is be reversed before being sent so that they match the order for the downstream FEC element for an S-PE, or the locally configured SAI:TAI for a remote T-PE of that MS-PW. Note that is SAI:TAI 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 SAI:TAI 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 SAI:TAI 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 SAI:TAI in the switching point TLV is maintained for the `vccv-ping` echo request message.

If no pseudowire switching point TLV was received, then the SAI:TAI 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 SAI:TAI 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 24: Dynamic MS-PW Example

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

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

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



OSSG115

Figure 25: VLL Resilience

If the Ethernet SAP on PE2 fails, PE2 notifies PE1 of the failure by either withdrawing the primary pseudowire label it advertised or by sending a pseudowire status notification with the code set to indicate a SAP defect. PE1 will receive it and will immediately switch its local SAP to forward over the secondary standby spoke SDP. 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 129](#), 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 26](#) illustrates the operation of master-slave pseudowire redundancy. In this scenario, an Epipe service is provided between CE1 and CE2. CE2 is dual homed to PE2 and PE3, 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.




Figure 26: Master-Slave Pseudowire Redundancy

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

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

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

The fact that the given spoke SDP has been blocked will be signaled to LDP peer through the pseudowire status bit (PW FWD Standby (0x20)). This will prevent traffic being sent over this spoke SDP by the remote peer, but obviously only in case that remote peer supports and reacts to pseudowire status notification. Previously, this applied only if the spoke SDP terminates on an IES, VPRN or VPLS. However, if standby-signaling-slave is enabled at the remote VLL endpoint then the Tx direction of the spoke SDP will also be blocked, according to the rules in [Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios on page 107](#).

Note that although master-slave operation provides bidirectional blocking of a standby spoke SDP during steady-state conditions, it is possible that the Tx directions of more than one slave endpoint can be active for transient periods during a fail-over operation. This is due to slave endpoints

transitioning a spoke SDP from standby to active receiving and/or processing a pseudowire preferential forwarding status message before those 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 27 shows an example for the case of Ethernet LMI.




Figure 27: 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 28 displays a VLL resilience path example. An sample configuration follows.




Figure 28: 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

```

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

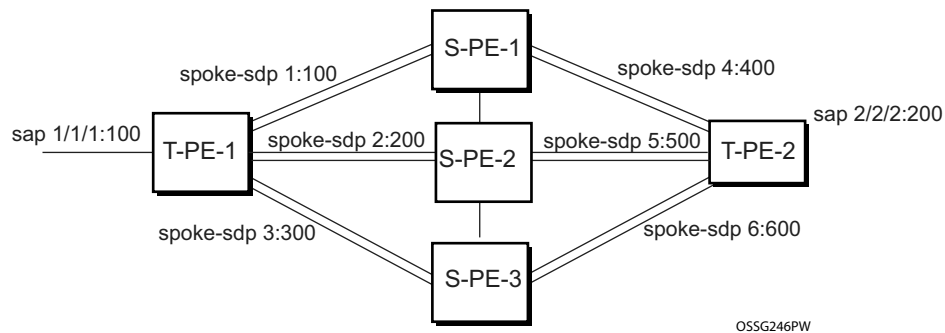


Figure 29: VLL Resilience with Pseudowire Switching

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
```

Epipe Using BGP-MH Site Support for Ethernet Tunnels

```
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
```

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 129](#)) 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.

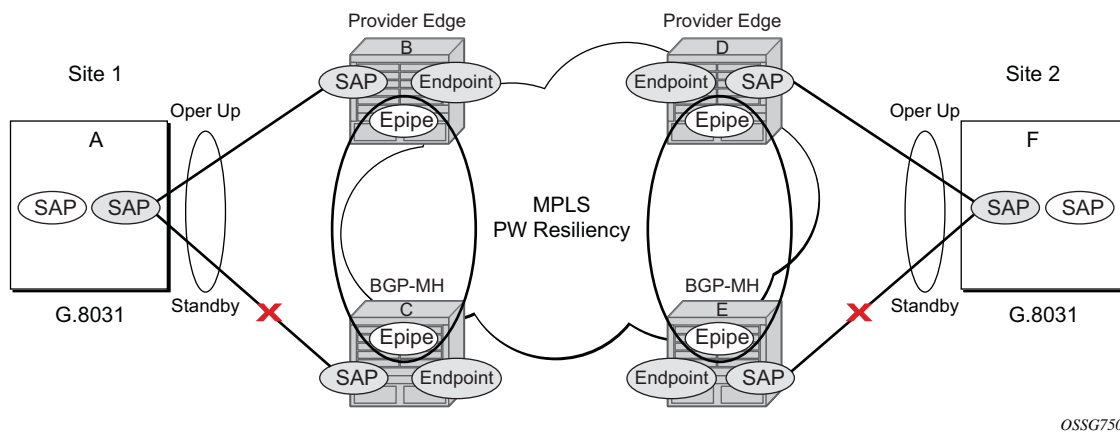


Figure 30: BGP-MH Site Support for Ethernet Tunnels

Figure 30 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 31](#) shows an Epipe service on a Customer Edge (CE) device that uses G.8031 with two paths and two MEPs. The Paths can use a single VLAN of DOT1Q or QinQ encapsulation.

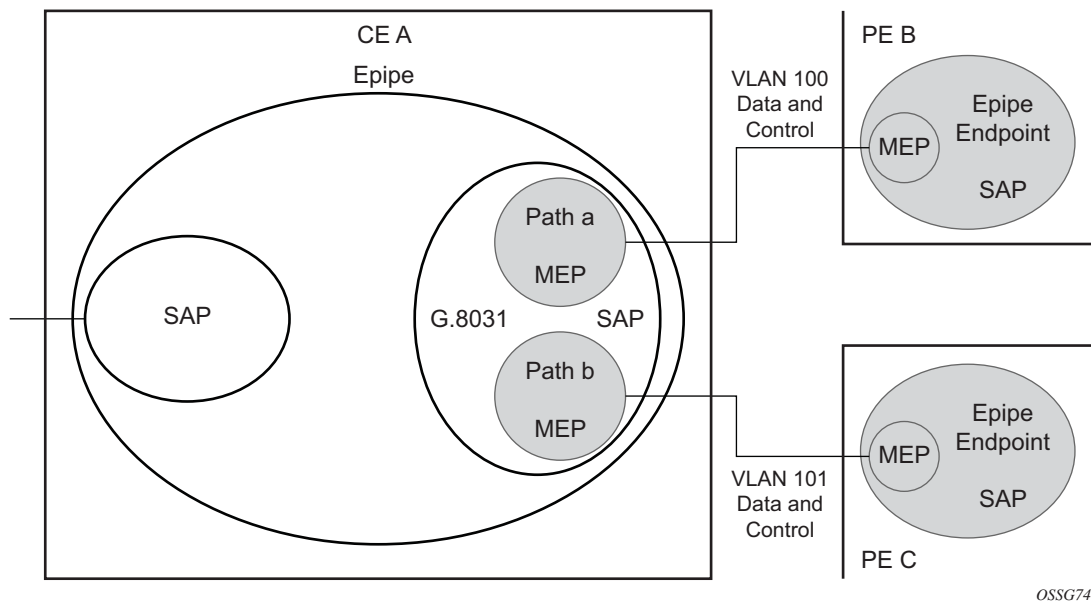


Figure 31: G.8031 for Slave Operation

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

To get fate sharing for multiple services with this feature, a dedicated G.8031 CE based service (one VLAN) is 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.


There are two configuration options:

- Every G.8031 service SAP terminates on a single Epipipe that has BGP-MH. These Epipipes may utilize endpoints with or without ICBs.
- A control Epipipe service that monitors a single SAP that is used for group control of fate shared CE services. In this case, the Epipipe 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 Epipipe 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 32: Full Redundancy G.8031 Epipipe & BGP-MH

[Table 7](#) lists the SAP MEP signaling shown in [Figure 32](#). For a description of the events shown in this sample operation, see [Events in Sample Operation on page 115](#).

Table 7: 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 32](#). This configuration uses operational groups. The nodes of interest are A, B and C listed in [Table 7](#).

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 7](#) 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 7](#) 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 7](#) 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 7](#) 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 7](#) 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 7](#) 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 494](#) 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 133](#) and [Figure 35 on page 131](#) 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 35](#) 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.

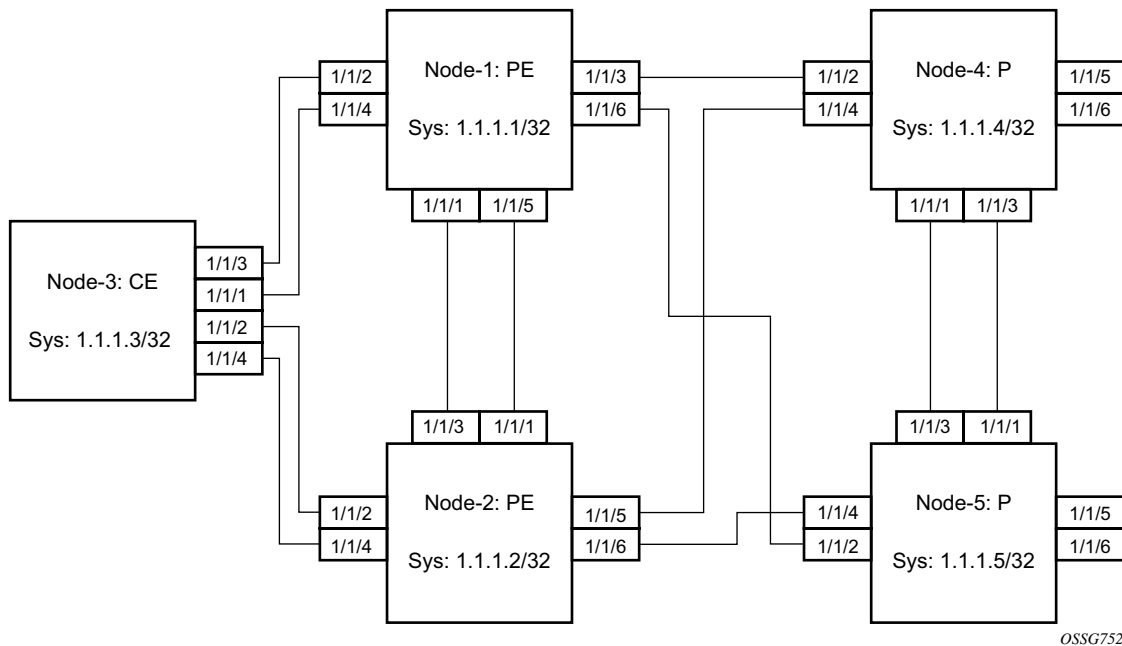


Figure 33: Sample Topology Full Redundancy

Refer to [Configuration Examples on page 121](#) for configuration examples derived from [Figure 33](#).

Configuration Examples

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

```
#-----
echo "Eth-CFM Configuration"
#-----

eth-cfm
  domain 100 format none level 3
    association 2 format icc-based name "node-3-site-1-0"
      bridge-identifier 1
      exit
      remote-mepid 310
    exit
    association 2 format icc-based name "node-3-site-1-1"
      bridge-identifier 100
      exit
      remote-mepid 311
    exit
  exit
exit

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

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

pw-template 1 create
  vc-type vlan
exit
oper-group "og-name-et" create
exit
oper-group "og-name-et100" create
exit
epipe 1 customer 1 create
  service-mtu 500
  bgp
    route-distinguisher 65000:1
    route-target export target:65000:1 import target:65000:1
  exit
  site "site-1" create
    site-id 1
    sap 1/1/2:1.1
    boot-timer 100
    site-activation-timer 2
    no shutdown
```

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

```

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
            exit
            association 3 format icc-based name "node-3-site-2-0"
                bridge-identifier 1
                exit
                ccm-interval 1
                remote-mepid 120
            exit
            association 3 format icc-based name "node-3-site-2-1"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 121
            exit
        exit
    exit

```

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

eth-tunnel 1
  description "Eth Tunnel loadsharing mode QinQ example"
  protection-type loadsharing
  ethernet
    encap-type qinq
  exit
  path 1
    member 1/1/3
    control-tag 1.1
    eth-cfm
      mep 310 domain 100 association 2
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
path 2
  member 1/1/4
  control-tag 1.2
  eth-cfm
    mep 320 domain 100 association 3
    ccm-enablepath
    control-mep
    no shutdown
  exit
exit
  no shutdown
exit
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
```

Node-3:

```

#-----
echo "Eth-CFM Configuration"
#-----

    eth-cfm
        domain 100 format none level 3
            association 1 format icc-based name "node-3-site-1-0"
                bridge-identifier 1
                exit
                ccm-interval 1
                remote-mepid 130
            exit
            association 2 format icc-based name "node-3-site-2-0"
                bridge-identifier 2
                exit
                ccm-interval 1
                remote-mepid 120
            exit
        exit
    exit

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

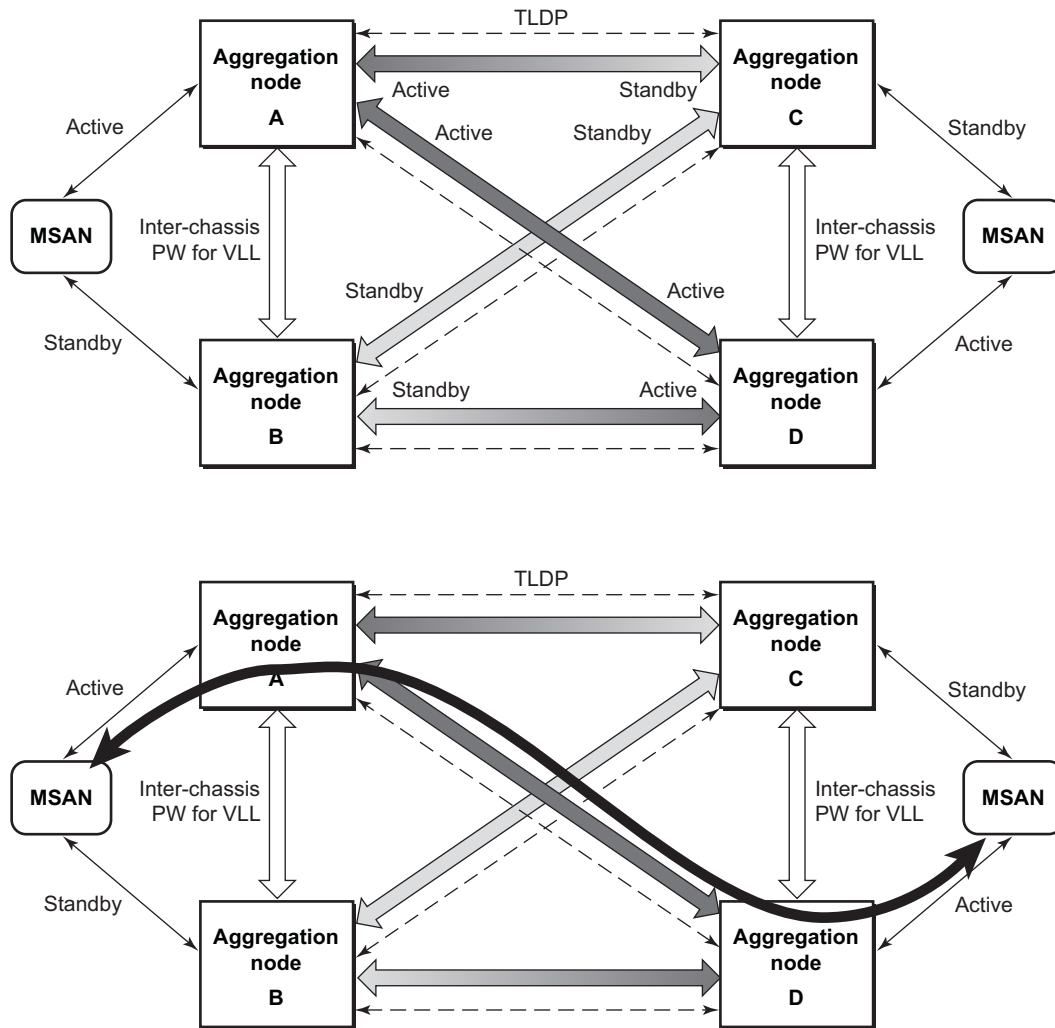
    eth-tunnel 2
        description "Eth Tunnel loadsharing mode QinQ example"
        protection-type loadsharing
        ethernet
            encap-type qinq
        exit
        path 1
            member 1/1/1
            control-tag 1.1
            eth-cfm
                mep 310 domain 100 association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

```

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

Access Node Resilience Using MC-LAG and Pseudowire Redundancy

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



OSSG116

Figure 34: Access Node Resilience

In this application, a new pseudowire status bit of active or standby indicates the status of the SAP in the MC-LAG instance in the SR-Series aggregation node. All spoke SDPs are of secondary type and there is no use of a primary pseudowire type in this mode of operation. Node A is in the active

state according to its local MC-LAG instance and 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.

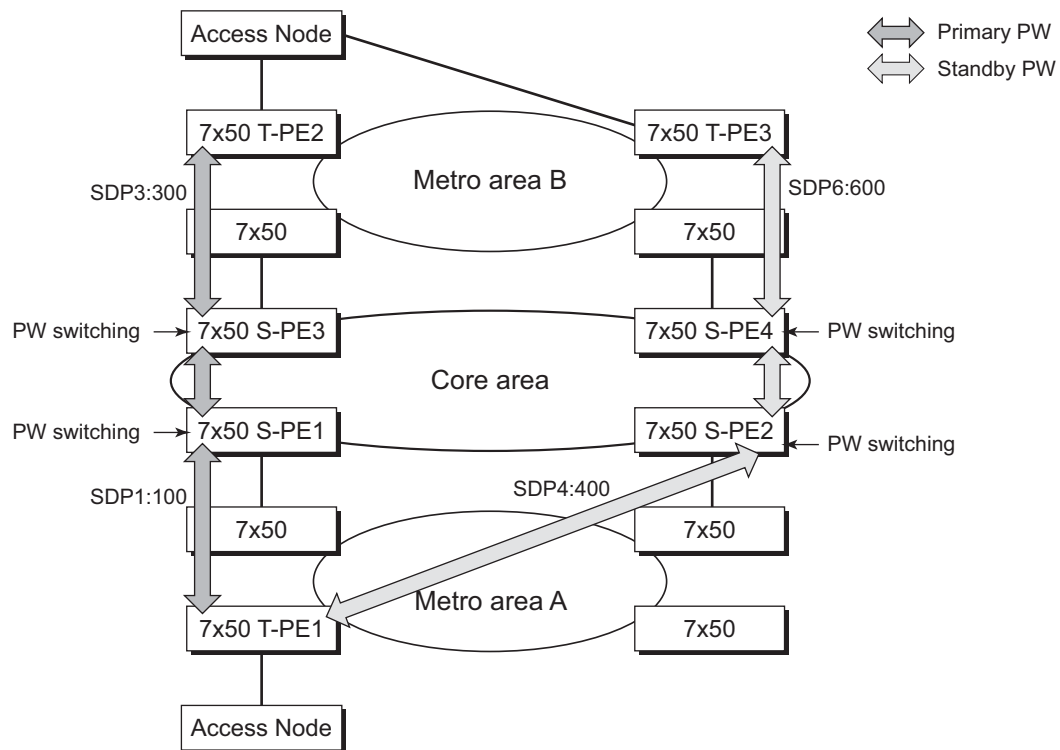
An SR-Series node selects a pseudowire as the active path for forwarding packets when both the local pseudowire status and the received remote pseudowire status indicate active status. However, an SR-Series device in standby status according to the SAP in its local MC-LAG instance is capable of processing 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 100](#). 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 35 illustrates the use of both pseudowire redundancy and pseudowire switching to provide a resilient VLL service across multiple IGP areas in a provider network.



OSSG114

Figure 35: VLL Resilience with Pseudowire Redundancy and Switching

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

Like in the application in [VLL Resilience with Two Destination PE Nodes on page 100](#), 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.

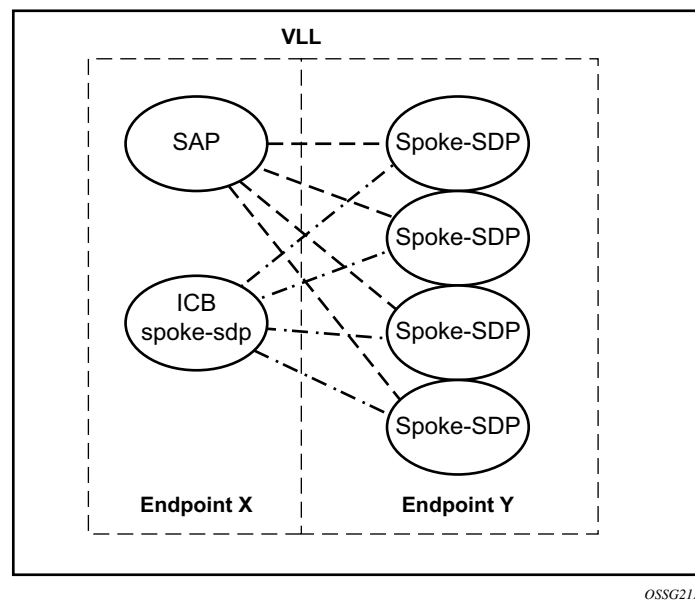
Pseudowire Redundancy Service Models

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 36](#) illustrates the model for a redundant VLL service based on the concept of endpoints.



OSSG211

Figure 36: 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 36](#).

Note that [Figure 36](#) 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 36 on page 133](#) 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 133](#) can be used on endpoints “X” and “Y”. The following sections refer to the specific combination objects in [Figure 36](#) 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

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.

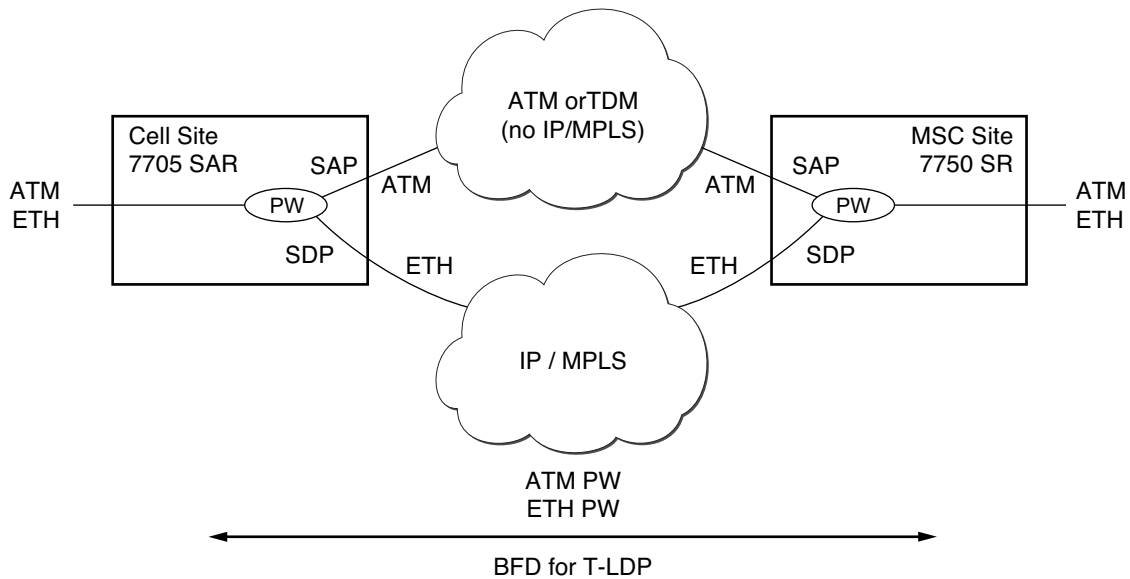
If a spoke SDP in endpoint “Y”, including the ICB spoke SDP, transitions locally to down state, the node must send T-LDP SDP-binding down status bits on this spoke SDP.

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.

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



OSSG483

Figure 37: HSDPA Off Load Fallback over ATM

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

For example, the service pseudowires provider may use a 7705 SAR with one or two ATM E1 uplinks for real-time voice traffic and an Ethernet uplink connected to a DSL model for NRT data traffic. At the RNC site, a 7750 SR or 7710 SR service router can be used, connected by ASAP (E1 IMA bundles) or STM-n ATM to the TDM/ATM network, and Ethernet to the DSL backhaul network.

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

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

Primary Spoke-SDP Fallback to Secondary SAP

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

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

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

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

Reversion to Primary Spoke SDP Path

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

MC-APS and MC-LAG

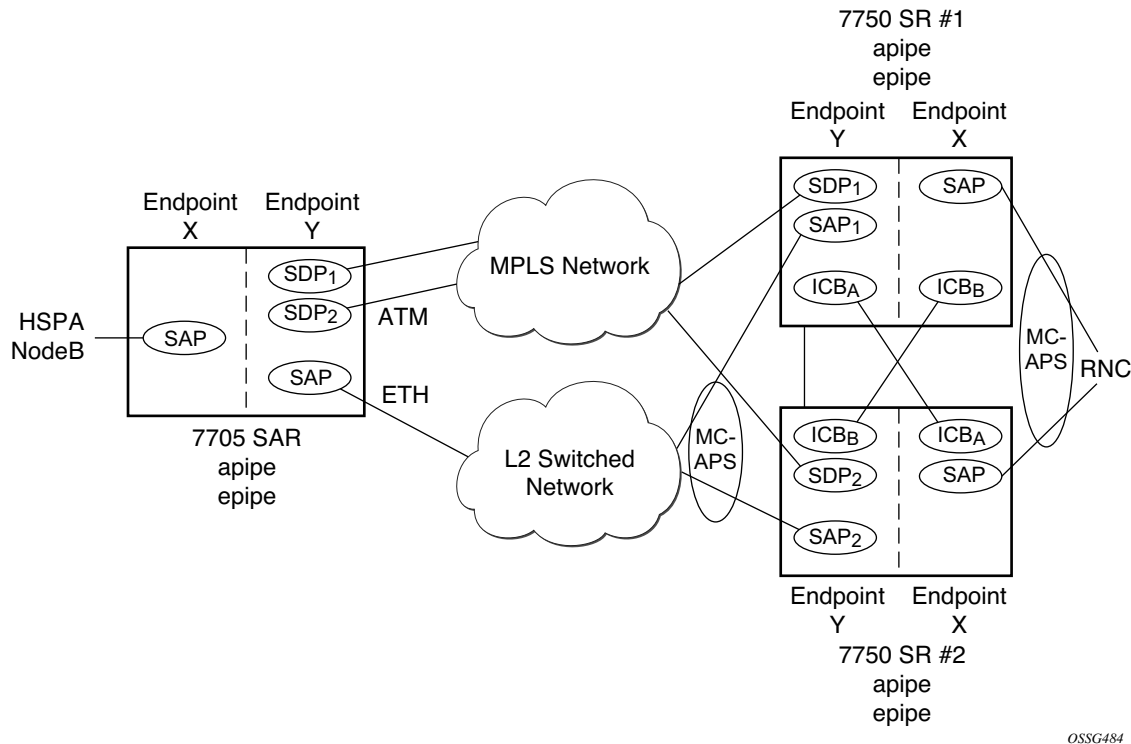


Figure 38: HSDPA Off Load Fallback with MC-APS

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

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

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

The following is an example configuration of Epipes mapping to [Figure 38](#). Note that a SAP can be added to an endpoint with a non-ICB spoke SDP only if the spoke's precedence is **primary**.

7750 SR#1:

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

7750 SR#2

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


Failure Scenarios

Following the before mentioned rules, the following are examples of a failure scenario operation. Assuming both links are active on 7750 SR#1 and the Ethernet connection to the cell site fails (most likely failure scenario as it would not be protected), SDP1 would go down and the secondary SAP would be used in 7750 SR#1 and 7705 SAR as shown in [Figure 39](#).

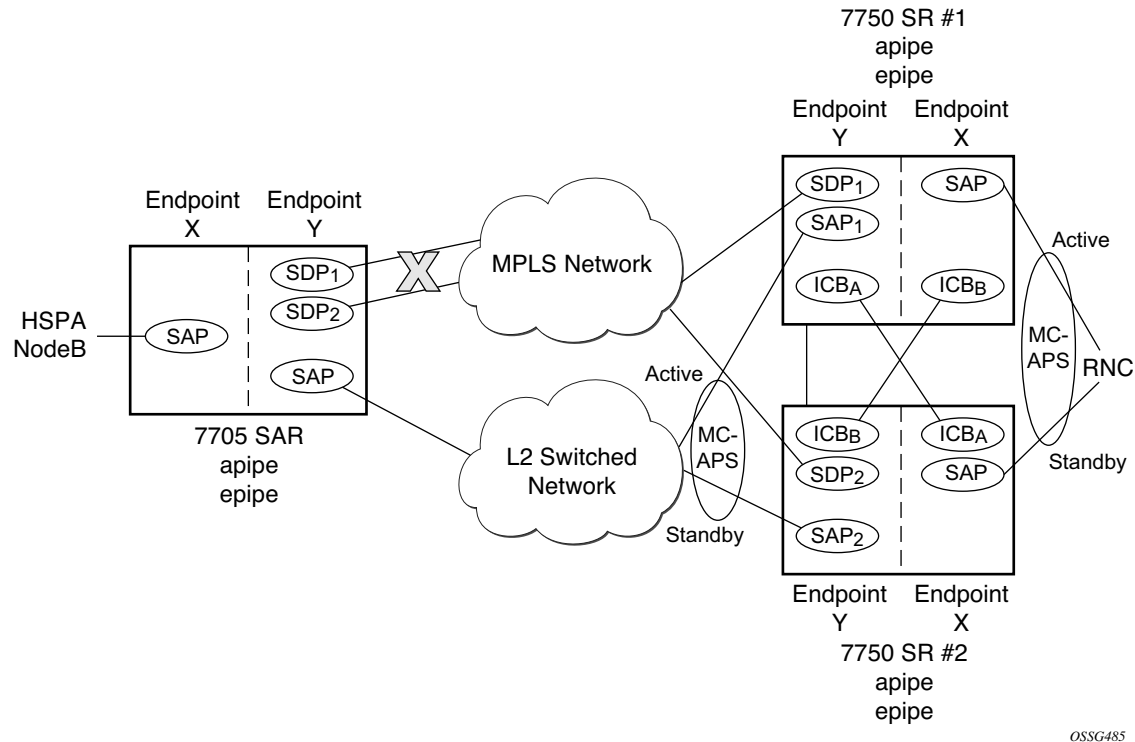


Figure 39: Ethernet Failure At Cell Site

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

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

VLL 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.

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

When using Ethernet Tunnels, the Ethernet Tunnel logical interface is created first. The Ethernet tunnel has member ports which are the physical ports supporting the links. The Ethernet tunnel control SAPs 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).

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

For further information, see the *Services Overview Guide*.

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 and a BGP signaled pseudowire, maintaining the Epipe properties such as no MAC learning. The pseudowire data plane uses a two label stack, the inner label is derived from the BGP signaling and identifies the Epipe service while the outer label is the tunnel label of an LSP transporting the traffic between the two end systems.

Figure 40 shows how this service would be used to provide a virtual lease-line service across an MPLS network between two sites, A and B.




Figure 40: Single-Homed BGP-VPWS Example

An Epipe is configured on PE1 and PE2 with BGP VPWS enabled. PE1 and PE2 are connected to site A and B, respectively, each using a SAP. The interconnection between the two PEs is achieved through a pseudowire 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 41](#).

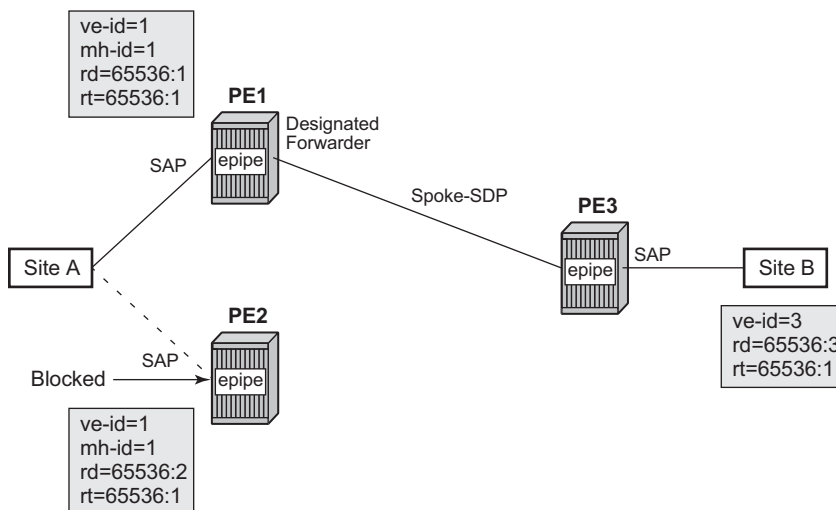


Figure 41: Dual-Homed BGP VPWS with Single Pseudowire

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

The pair of dual-homed PEs perform a designated forwarder election, which is influenced by BGP route selection, the site state, and by configuring the site-preference. A site will only be eligible to be the designated forwarder if it is up (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 42](#).

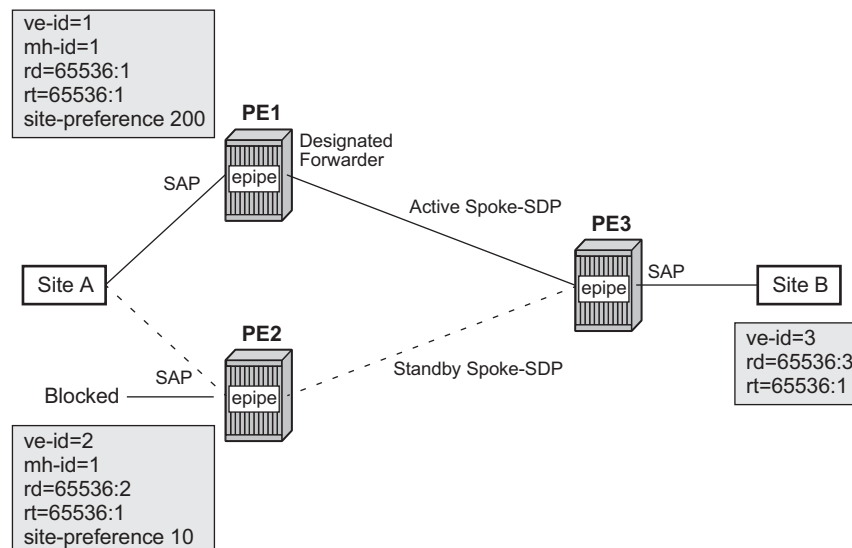


Figure 42: Dual-homed BGP VPWS with Active/Standby Pseudowires

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

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

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

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 43 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 43: BGP VPWS Update Extended Community Format

- Extended community type — The value allocated by IANA for this attribute is 0x800A
- Encaps Type — Encapsulation type, identifies the type of pseudowire encapsulation. Ethernet VLAN (4) and Ethernet Raw mode (5), as described in RFC 4448, are the only values supported. If there is a mismatch between the Encaps Type signaled and the one received, the pseudowire is created but with the 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|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 44](#).




Figure 44: 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 45](#).




Figure 45: BGP VPWS NLRI TLV Extension Format




Figure 46: Circuit Status Vector TLV Type

The circuit status vector is used to indicate the status of both the SAP 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 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 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
- 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

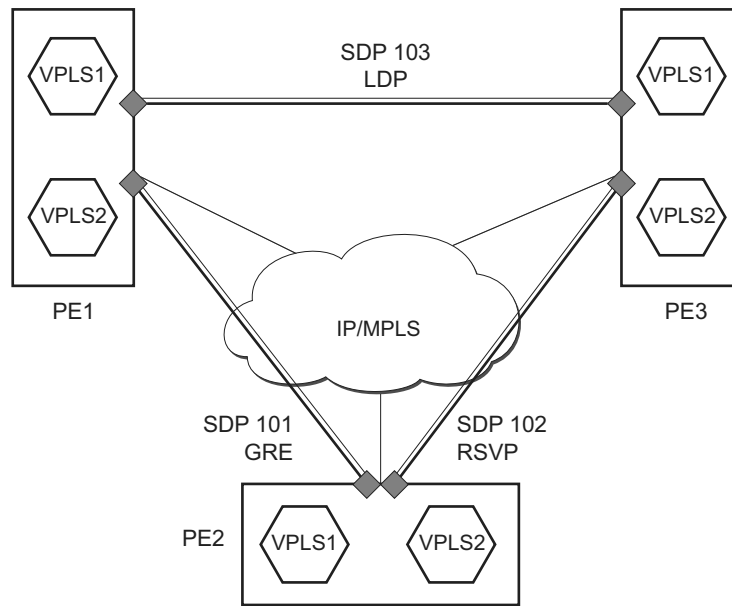
This section describes various of the general 7450 ESS service features and any special capabilities or considerations as they relate to VLL services.

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 47: SDP Statistics for VPLS and VLL Services

The simple three-node network described in [Figure 47](#) 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 7450 ESS Epipe service:

- Ethernet null
- Ethernet dot1q
- QinQ
- SONET/SDH BCP-null
- SONET/SDH BCP-dot1q
- FR VC with RFC 2427 Ethernet-bridged encapsulation

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 7450 ESS 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

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

MAC Resources

Epipe services are point-to-point layer 2 VPNs capable of carrying any Ethernet payloads. Although an Epipe is a Layer 2 service, the 7450 ESS 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 160](#)
- [Common Configuration Tasks on page 160](#)
 - [Configuring VLL Components on page 161](#)
 - [Creating an Epipe Service on page 162](#)
 - [Creating an Ipipe Service on page 173](#)
 - [Using Spoke SDP Control Words on page 177](#)
 - [Same Fate Epipe VLANs Access Protection on page 178](#)
 - [Configuring Pseudowire Scenarios](#)
 - [Pseudowire Configuration Notes on page 180](#)
 - [Configuring Two VLL Paths Terminating on T-PE2 on page 182](#)
 - [Configuring VLL Resilience on page 185](#)
 - [Configuring VLL Resilience for a Switched Pseudowire Path on page 186](#)
 - [Configuring BGP Virtual Private Wire Service \(VPWS\) on page 188](#)
- [Service Management Tasks on page 196](#)
 - Epipe:
 - [Modifying Epipe Service Parameters on page 197](#)
 - [Disabling an Epipe Service on page 197](#)
 - [Re-Enabling an Epipe Service on page 198](#)
 - [Deleting an Epipe Service on page 198](#)
 - Ipipe
 - [Modifying Ipipe Service Parameters on page 199](#)
 - [Disabling an Ipipe Service on page 200](#)
 - [Re-enabling an Ipipe Service on page 201](#)
 - [Deleting an Ipipe Service on page 201](#)

Basic Configurations

- [Creating an Epipe Service on page 162](#)
 - [Creating an Ipipe Service on page 173](#)
 - [Using Spoke SDP Control Words on page 177](#)
 - [Pseudowire Configuration Notes on page 180](#)
 - [Configuring Two VLL Paths Terminating on T-PE2 on page 182](#)
 - [Configuring VLL Resilience on page 185](#)
 - [Configuring VLL Resilience for a Switched Pseudowire Path on page 186](#)
-

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 162](#)
 - [Configuring Epipe SAP Parameters on page 163](#)
 - [Local Epipe SAPs on page 164](#)
 - [Distributed Epipe SAPs on page 166](#)
 - [Configuring Ingress and Egress SAP Parameters on page 169](#)

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

```

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

```

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 164](#)
- [Distributed Epipe SAPs on page 166](#)

CLI Syntax:

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

Local Epipe SAPs

Table 8: 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 170](#).

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

The following example 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 2/1/3:21 create
            ingress
                qos 555
                filter ip 1
            exit
            egress
                scheduler-policy "alpha"
                qos 627
            exit
        exit
    spoke-sdp 2:123 create
        ingress
            vc-label 6600
        exit
        egress
            vc-label 5500
        exit
    exit
    no shutdown
    exit
-----
A:ALA-1>config>service#
```

Configuring SDP Bindings

Figure 48 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 48: 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 166](#)).

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 2/1/3:21 create
            ingress
                qos 555
                filter ip 1
            exit
            egress
                scheduler-policy "alpha"
                qos 627
            exit
        exit
        spoke-sdp 2:123 create
            ingress
                vc-label 6600
            exit
            egress
                vc-label 5500
            exit
        exit
        no shutdown
    exit
...
-----
```

A:ALA-1>config>service#

A:ALA-2>config>service# info

```
-----
```

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#
```


Creating an Ipipe Service

Use the following CLI syntax to create an Ipipe service.

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

The following example displays an Ipipe configuration example:

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

Configuring Ipipe SAP Parameters

The following displays an Ipipe SAP configuration example:

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

The following displays a Frame Relay to Ethernet local Ipipe example:

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

The following displays the output:

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

The following displays a PPP to Ethernet local Ipipe example:

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

The following displays the output:

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

Configuring Ipipe SDP Bindings

The following displays an Ipipe SDP configuration example:

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

Using Spoke SDP Control Words

The control word command provides the option to add a control word as part of the packet encapsulation for PW types for which the control word is optional. These are Ethernet pseudowire (Epipe), ATM N:1 cell mode pseudowires (Apipe vc-types atm-vcc and atm-vpc) and VT pseudowire (Apipe vc-type atm-cell). The control word might be needed because when ECMP is enabled on the network, packets of a 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#
```

Same Fate Epipe VLANs Access Protection

The following displays a G.8031 Ethernet Tunnel for Epipe protection configuration example using same-fate SAPs for each Epipe access (two ethernet member ports 1/1/1 and 2/1/1 are used):

```
*A:7750_ALU>config>eth-tunnel 1
-----
description "Protection is APS"
protection-type 8031_lto1
ethernet
    mac 00:11:11:11:11:12
    encap-type dot1q
exit
ccm-hold-time down 5 up 10 // 50 ms down, 1 second up
path 1
    member 1/1/1
    control-tag 5 // primary control vlan 5
    precedence primary
    eth-cfm
        mep 2 domain 1 association 1
        ccm-enable
        control-mep
        no shutdown
    exit
exit
no shutdown
exit
path 2
    member 2/1/1
    control-tag 105 //secondary control vlan 105
    eth-cfm
        mep 2 domain 1 association 2
        ccm-enable
        control-mep
        no shutdown
    exit
exit
no shutdown
exit
no shutdown
-----
# Configure Ethernet tunnel SAPs
-----
*A:7750_ALU>config>service epipe 10 customer 5 create
    sap eth-tunnel-1 create // Uses control tags from the Ethernet tunnel port
        description "g8031-protected access ctl/data SAP for eth-tunnel 1"

    exit
    no shutdown
-----
*A:7750_ALU>config>service epipe 11 customer 5 create
    sap eth-tunnel-1:1 create
        description "g8031-protected access same-fate SAP for eth-tunnel 1"

    // must specify tags for each corresponding path in Ethernet tunnel port
    eth-tunnel path 1 tag 6
    eth-tunnel path 2 tag 106
exit
...
```

```
-----  
*A:7750_ALU>config>service epipe 10 customer 5 create  
  sap eth-tunnel-1:3 create  
    description "g8031-protected access same-fate SAP for eth-tunnel 1"  
    // must specify tags for each path for same-fate SAPs  
    eth-tunnel path 1 tag 10  
    eth-tunnel path 2 tag 110  
  exit  
  ...  
-----
```

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: SVCNMR #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# apipe service-id [customer customer-id] [vpn vpn-id] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-cell}] [vc-switching] description description-string spoke-sdp sdp-id:vc-id

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

CLI Syntax: config>service# fpipe service-id [customer customer-id][vpn vpn-id] [vc-type {fr-dlci}] [vc-switching] description description-string spoke-sdp sdp-id:vc-id

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

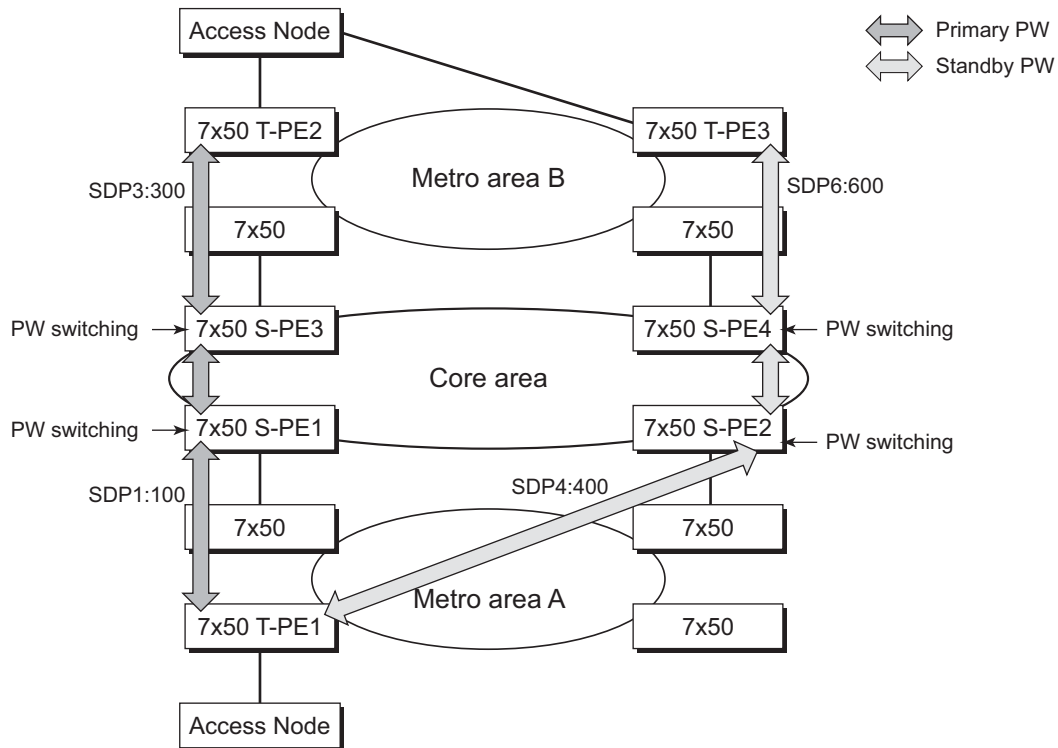
The following displays an example of the command usage to configure VLL pseudowire switching services:

Example:config>service# apipe 1 customer 1 vpn 1 vc-switching create
config>service>apipe\$ description "Default apipe description for service id 100"
config>service>apipe# spoke-sdp 3:1 create
config>service>apipe>spoke-sdp# exit
config>service>apipe# spoke-sdp 6:200 create
config>service>apipe>spoke-sdp# exit
config>service>apipe# no shutdown

The following example displays configurations for each service:

```
*A:ALA-48>config>service# info
-----
...
    apipe 100 customer 1 vpn 1 vc-switching create
        description "Default apipe description for service id 100"
        spoke-sdp 3:1 create
        exit
        spoke-sdp 6:200 create
        exit
        no shutdown
    exit
...
    epipe 107 customer 1 vpn 107 vc-switching create
        description "Default epipe description for service id 107"
        spoke-sdp 3:8 create
        exit
        spoke-sdp 6:207 create
        exit
        no shutdown
    exit
...
    ipipe 108 customer 1 vpn 108 vc-switching create
        description "Default ipipe description for service id 108"
        spoke-sdp 3:9 create
        exit
        spoke-sdp 6:208 create
        exit
        no shutdown
    exit
...
    fpipe 109 customer 1 vpn 109 vc-switching create
        description "Default fpipe description for service id 109"
        spoke-sdp 3:5 create
        exit
        spoke-sdp 6:209 create
        exit
        no shutdown
    exit
...
-----
*A:ALA-48>config>service#
```

Configuring Two VLL Paths Terminating on T-PE2



OSSG114

Figure 49: VLL Resilience with Pseudowire Redundancy and Switching

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 2/2/2:200 endpoint "x" create
        exit
        spoke-sdp 3:300 endpoint "y" create
            precedence primary
            revert-time 0
        exit
        spoke-sdp 6:600 endpoint "y" create
            precedence 0
        exit
        no shutdown
-----
*A:ALA-T-PE2>config>service>epipe#

```

S-PE1: 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
-----
...
        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 50 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.

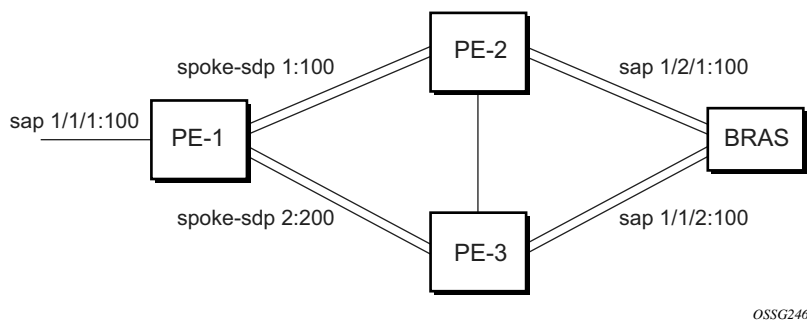


Figure 50: VLL Resilience

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

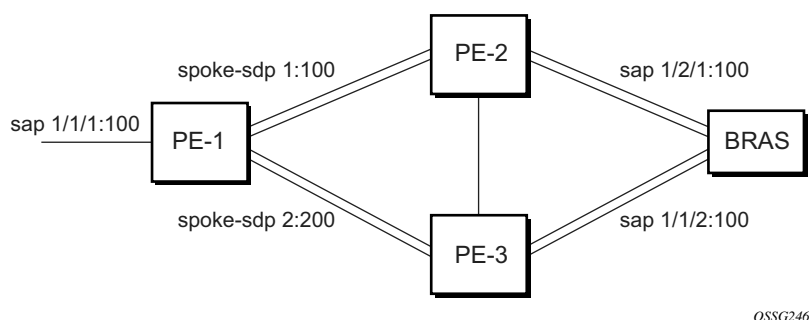


Figure 51: VLL Resilience with Pseudowire Switching

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 52 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 52: 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
```



```

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 53 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 PE and the remote PE.

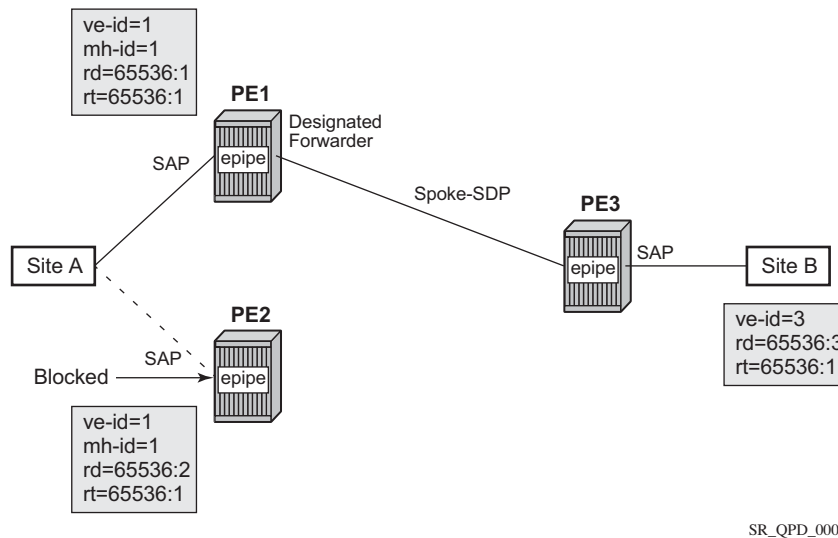


Figure 53: Example of Dual-Homed BGP VPWS with Single Pseudowire

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
```

```

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

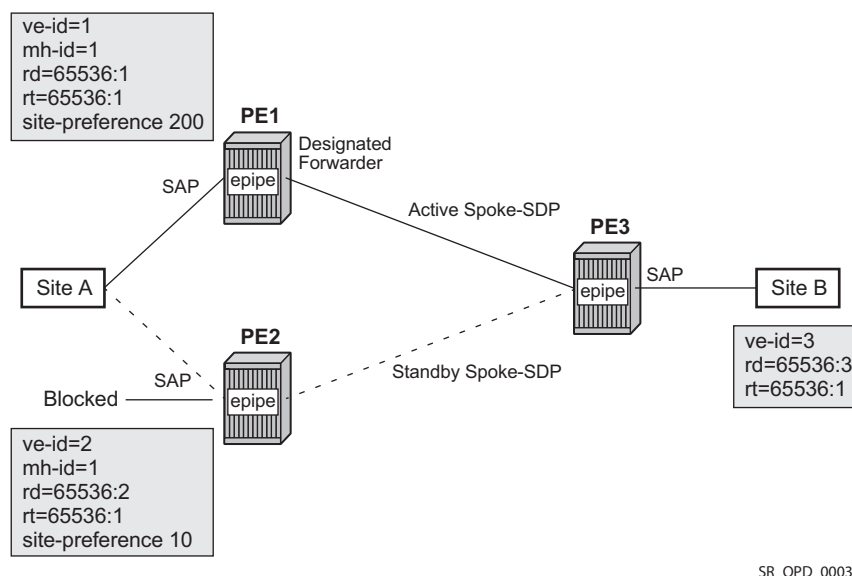
```

Configuring BGP Virtual Private Wire Service (VPWS)

```
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
    exit
exit
bgp-vpws
    ve-name PE3
        ve-id 3
    exit
    remote-ve-name PE1orPE2
        ve-id 1
    exit
    no shutdown
exit
sap 1/1/1:1 create
exit
no shutdown
exit
```

Active/Standby Pseudowire Example:

Figure 54 shows an example topology for a dual-homed BGP VPWS service used to create a virtual lease-line across an MPLS network between two sites, A and B. 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.



SR_QPD_0003

Figure 54: 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
```

Configuring BGP Virtual Private Wire Service (VPWS)

```
exit
bgp-vpws
  ve-name PE1
  ve-id 1
  exit
  remote-ve-name PE3
  ve-id 3
  exit
  no shutdown
exit
sap 1/1/1:1 create
exit
site "siteA" create
  site-id 1
  sap 1/1/1:1
  boot-timer 20
  site-activation-timer 5
  site-preference 200
  no shutdown
exit
no shutdown
exit
```

PE2:

```
pw-template 1 create
exit
epipe 1 customer 1 create
  bgp
    route-distinguisher 65536:2
    route-target export target:65536:1 import target:65536:1
    pw-template-binding 1
  exit
exit
bgp-vpws
  ve-name PE2
  ve-id 2
  exit
  remote-ve-name PE3
  ve-id 3
  exit
  no shutdown
exit
sap 1/1/1:1 create
exit
site "siteA" create
  site-id 1
  sap 1/1/1:1
  boot-timer 20
  site-activation-timer 5
  site-preference 10
  no shutdown
exit
no shutdown
exit
```

PE3:

```
pw-template 1 create
exit
```

```
epipe 1 customer 1 create
  bgp
    route-distinguisher 65536:3
    route-target export target:65536:1 import target:65536:1
    pw-template-binding 1
  exit
exit
bgp-vpws
  ve-name PE3
  ve-id 3
  exit
  remote-ve-name PE1
  ve-id 1
  exit
  remote-ve-name PE2
  ve-id 2
  exit
  no shutdown
exit
sap 1/1/1:1 create
exit
no shutdown
exit
```

Service Management Tasks

This section discusses the following Epipe service management tasks:

- [Modifying Epipe Service Parameters on page 197](#)
- [Disabling an Epipe Service on page 197](#)
- [Re-Enabling an Epipe Service on page 198](#)
- [Deleting an Epipe Service on page 198](#)

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 2/1/3:21
        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 2/1/3:21 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 2/1/3:21
config>service>epipe>sap# shutdown
config>service>epipe>sap# exit
config>service>epipe# no sap 2/1/3:21
config>service>epipe# spoke-sdp 2:6000
config>service>epipe>spoke-sdp# shutdown
config>service>epipe>spoke-sdp# exit
config>service>epipe# no spoke-sdp 2:6000
config>service>epipe# epipe 2
config>service>epipe# shutdown
config>service>epipe# exit
config>service# no epipe 2

Modifying Ipipe Service Parameters

The following example displays command usage to modify Ipipe parameters:

Example:

```
config>service# ipipe 202
config>service>ipipe# sap 1/1/2:444
config>service>ipipe>sap# shutdown
config>service>ipipe>sap# exit
config>service>ipipe# no sap 1/1/2:444
config>service>ipipe# sap 1/1/2:555 create
config>service>ipipe>sap$ description "eth_ipipe"
config>service>ipipe>sap$ ce-address 31.31.31.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# info
```

```
A:ALA-48>config>service# info
-----
...
    ipipe 202 customer 1 create
        sap 1/1/2:445 create
            description "eth_ipipe"
            ce-address 31.31.31.2
        exit
        sap 1/1/2:555 create
            description "eth_ipipe"
            ce-address 31.31.31.1
        exit
        no shutdown
    exit
...
-----
A:ALA-48>config>service#
```

Disabling an Ipipe Service

An Ipipe service can be shut down without deleting any service parameters.

CLI Syntax: config>service#
 ipipe service-id
 shutdown

Example: A:ALA-41>config>service# ipipe 202
 A:ALA-41>config>service>ipipe# shutdown

```
A:ALA-48>config>service# info
-----
...
    ipipe 202 customer 1 create
        shutdown
    sap 1/1/2:445 create
        description "eth_ipipe"
        ce-address 31.31.31.2
    exit
    sap 1/1/2:555 create
        description "eth_ipipe"
        ce-address 31.31.31.1
    exit
exit
...
-----
A:ALA-48>config>service#
```

Re-enabling an Ipipe Service

To re-enable an Ipipe service that was shut down.

CLI Syntax: config>service#
 ipipe service-id
 no shutdown

Example: A:ALA-41>config>service# ipipe 202
 A:ALA-41>config>service>ipipe# no shutdown

Deleting an Ipipe Service

An Ipipe service cannot be deleted until the SAP is shut down. If protocols and/or a spoke-SDP are defined, they must be shut down and removed from the configuration as well.

Use the following CLI syntax to delete an Ipipe service:

CLI Syntax: config>service#
 no ipipe service-id
 shutdown
 no sap sap-id
 shutdown
 no spoke-sdp [sdp-id:vc-id]
 shutdown

Example: config>service# ipipe 207
 config>service>ipipe# sap 1/1/2:449
 config>service>ipipe>sap# shutdown
 config>service>ipipe>sap# exit
 config>service>ipipe# no sap 1/1/2:449
 config>service>ipipe# spoke-sdp 16:516
 config>service>ipipe>spoke-sdp# shutdown
 config>service>ipipe>spoke-sdp# exit
 config>service>ipipe# no spoke-sdp 16:516
 config>service>ipipe# exit
 config>service# no ipipe 207
 config>service#

VLL Services Command Reference

Command Hierarchies

- [Apipe Service Configuration Commands on page 203](#)
- [Epipe Service Configuration Commands on page 208](#)
- [Ipipe Service Configuration Commands on page 219](#)

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
        — [no] collect-stats
        — 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
— frame-relay
— scheduling-class class-id
— no scheduling-class

```


- **ingress**
 - **policer-control-override** **[create]**
 - **no policer-control-override**
 - **max-rate** { *rate* | **max** }
 - **priority-mbs-thresholds**
 - **min-thresh-separation**
 - **[no] priority** *level*
 - **mbs-contribution** *size* [bytes | kilobytes]
 - **[no] policer-override**
 - **policer** *policer-id* **[create]**
 - **no policer** *policer-id*
 - **cbs** *size* [bytes | kilobytes]
 - **no cbs**
 - **mbs** *size* [bytes | kilobytes]
 - **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**
 - **[no] control-word**
 - **egress**
 - **qos** *network-policy-id* **port-redirect-group** *queue-group-name* [**instance** *instance-id*]
 - **no qos**
 - **vc-label** *ingress-vc-label*
 - **no vc-label** [*ingress-vc-label*]
 - **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*
- **[no] control-channel-status**
 - **[no] acknowledgment**
 - **refresh-timer** *value*
 - **no refresh-timer**
 - **request-timer** *timer1* **retry-timer** *timer2* [**timeout-multiplier** *multiplier*]
 - **no request-timer**
- **[no] control-word**
- **[no] pw-path-id**
 - **agi** *agi*
 - **no agi**
 - **saii-type2** *global-id:node-id:ac-id*
 - **no saii-type2**
 - **taii-type2** *global-id:node-id:ac-id*
 - **no taii-type2**

Epipe Service Configuration Commands

- [Epipe Global Commands on page 208](#)
- [Epipe SAP Configuration Commands on page 210](#)
- [Epipe Spoke SDP Configuration Commands on page 215](#)

Epipe Global Commands

```

config
— service
— [no] epipe service-id [customer customer-id] [test] [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
— bfd-enable
— bfd-template name
— no bfd-template
— shutdown
— route-distinguisher auto-rd
— no route-distinguisher
— route-distinguisher rd
— route-target { ext-community | { [export ext-community] [import ext-community] } }
— no route-target
— [no] bgp-vpws
— [no] remote-ve-name name
— ve-id value
— no ve-id
— [no] shutdown
— [no] ve-name name
— ve-id value
— no ve-id
— description description-string
— no description
— [no] endpoint endpoint-name
— active-hold-delay active-endpoint-delay
— no active-hold-delay
— description description-string
— no description
— revert-time [revert-time / infinite]
— no revert-time
— [no] standby-signaling-master
— [no] standby-signaling-slave
— load-balancing
— [no] per-service-hashing
— 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
    — [no] control-channel-status
        — [no] acknowledgment
        — refresh-timer value
        — no refresh-timer
        — request-timer timer1 retry-timer timer2 [timeout-multiplier
            multiplier]
        — no request-timer
    — [no] control-word
    — hash-label
    — no hash-label
    — [no] standby-signaling-slave
    — [no] pw-path-id
        — agi agi
        — no agi
        — saii-type2 global-id:node-id:ac-id
        — no saii-type2
        — taii-type2 global-id:node-id:ac-id
        — no taii-type2

```

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
— app-profile app-profile-name
— no app-profile
— [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]
— [no] hsmda-queue-override
— packet-byte-offset { add add-bytes | subtract sub-bytes }
— no packet-byte-offset
— queue queue-id
— no queue queue-id
— mbs size { [bytes | kilobytes] | default }
— no mbs
— rate pir-rate
— no rate
— slope-policy hsmda-slope-policy-name allowable
— no slope-policy
— wrr-weight weight
— no wrr-weight
— secondary-shaper secondary-shaper-name
— no secondary-shaper
— wrr-policy hsmda-wrr-policy-name
— no wrr-policy
— policer-control-override [create]
— no policer-control-override
— max-rate { rate | max }
— priority-mbs-thresholds
— min-thresh-separation
— [no] priority level
— mbs-contribution size [bytes | kilobytes]
— policer-control-policy policy-name
— no policer-control-policy
— [no] policer-override

```

```

— policer 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 {[weightweight] [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]
- **eth-tunnel**
 - **path** *path-index* **tag** *qtag* [*.qtag*]
 - **no path** *path-index*
- **ethernet**
 - [no] **llf**
- **frame-relay**
 - [no] **frf-12**
 - **ete-fragment-threshold** *threshold*
 - **no ete-fragment-threshold**
 - [no] **interleave**
 - **scheduling-class** *class-id*
 - **no scheduling-class**
- [no] **ignore-oper-down**
- **ingress**
 - **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
        — app-profile app-profile-name
        — no app-profile
        — bandwidth bandwidth
        — no bandwidth
        — [no] bfd-enable
        — bfd-template name
        — no bfd-template
        — [no] collect-stats
        — [no] control-word
        — [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] snelch-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*
- **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
    — ipipe service-id [customer customer-id] [vpn vpn-id] [vc-switching]
    — no ipipe service-id
      — ce-address-discovery [ipv6] [keep]
      — [no] ce-address-discovery
      — description description-string
      — no description
      — [no] endpoint endpoint-name
        — active-hold-delay active-endpoint-delay
        — no active-hold-delay
        — description description-string
        — no description
        — revert-time revert-time
        — no revert-time
      — eth-legacy-fault-notification
        — recovery-timer timer-value
        — [no] recovery-timer
        — [no] shutdown
      — sap sap-id [no-endpoint]
      — sap sap-id endpoint endpoint-name
      — [no] sap eth-tunnel-tunnel-id [:eth-tunnel-sap-id] [create]
      — no sap sap-id
        — accounting-policy acct-policy-id
        — no accounting-policy
        — ce-address ip-address
        — no ce-address
        — collect-stats
        — no collect-stats
        — description description-string
        — no description
        — 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
— [no] use-broadcast-mac
— service-mtu octets
— no service-mtu
— service-name service-name
— no service-name
— [no] shutdown

```

```

— spoke-sdp [sdp-id[:vc-id]] [no-endpoint]
— spoke-sdp [sdp-id[:vc-id]] endpoint endpoint-name [icb]
— no spoke-sdp sap-id
    — bandwidth bandwidth
    — no bandwidth
    — bfd-enable
    — no bfd-enable
    — bfd-template name
    — no bfd-template
    — ce-address ip-address
    — no ce-address
    — [no] control-word
    — egress
        — filter {ip ip-filter-id / ipv6 ipv6-filter-id}
        — no filter {ip ip-filter-id / ipv6 ipv6-filter-id}
        — qos network-policy-id port-redirect-group queue-group-name
          [instance instance-id]
        — no qos
        — [no] vc-label vc-label
    — 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
— [no] stack-capability-signaling

```

VLL Service Configuration Commands

- [Generic Commands on page 223](#)
- [VLL Global Commands on page 228](#)
- [VLL SAP Commands on page 244](#)
- [VLL Frame Relay Commands on page 311](#)
- [VLL SDP Commands on page 313](#)

Generic Commands

shutdown

| | |
|----------------------|---|
| Syntax | [no] shutdown |
| Context | 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 |
| 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> |

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

description

| | |
|--------------------|---|
| Syntax | description <i>description-string</i> no description |
| Context | 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 { <i>atm-vcc</i> <i>atm-sdu</i> <i>atm-vpc</i> <i>atm-cell</i> }] [vc-switching] [test] [create]] 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 7750 SR, 7450 ESS and 7710 SR 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> <p>test — Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs.</p> <p>test — Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs.</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> [test] [create] no epipe <i>service-id</i> | | | | | | | | | | | | |
| Context | config>service <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 ESS-Series 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.</p> <p>No MAC learning or filtering is provided on an Epipe.</p> <p>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 epipe services exist until they are explicitly created with this command.</p> <p>The no form of this command deletes the epipe service instance with the specified <i>service-id</i>. The service cannot be deleted until the service has been shutdown.</p> <p>Cpipe services are enabled on the ESS-Series in mixed mode.</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 ESS-Series on which this service is defined.</p> <table><tr><td>Values</td><td><i>service-id:</i></td><td>1 — 2147483648</td></tr><tr><td></td><td><i>svc-name:</i></td><td>64 characters maximum</td></tr></table> <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> <table><tr><td>Values</td><td>1 — 2147483647</td></tr></table> <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> <table><tr><td>Values</td><td>1 — 2147483647</td></tr><tr><td>Default</td><td>null (0)</td></tr></table> <p>vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.</p> | Values | <i>service-id:</i> | 1 — 2147483648 | | <i>svc-name:</i> | 64 characters maximum | Values | 1 — 2147483647 | Values | 1 — 2147483647 | Default | null (0) |
| Values | <i>service-id:</i> | 1 — 2147483648 | | | | | | | | | | | |
| | <i>svc-name:</i> | 64 characters maximum | | | | | | | | | | | |
| Values | 1 — 2147483647 | | | | | | | | | | | | |
| Values | 1 — 2147483647 | | | | | | | | | | | | |
| Default | null (0) | | | | | | | | | | | | |

test — Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs.

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

ipipe

| | |
|--------------------|---|
| Syntax | ipipe <i>service-id</i> [customer <i>customer-id</i>] [create] [vpn <i>vpn-id</i>] [vc-switching] no ipipe <i>service-id</i> |
| Context | config>service |
| Description | This command configures an IP-Pipe service. |
| 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 7750 SR, 7450 ESS and 7710 SR 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-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.</p> <p>create — Keyword used to create the Ipipe service instance. The create keyword requirement can be enabled/disabled in the environment>create context.</p> |

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 | <p>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.</p> <p>The pw-template-binding applies to BGP-VPWS when enabled in the Epipe.</p> <p>For BGP VPWS, the following additional rules govern the use of pseudowire-template:</p> <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. <p>The tools perform commands can be used to control the application of changes in pw-template for BGP-VPWS.</p> <p>The no form of the command removes the values from the configuration.</p> |
| 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. An RD value must be configured under BGP node. Alternatively, the auto-rd option allows the system to automatically generate an RD based on the bgp-auto-rd-range command configured at the service level. Format: Six bytes, other 2 bytes of type will be automatically generated. | | |
| Parameters | <i>ip-addr:comm-val</i> — Specifies the IP address. Values ip-addr a.b.c.d comm-val 0 — 65535 as-number: <i>as-number:ext-comm-val</i> — Specifies the AS number. Values as-number 1 — 65535 ext-comm-val 0 — 4294967295 auto-rd — The system will generate an RD for the service according to the IP address and range configured in the bgp-auto-rd-range command. <i>rd</i> — Specifies the route distinguisher. Values <rd> <ip-addr:comm-val> <2byte-asnumber:ext-comm-val> <4byte-asnumber:comm-val> ip-addr a.b.c.d comm-val [0..65535] 2byte-asnumber [1..65535] ext-comm-val [0..4294967295] 4byte-asnumber [0..4294967295] | | |

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 ext-comm 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 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 *value* — Specifies the site identifier.

Values 1 — 65535

site-preference

Syntax **site-preference** *preference-value*
no site-preference

Context config>service>epipe>site

Description This command defines the value to advertise in the VPLS preference field of the BGP VPWS and BGP Multi-homing NLRI extended community. This value can be changed without having to shutdown the site itself. The site-preference is only applicable to VPWS services.

When not configured, the default is zero, indicating that the VPLS preference is not in use.

Default no site-preference, value=0

Parameters *preference-value* — Specifies the preference value to advertise in the NLRI L2 extended community for this site.

Values 1 — 65535

Parameters **primary** — Sets the site-preference to 65535.

backup — Sets the site-preference to 1.

ce-address-discovery

Syntax [**no**] **ce-address-discovery** [**ipv6**] [**keep**]

Context config>service>ipipe

Description This command specifies whether the service will automatically discover the CE IP addresses.

When enabled, the addresses will be automatically discovered on SAPs that support address discovery, and on the spoke SDPs. When enabled, addresses configuration on the Ipipe SAP and spoke SDPs will not be allowed.

If disabled, CE IP addresses must be manually configured for the SAPs to become operationally up.

Default no ce-address-discovery

Parameters **ipv6** — The **ipv6** keyword enables IPv6 CE address discovery support on the Ipipe so that both IPv4 and IPv6 address discovery are supported. If the **ipv6** keyword is not included, then only IPv4 address discovery is supported and IPv6 packets are dropped. This feature requires IOM2 or better. For the 7450 ESS platforms, it requires mixed mode support and network chassis mode D to be enabled.

keep — The **keep** keyword is only applicable to eth-legacy-fault-notification. This option maintains the CE address discovered even when the SAP on which the address was learned fails. The ARP entry will not be maintained if the SAP is administratively shutdown, the clear service id x {arp | neighbor} is used to remove the ARP entry or the node reboots.

stack-capability-signaling

| | |
|--------------------|---|
| Syntax | [no] stack-capability-signaling |
| Context | config>service>ipipe |
| Description | <p>This command enables stack capability signaling in the initial label mapping message of the ipipe PW to indicate that IPv6 is supported.</p> <p>When enabled, the 7750 includes the stack capability TLV with the IPv6 stack bit set according to the ce-address-discovery ipv6 keyword, and also checks the value of the stack-capability TLV received from the far end.</p> <p>This command must be blocked if no ce-address-discovery is specified, or the ipv6 keyword is not included with the ce-address-discovery command.</p> <p>This command is only applicable to the ipipe service and must be blocked for all other services.</p> <p>This command has no effect if both SAPs on the ipipe service are local to the node.</p> <p>This feature requires IOM2 or better. For the 7450 ESS platforms, it requires mixed mode support and network chassis mode D to be enabled.</p> |
| Default | no stack-capability-signaling |

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. |

load-balancing

| | |
|--------------------|--|
| Syntax | load-balancing |
| Context | config>service>epipe> |
| Description | <p>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.</p> |
| Default | not applicable |

per-service-hashing

| | |
|--------------------|---|
| Syntax | [no] per-service-hashing |
| Context | config>service>epipe>load-balancing |
| Description | <p>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.</p> <p>The following algorithm describes the hash-key used for hashing when the new option is enabled:</p> <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> – 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 <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> – Unicast, BUM flooded without MMRP or with MMRP, IGMP snooped <p>The no form of this command implies the use of existing hashing options.</p> |
| Default | no per-service-hashing |

tunnel

| | | | | | |
|--------------------|--|--------------------|----------------|------------------|-----------------------|
| Syntax | tunnel service-id backbone-dest-mac ieee-address isid ISID no tunnel | | | | |
| Context | config>service>epipe>pbb | | | | |
| Description | This command configures a Provider Backbone Bridging (PBB) tunnel for Backbone VPLS (B-VPLS) service information. | | | | |
| Parameters | <p><i>service-id</i> — Specifies the B-VPLS service for the PBB tunnel associated with this service.</p> <p>Values</p> <table> <tr> <td><i>service-id:</i></td><td>1 — 2147483648</td></tr> <tr> <td><i>svc-name:</i></td><td>64 characters maximum</td></tr> </table> <p>backbone-dest-mac ieee-address — Specifies the backbone destination MAC-address for PBB packets.</p> <p>isid ISID — Specifies a 24 bit service instance identifier for the PBB tunnel associated with this service. As part of the PBB frames, it is used at the destination PE as a demultiplexor field.</p> <p>Values 0 — 16777215</p> | <i>service-id:</i> | 1 — 2147483648 | <i>svc-name:</i> | 64 characters maximum |
| <i>service-id:</i> | 1 — 2147483648 | | | | |
| <i>svc-name:</i> | 64 characters maximum | | | | |

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 config>service>ipipe>endpoint |
| Description | <p>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.</p> <p>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.</p> <p>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.</p> |
| 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 config>service>ipipe>endpoint |
| Description | This command configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP. |
| Parameters | <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.

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 to 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>ipipe
 config>service>epipe

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a 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
 config>service>ipipe

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 on which the SAP was created. If the required

payload is larger than the port MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port 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 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 Epipe 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

ipipe: 1500

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 |
| 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 <i>encap-value</i> |
| Context | config>connection-profile |
| Description | <p>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.</p> <p>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.</p> <p>The no form of this command deletes the member from the configuration..</p> |
| Default | none |
| Parameters | <i>encap-value</i> — 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 | 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> |
| Context | config>service>apipe config>service>ipipe config>service>epipe |
| 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.</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. 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> |

sap-id — Specifies the physical port identifier portion of the SAP. See [Common CLI Command Descriptions on page 1319](#) for command syntax.

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, 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>ipipe 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 <i>link-map-profile-id</i> 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 <i>group-name</i> 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** *agg-rate*
no agg-rate-limit

Context config>service>epipe>sap>ingress

Description 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.

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.

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.

If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.

The **no** form of the command removes the aggregate rate limit from the SAP or multi-service site.

Parameters *agg-rate* — Defines the rate, in kilobits-per-second, that the maximum aggregate rate the queues on the SAP or MSS can operate.

Values 1 — 40000000, max

agg-rate

Syntax [**no**] **agg-rate**

Context config>service>cpipe>sap>egress
 config>service>epipe>sap>egress
 config>service>ipipe>sap>egress

Description This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, and **queue-frame-based-accounting**.

rate

| | |
|--------------------|--|
| Syntax | rate {max rate} no rate |
| Context | config>service>cpipe>sap>egress>agg-rate config>service>epipe>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>cpipe>sap>egress>agg-rate config>service>epipe>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>cpipe>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 HSMDB Ethernet ports. |

policer-control-override

| | |
|--------------------|--|
| Syntax | policer-control-override [create] no policer-control-override |
| Context | config>service>epipe>sap>egress |
| Description | <p>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.</p> <p>The no form of the command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.</p> |

| | |
|-------------------|---|
| 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 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 <i>level</i> |
| 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><i>level</i> — 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 <i>size</i> [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 | <p>bytes — This keyword signifies that size is expressed in bytes.</p> <p>kilobytes — The optional kilobytes keyword signifies that size is expressed in kilobytes.</p> |

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. 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 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>epipe>sap>egress |
| 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>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 | <p>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.</p> <p>The no form of the command is used to restore the policer's mbs setting to the policy defined value.</p> |
| Default | no mbs |
| Parameters | <p>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.</p> <p>Values 0 – 16777216 or default</p> <p>kilobytes — When kilobytes is defined, the value given for size is interpreted as the policer's MBS value given in kilobytes.</p> |

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 | <p>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.</p> <p>The no packet-byte-offset command is used to restore the policer's packet-byte-offset setting to the policy defined value.</p> |
| Default | no packet-byte-offset |
| Parameters | <p>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.</p> <p>Values 1 — 31</p> <p>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.</p> <p>Values 1 — 64</p> |

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 | <p><i>pir-rate</i> — The pir-percent parameter is used to express the policer's PIR as a percentage of the policers's parent arbiter rate.</p> <p>Values Percentage ranging from 0.01 to 100.00. The default is 100.00.</p> <p><i>cir</i> <i>cir-rate</i> — Configures the administrative CIR specified by the user.</p> <p>Values 0 — 200000000, max</p> |

percent-rate

| | |
|--------------------|---|
| Syntax | percent-rate <i>pir-percent</i> [cir <i>cir-percent</i>] [no percent-rate] |
| Context | config>service>epipe>sap>egress>queue-override>queue |
| Description | <p>The percent-rate command within the SAP ingress and egress QoS policy enables supports for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate or of its parent scheduler's rate.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> |

- 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> <p>When collect-stats is enabled, the lack of counters causes the system to generate the following statistics:</p> |

| | |
|----------------|-----|
| 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 <i>ip-address</i> no ce-address |
| Context | config>service>ipipe>sap config>service>ipipe>spoke-sdp |
| Description | <p>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.</p> <p>This command specifies the IP address of the CE device associated with an Ipipe SAP. In the case of a SAP, it is the address of the CE device directly attached to the SAP. The address must be a host address (no subnet addresses are accepted) as there must be only one CE device attached to an Ipipe SAP. The CE address specified at one end of an Ipipe will be used in processing ARP messages at the other endpoint, as the router acts as a proxy for ARP messages.</p> |
| Parameters | <i>ip-address</i> — 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 |

| | |
|--------------------|---|
| 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 <i>customer-site-name</i> no multi-service-site |
| Context | config>service>ipipe>sap config>service>epipe>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>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> |

ring-node

| | |
|--------------------|--|
| Syntax | ring-node <i>ring-node-name</i> no ring-node |
| Context | config>service>epipe>sap |
| Description | <p>This command configures a multi-chassis ring-node for this SAP.</p> <p>The no form of the command removes the name from the configuration.</p> |
| Default | none |

tod-suite

| | |
|---------------|---|
| Syntax | tod-suite <i>tod-suite-name</i> 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 prefix-aasub-policy-id no transit-policy |
| Context | config>service>epipe>sap |
| | 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 |

use-broadcast-mac

| | |
|--------------------|--|
| Syntax | [no] use-broadcast-mac |
| Context | config>service>ipipe>sap |
| Description | This command enables the user of a of broadcast MAC on SAP. An Ipipe VLL service with the ce-address-discovery command enabled forwards unicast IP packets using the broadcast MAC address until the ARP cache is populated with a valid entry for the CE IP and MAC addresses. The no form of this command enables the user of a of broadcast MAC on SAP. |
| Default | no use-broadcast-mac |

mac

| | |
|--------------------|---|
| Syntax | [no] mac ieee-address |
| Context | config>service>ipipe>sap |
| Description | This command assigns a specific MAC address to an Ipipe SAP. The no form of this command returns the MAC address of the SAP to the default value. |

| | |
|-------------------|---|
| Default | The physical MAC address associated with the Ethernet interface where the SAP is configured. |
| Parameters | <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 |
| Description | <p>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.</p> <p>The no form of this command restores mac-refresh to the default value.</p> |
| 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 config>service>ipipe |
| Description | <p>This command creates the accounting policy context that can be applied to a SAP.</p> <p>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.</p> <p>A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the config>log context.</p> <p>The no form of this command removes the accounting policy association from the SAP, and the accounting policy reverts to the default.</p> |
| 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 |

app-profile

| | |
|--------------------|--|
| Syntax | app-profile <i>app-profile-name</i> no app-profile |
| Context | config>service>epipe>sap config>service>epipe>spoke-sdp |
| 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. |

bandwidth

| | |
|--------------------|---|
| Syntax | bandwidth <i>bandwidth</i> no bandwidth |
| Context | config>service>epipe>spoke-sdp config>service>ipipe>spoke-sdp |
| Description | <p>This command specifies the bandwidth to be used for VLL bandwidth accounting by the VLL CAC feature.</p> <p>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.</p> <p>If an LSP consists of a primary and many secondary standby LSPs, then the bandwidth used in the maximum SDP available bandwidth is that of the active path. Any change to and LSP active path bandwidth will update the maximum SDP available bandwidth. Note however that a change to any constituent LSP bandwidth due to re-signaling of the primary LSP path or the activation of a secondary path which causes overbooking of the maximum SDP available bandwidth causes a warning and a trap to be issued but no further action is taken. The activation of a bypass or detour LSP in the path of the primary LSP does not change the maximum SDP available bandwidth.</p> <p>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.</p> <p>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.</p> <p>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.</p> |

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 | bdf-template <i>name</i> 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 | <i>name</i> — 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 | <p>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:</p> <table> <tr> <td>0x00000001</td><td>Pseudowire Not Forwarding</td></tr> <tr> <td>0x00000002</td><td>Local Attachment Circuit (ingress) Receive Fault</td></tr> <tr> <td>0x00000004</td><td>Local Attachment Circuit (egress) Transmit Fault</td></tr> <tr> <td>0x00000008</td><td>Local PSN-facing PW (ingress) Receive Fault</td></tr> <tr> <td>0x00000010</td><td>Local PSN-facing PW (egress) Transmit Fault</td></tr> </table> <p>The transmit direction is unblocked when the following PW status code is received:</p> <table> <tr> <td>0x00000000</td><td>Pseudowire forwarding (clear all failures)</td></tr> </table> <p>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.</p> | 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 | 0x00000000 | Pseudowire forwarding (clear all failures) |
| 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 | | | | | | | | | | | | |
| 0x00000000 | Pseudowire forwarding (clear all failures) | | | | | | | | | | | | |
| Default | no block-on-peer-fault | | | | | | | | | | | | |

cflowd

| | |
|--------------------|---|
| Syntax | [no] cflowd |
| Context | config>service>epipe>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 L2 services, only ingress sampling is supported.</p> |

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 IOM 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

dist-cpu-protection

Syntax dist-cpu-protection *policy-name*
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 | <p>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.</p> <p>The Ethernet port must be configured for null encapsulation.</p> |

ETH-CFM Service Commands

eth-cfm

| | |
|--------------------|--|
| Syntax | eth-cfm |
| Context | config>service>epipe>spoke-sdp config>service>epipe config>service>epipe>sap config>service>ipipe>sap |
| Description | This command 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-meg-level configuration option has been included. Changing this parameter will evaluate the MEP for AIS triggers based on the new criteria. |
| Parameters | allDef — Keyword that includes any CCM defect condition to trigger AIS generation macRemErrXcon — Keyword that excludes RDI CCM Defect condition to trigger AIS generation. |

collect-lmm-stats

| | |
|--------------------|---|
| Syntax | collect-lmm-stats no collect-lmm-stats |
| Context | config>service>epipe>sap>eth-cfm config>service>epipe>spoke-sdp>eth-cfm config>service>vpls>sap>eth-cfm config>service>vpls>spoke-sdp>eth-cfm config>service>vpls>mesh-sdp>eth-cfm config>service>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 | <p>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.</p> <p>The no form of the command disables and deletes the counters for this SAP or MPLS SDP binding.</p> |
| 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 | <p>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.</p> |
| 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. <table> <tr> <td>Values</td><td>1 — 7</td></tr> <tr> <td>Default</td><td>1</td></tr> </table> | Values | 1 — 7 | Default | 1 |
| 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. <table> <tr> <td>Default</td><td>1</td></tr> </table> | Default | 1 |
| 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. <table> <tr> <td>Values</td><td>0 — 7</td></tr> <tr> <td>Default</td><td>1</td></tr> </table> | Values | 0 — 7 | Default | 1 |
| 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 { up 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 up | 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

up — Sends ETH-CFM messages towards the MAC relay entity.

ccm-enable

| | |
|--------------------|--|
| Syntax | [no] ccm-enable |
| Context | config>service>epipe>spoke-sdp>eth-cfm>mep config>service>epipe>sap>eth-cfm>mep config>service>ipipe>sap>eth-cfm>mep |
| Description | This command enables the generation of CCM messages. The no form of the command disables the generation of CCM messages. |

ccm-ltm-priority

| | |
|--------------------|---|
| Syntax | ccm-ltm-priority <i>priority</i> no ccm-ltm-priority |
| Context | config>service>epipe>spoke-sdp>eth-cfm>mep config>service>epipe>sap>eth-cfm>mep config>service>ipipe>sap>eth-cfm>mep |
| Description | This command specifies the priority value for CCMs and LTMs transmitted by the MEP. The no form of the command removes the priority value from the configuration. |
| Default | The highest priority on the bridge-port. |
| Parameters | <i>priority</i> — Specifies the priority of CCM and LTM messages. Values 0 — 7 |

ccm-padding-size

| | |
|--------------------|--|
| Syntax | ccm-padding-size <i>ccm-padding</i> no ccm-padding-size <i>ccm-padding</i> |
| Context | config>service>epipe>sap>eth-cfm>mep config>service>ipipe>sap>eth-cfm>mep config>service>epipe>sdp>eth-cfm>mep config>service>epipe>spoke-sdp>eth-cfm>mep config>service>vpls>sap>eth-cfm>mep config>service>vpls>spoke-sdp>eth-cfm>mep config>service>vpls>mesh-sdp>eth-cfm>mep config>service>vpls>sap>eth-cfm>mep config>service>vpls>spoke-sdp>eth-cfm>mep config>service>vpls>mesh-sdp>eth-cfm>mep config>service>ies>if>sap>eth-cfm>mep> config>service>ies>if>spoke-sdp>eth-cfm>mep config>service>ies>sub-if>grp-if>sap>eth-cfm>mep config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>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>eth-cfm>mep config>service>ipipe>sap>eth-cfm>mep |
| Description | For this test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands: oam eth-cfm eth-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [priority <i>priority</i>] [data-length <i>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>eth-cfm>mep>eth-test-enable config>service>ipipe>sap>eth-cfm>mep>eth-test-enable |
| Description | This command configures the test pattern for eth-test frames. The no form of the command removes the values from the configuration. |
| Default | all-zeros |
| Parameters | all-zeros — Specifies to use all zeros in the test pattern. all-ones — Specifies to use all ones in the test pattern. crc-enable — Generates a CRC checksum. |

fault-propagation-enable

| | |
|--------------------|---|
| Syntax | fault-propagation-enable {use-if-tlv suspend-ccm} no fault-propagation-enable |
| Context | config>service>epipe>sap>eth-cfm>mep config>service>epipe>spoke-sdp>eth-cfm>mep config>service>ipipe>sap>eth-cfm>mep |
| Description | This command configures the fault propagation for the MEP. |
| Parameters | use-if-tlv — Specifies to use the interface TLV. suspend-ccm — Specifies to suspend 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 |

```
cconfig>service>ipipe>sap>eth-cfm>mep
```

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>epipe>spoke-sdp>eth-cfm>mep config>service>epipe>sap>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 SDP). |
| Parameters | <i>mac-address</i> — 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 <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 |

mip

| | |
|--------------------|---|
| Syntax | mip [mac <i>mac-address</i>] primary-vlan-enable [vlan <i>vlan-id</i>] 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 | <p>mac — provides a method for manually configuring the MIP MAC.</p> <p><i>mac-address</i> — Specifies the MAC address of the MIP.</p> <p>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.</p> <p>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.</p> <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> |
| Default | no mip |

squelch-ingress-levels

| | |
|----------------|--|
| Syntax | squelch-ingress-levels [<i>md-level</i> [<i>md-level</i> ...]] no squelch-ingress-levels |
| Context | config>service>epipe>sap>eth-cfm config>service>epipe>spoke-sdp>eth-cfm config>service>vpls>sap>eth-cfm config>service>vpls>spoke-sdp>eth-cfm config>service>vpls>mesh-sdp>eth-cfm config>service>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 | <p>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.</p> <p>The no form of the command removes the silent discarding of previously matching ETH-CFM PDUs.</p> |
| Default | no squelch-ingress-levels |
| Parameters | <i>md-level</i> — Identifies the level. |
| | Values [0..7] |

tunnel-fault

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

Service Filter and QoS Policy Commands

egress

| | |
|--------------------|---|
| Syntax | egress |
| Context | config>service>epipe>spoke-sdp config>service>ipipe>sap config>service>epipe>sap |
| Description | This command enables the context to configure egress SAP parameters. If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. |

force-qinq-vc-forwarding

| | |
|--------------------|--|
| Syntax | [no] force-qinq-vc-forwarding |
| Context | config>service>epipe>spoke-sdp config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp config>service>pw-template |
| Description | <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
```

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>ipipe>sap
 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 [**ipv6** *ipv6-filter-id*]
filter [**mac** *mac-filter-id*]
no filter [**ip** *ip-filter-id*]
no filter [**ipv6** *ipv6-filter-id*]
no filter [**mac** *mac-filter-id*]

Context config>service>epipe>sap>egress
 config>service>epipe>sap>ingress
 config>service>epipe>spoke-sdp>egress
 config>service>epipe>spoke-sdp>ingress
 config>service>ipipe>spoke-sdp>egress
 config>service>ipipe>sap>ingress
 config>service>ipipe>sap>egress
 config>service>ipipe>spoke-sdp>ingress
 config>service>epipe>sap>egress
 config>service>epipe>sap>ingress

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP matching criteria. Only one filter can be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *filter-id* with an ingress or egress SAP. The *filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Note that IPv6 filters are not supported on a Layer 2 SAP that is configured with QoS MAC criteria. Also, MAC filters are not supported on a Layer 2 SAP that is configured with QoS IPv6 criteria.

| | |
|----------------------|--|
| Special Cases | Epipe — Both MAC and IP filters are supported on an Epipe service SAP. |
| Parameters | <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> |

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 | This command configures the RX/TX cookie for L2TPv3 spoke-SDPs for EPipe services. The RX cookie must match the configured TX cookie on a far-end node, while the TX cookie must match the configured RX cookie on a far-end node. If a mismatch is detected between the configured (far-end |

binding cookie) to what is received by the local IP address of the SDP a flag is set and must be manually cleared by an operator.

The purpose of the cookie is to provide validation against misconfiguration of service endpoints, and to ensure that the right service egress is being used.

One egress cookie and up to two ingress cookies may be configured per spoke-SDP binding. One or two cookies can be configured for matching ingress packets from the far-end node, in order to support cookie rollover without dropping packets. When a cookie is not configured, SR-OS assumes a value of 00:00:00:00:00:00:00:00.

A cookie is not mandatory. An operator may delete an egress cookie or either or both ingress cookies.

| | |
|-------------------|---|
| Default | no cookie1 cookie2 |
| Parameters | <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 HSMDA 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

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 20 bytes may be added to the packet and up to 43 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As mentioned 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.

The packet-byte-offset value can be overridden for the HSMDA queue at the SAP or subscriber profile level.

The **no** form of the command removes any accounting size changes to packets handled by the queue. The command does not effect overrides that may exist on SAPs or subscriber profiles associated with the queue.

Parameters

add *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>hsmda-queue-over

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

Description

This command, within the QoS policy hsmda-queue context, is a container for the configuration parameters controlling the behavior of an HSMDA queue. Unlike the standard QoS policy queue command, this command is not used to actually create or dynamically assign the queue to the object which the policy is applied. The queue identified by queue-id always exists on the SAP or subscriber context whether the command is executed or not. In the case of HSMDA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDA queue group to the object (both ingress and egress).

Best-Effort, Expedited and Auto-Expedite Queue Behavior Based on Queue-ID

With standard service queues, the scheduling behavior relative to other queues is based on two items, the queues Best-Effort or Expedited nature and the dynamic rate of the queue relative to the defined CIR. HSMDA queues are handled differently. The create time auto-expedite and explicit expedite and best-effort qualifiers have been eliminated and instead the scheduling behavior is based solely on the queues identifier. Queues with a queue-id equal to 1 are placed in scheduling class 1. Queues with queue-id 2 are placed in scheduling class 2. And so on up to scheduling class 8. Each scheduling class is either mapped directly to a strict scheduling priority level based on the class ID, or the class may be placed into a weighted scheduling class group providing byte fair weighted round robin scheduling between the members of the group. Two weighted groups are supported and each may contain up to three consecutive scheduling classes. The weighed group assumes its highest member class is inherent strict scheduling level for scheduling purposes. Strict priority level 8 has the highest priority while strict level 1 has the lowest. When grouping of scheduling classes is defined, some of the strict levels will not be in use.

Single Type of HSMDA Queues

Another difference between HSMDA queues and standard service queues is the lack of Multipoint queues. At ingress, an HSMDA SAP or subscriber does not require Multipoint queues since all forwarding types (broadcast, multicast, unicast and unknown) forward to a single destination 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 <i>pir-rate</i> 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 | <i>pir-rate</i> — 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>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 | <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>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 | <i>hsmda-wrr-policy-name</i> — Specifies the existing HSMDA WRR policy name to associate to the queue. |

slope-policy

| | |
|---------------|--|
| Syntax | slope-policy <i>hsmda-slope-policy-name</i> no slope-policy |
|---------------|--|

| | |
|--------------------|---|
| Context | config>service>epipe>sap>egress>hsmda-queue-over config>service>ipipe>sap>egress>hsmda-queue-over |
| Description | <p>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.</p> <p>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.</p> <p>Default HSMDA Slope Policy</p> <p>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.</p> <p>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.</p> |
| Parameters | <i>hsmda-slope-policy-name</i> — 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. |

filter

| | |
|----------------|--|
| Syntax | filter [ip <i>ip-filter-id</i>] filter [ipv6 <i>ipv6-filter-id</i>] no filter [ip <i>ip-filter-id</i>] [ipv6 <i>ipv6-filter-id</i>] no filter [ip <i>ip-filter-id</i>] |
| Context | config>service>cpipe>spoke-sdp>egress config>service>cpipe>spoke-sdp>ingress config>service>ipipe>spoke-sdp>egress config>service>ipipe>sap>ingress config>service>ipipe>sap>egress |

```
config>service>ipipe>spoke-sdp>ingress
```

| | |
|--------------------|--|
| Description | <p>This command associates a 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 matching criteria. Only one filter can be applied to a SAP at a time.</p> <p>The filter command is used to associate a filter policy with a specified <i>ip-filter-id</i> with an ingress or egress SAP. The <i>ip-filter-id</i> 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>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.</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> |
| 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> |

qos

| | |
|--------------------|---|
| Syntax | <p>qos <i>policy-id</i> [shared-queuing] [fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i>]</p> <p>no qos</p> |
| Context | <pre>config>service>apipe>sap>ingress config>service>fpipe>sap>ingress config>service>ipipe>sap>ingress config>service>epipe>sap>ingress</pre> |
| 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> |

By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

| | |
|-------------------|---|
| Default | none |
| Parameters | <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-queueing — 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> <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.</p> <p>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.</p> <p>The multipoint keyword can be entered in the command line on a pre-existing multipoint queue to edit queue-id parameters.</p> <p>Default Present (the queue is created as non-multipoint).</p> <p>Values Multipoint or not present.</p> <p>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 <i>instance-id</i> are mandatory parameters when executing the command.</p> <p><i>queue-group-name</i> — 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 <i>queue-group-name</i> must correspond to a valid ingress forwarding plane queue group, created under <i>config>card>fp>ingress>access</i>.</p> <p>instance <i>instance-id</i> — Specifies the instance of the named queue group on the IOM/IMM/XMA ingress forwarding plane.</p> |

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> <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 egress policy ID to associate with SAP on egress. The policy ID must already exist.</p> <p>Values 1 — 65535</p> <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>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>.</p> <p>instance instance-id — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.</p> <p>Values 1 — 40960</p> <p>Default 1</p> |

queue-override

| | |
|----------------|---|
| Syntax | [no] queue-override |
| Context | <pre> config>service>apipe>sap>egress config>service>apipe>sap>ingress config>service>cpipe>sap>egress config>service>cpipe>sap>ingress config>service>fpipesap>egress config>service>fpipesap>ingress config>service>ipipesap>egress config>service>ipipesap>ingress config>service>epipesap>egress config>service>epipesap>ingress </pre> |

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>ipipe>sap>egress>queue-override>queue
 config>service>ipipe>sap>ingress>queue-override>queue
 config>service>epipe>sap>egress>queue-override>queue
 config>service>epipe>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default no adaptation-rule

Parameters **pir** — The **pir** parameter defines the constraints enforced when adapting the PIR rate defined within the **queue queue-id rate** command. The **pir** parameter requires a qualifier that defines the

constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values

max — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

| | |
|--------------------|---|
| Syntax | avg-frame-overhead percent no avg-frame-overhead |
| Context | config>service>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> |

- **Offered-load** — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- **Frame encapsulation overhead** — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 0.1 or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50 x 20 or 1000 octets.

- **Frame based offered-load** — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- **Packet to frame factor** — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be $1000 / 10000$ or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- **Frame based CIR** — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500×1.1 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 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.

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 | <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 size [byte kilobyte]} no burst-limit |
| Context | config>service>ipipe>sap>egress>queue-override>queue config>service>ipipe>sap>ingress>queue-override>queue config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue |
| Description | <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 sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates 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 burst-limit default within the QoS policies or queue group templates. When specified within a queue-override 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 burst-limit default is executed, the queue is returned to the system default value.</p> <p><i>size</i> — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in Kilobytes. If the value is intended to be interpreted in bytes, the 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>ipipe>sap>egress>queue-override>queue config>service>ipipe>sap>ingress>queue-override>queue config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue |
| Description | <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> <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.</p> <p>The no form of this command returns the CBS size to the default value.</p> |
| Default | no cbs |
| Parameters | <p><i>size-in-kbytes</i> — 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).</p> <p>Values 0 — 131072, default</p> |

high-prio-only

| | |
|--------------------|--|
| Syntax | high-prio-only <i>percent</i> no high-prio-only |
| Context | config>service>ipipe>sap>egress>queue-override>queue config>service>ipipe>sap>ingress>queue-override>queue config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue |
| Description | <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> |

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>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue
```

Description This command can be used to override specific attributes of the specified queue's 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 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.

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 | <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. |
| Values | 0 — 1073741824 or default in expressed bytes or kilobytes |

monitor-depth

| | |
|--------------------|--|
| Syntax | [no] monitor-depth |
| Context | config>service>apipe>sap>egress>queue-override>queue config>service>apipe>sap>ingress>queue-override>queue config>service>cpipe>sap>egress>queue-override>queue config>service>cpipe>sap>ingress>queue-override>queue config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue config>service>fpipe>sap>egress>queue-override>queue config>service>fpipe>sap>ingress>queue-override>queue config>service>ipipe>sap>egress>queue-override>queue config>service>ipipe>sap>ingress>queue-override>queue 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 <i>weight</i>] [cir-weight <i>cir-weight</i>]} no parent |
| Context | config>service>epipe>sap>egress>queue-override>queue config>service>epipe>sap>ingress>queue-override>queue |
| Description | This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the weight or level parameters define how this queue contends with the other children for the parent's bandwidth. Checks are not performed to see if a <i>scheduler-name</i> exists when the parent command is defined on the queue. Scheduler names are configured in the config>qos>scheduler-policy>tier <i>level</i> context. Multiple schedulers can exist with the <i>scheduler-name</i> 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 <i>scheduler-name</i> is dependent on a scheduler policy containing the <i>scheduler-name</i> being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the <i>scheduler-name</i> 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** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the

intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile 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 | <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> <p><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> <p>Values 0 — 3200000000, max, sum kbps</p> <p>Default 0</p> |

scheduler-override

| | |
|--------------------|--|
| Syntax | [no] scheduler-override |
| Context | <pre>config>service>ipipe>sap>egress config>service>ipipe>sap>ingress config>service>epipe>sap>egress config>service>epipe>sap>ingress</pre> |
| 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 <i>scheduler-name</i> |
| Context | config>service>ipipe>sap>egress>sched-override config>service>ipipe>sap>ingress>sched-override config>service>epipe>sap>egress>sched-override config>service>epipe>sap>ingress>sched-override |
| Description | <p>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.</p> <p>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 <i>scheduler-name</i> 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).</p> <p>If the <i>scheduler-name</i> 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.</p> <p>If the <i>scheduler-name</i> 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:</p> <ol style="list-style-type: none"> 1. The maximum number of schedulers has not been configured. 2. The provided <i>scheduler-name</i> 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). <p>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.</p> <p>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.</p> |
| Parameters | <p><i>scheduler-name</i> — The name of the scheduler.</p> <p>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.</p> <p>Default None. Each scheduler must be explicitly created.</p> <p><i>create</i> — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given <i>scheduler-name</i>. 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</p> |

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 <i>weight</i>] [cir-weight <i>cir-weight</i>] 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 | <p>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.</p> <p>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.</p> <p>The no form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.</p> |
| 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</p> |

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 <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 maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.</p> <p>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.</p> <p>The no form of this command returns all queues created with this <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters.</p> |
| Parameters | <p><i>pir-rate</i> — 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.</p> <p>To calculate the actual PIR rate, the rate described by the queue's rate is multiplied by the <i>pir-rate</i>.</p> <p>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.</p> |

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 — 100000000, **max**, **sum**

Default sum

scheduler-policy

| | |
|--------------------|---|
| Syntax | scheduler-policy <i>scheduler-policy-name</i> no scheduler-policy |
| Context | config>service>ipipe>sap>ingress config>service>ipipe>sap>egress config>service>epipe>sap>ingress config>service>epipe>sap>egress |
| Description | <p>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.</p> <p>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.</p> <p><i>scheduler-policy-name</i> — The <i>scheduler-policy-name</i> parameter applies an existing scheduler policy that was created in the config>qos>scheduler-policy <i>scheduler-policy-name</i> 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.</p> |

vlan-translation

| | |
|--------------------|--|
| Syntax | vlan-translation { <i>vlan-id</i> copy-outer } no vlan-translation |
| Context | config>service>epipe>sap>ingress |
| Description | <p>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.</p> <p>The no version of the command sets the default value and no action will be taken.</p> |
| Default | Per default, the preserved VLAN values will not be overwritten. |
| Parameters | <p><i>vlan-id</i> — Specifies that the preserved vlan-id will be overwritten with this value.</p> <p>Values 0 — 4094</p> <p>outer-copy — Keyword specifies to use the outer VLAN ID.</p> |

match-qinq-dot1p

| | |
|--------------------|---|
| Syntax | match-qinq-dot1p { top bottom } no match-qinq-dot1p de |
| Context | config>service>ipipe>sap>ingress config>service>epipe>sap>ingress |
| Description | <p>This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.</p> <p>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.</p> <p>The setting also applies to classification based on the DE indicator bit.</p> <p>The no form of this command reverts the dot1p and de bits matching to the default tag.</p> <p>By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 9 defines the default behavior for Dot1P evaluation.</p> |

Table 9: 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 |

Table 9: Default QinQ and TopQ SAP Dot1P Evaluation (Continued)

| Port / SAP Type | Existing Packet Tags | PBits Used for Match |
|-----------------|-------------------------------|----------------------|
| 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 |
| 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 10: 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 |
| 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 Frame Relay Commands

frame-relay

| | |
|--------------------|---|
| Syntax | frame-relay |
| Context | config>service>ipipe>sap config>service>epipe>sap |
| Description | This command enables the context to configure Frame Relay parameters. |

frf-12

| | |
|--------------------|--|
| Syntax | [no] frf-12 |
| Context | config>service>ipipe>sap>frame-relay config>service>epipe>sap>frame-relay |
| Description | This command enables the use of FRF12 headers. The no form of the command disables the use of FRF12 headers. |

ete-fragment-threshold

| | | | | | |
|--------------------|---|---------------|-----------|----------------|---|
| Syntax | ete-fragment-threshold <i>threshold</i> no ete-fragment-threshold | | | | |
| Context | config>service>ipipe>sap>frame-relay>frf-12 config>service>epipe>sap>frame-relay>frf-12 | | | | |
| Description | This command specifies the maximum length of a fragment to be transmitted. The no form of the command reverts to the default. | | | | |
| Parameters | <i>threshold</i> — The maximum length of a fragment to be transmitted. <table> <tr> <td>Values</td><td>128 — 512</td></tr> <tr> <td>Default</td><td>0</td></tr> </table> | Values | 128 — 512 | Default | 0 |
| Values | 128 — 512 | | | | |
| Default | 0 | | | | |

interleave

| | |
|----------------|--|
| Syntax | [no] interleave |
| Context | config>service>epipe>sap>frame-relay>frf.12 config>service>ipipe>sap>frame-relay>frf.12 |

| | |
|--------------------|--|
| Description | <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> | | | | |
| Description | <pre>config>service>ipipe>sap>frame-relay</pre> <pre>config>service>epipe>sap>frame-relay</pre> | | | | |
| Description | This command specifies the scheduling class to use for this SAP. | | | | |
| Parameters | <p><i>class-id</i> — Specifies the scheduling class to use for this sap.</p> <table> <tr> <td>Values</td><td>0 — 3</td></tr> <tr> <td>Default</td><td>0</td></tr> </table> | Values | 0 — 3 | Default | 0 |
| Values | 0 — 3 | | | | |
| Default | 0 | | | | |

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 **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding.

vlan — Defines the VC type as VLAN. The top VLAN tag, if a VLAN tag is present, is stripped from traffic received on the pseudowire, and a vlan-tag is inserted when forwarding into the pseudowire. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.
The VLAN VC-type requires at least one dot1Q tag within each encapsulated Ethernet packet transmitted to the far end.

Note: The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

no endpoint — Removes the association of a spoke SDP with an explicit endpoint name.

endpoint *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 config>service>ipipe |
| 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> |

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a service. If the **sdp sdp-id** is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end SR/ESS devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default No *sdp-id* is bound to a service.

Parameters *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 to 17407 for existing SDPs.

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

no endpoint — Adds or removes a spoke SDP association.

endpoint endpoint-name — 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 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.

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 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 ESS-Series 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 ESS-Series 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

Context config>service>cpipe>spoke-sdp


```
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp
```

Description

The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe).

The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.

The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an “Illegal C-bit” status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes. The control word must be enabled to allow MPLS-TP OAM to be used on a static spoke-sdp in a apipe, epipe and cpipe service.

pw-path-id

Syntax [no] pw-path-id

Context config>service>epipe>spoke-sdp
 config>service>cpipe>spoke-sdp
 config>service>apipe>spoke-sdp
 config>service>vpls>spoke-sdp
 config>service>ies>interface>spoke-sdp
 config>service>vprn>interface>spoke-sdp

Description

This command enables the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID.

For an IES or VPRN spoke-sdp, the pw-path-id is only valid for ethernet spoke-sdps.

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

- The system is using network chassis mode D
- SDP signaling is off
- control-word is enabled (control-word is disabled by default)
- the service type is epipe, vpls, cpipe, apipe, or IES/VPRN interface
- mate SDP signaling is off for vc-switched services

The **no** form of the command deletes the PW path ID.

Default no pw-path-id

agi

| | |
|--------------------|--|
| Syntax | agi <i>agi</i> no agi |
| Context | config>service>epipe>spoke-sdp>pw-path-id config>service>cpipe>spoke-sdp>pw-path-id config>service>apipe>spoke-sdp>pw-path-id config>service>vpls>spoke-sdp>pw-path-id config>service>ies>interface>>spoke-sdp>pw-path-id config>service>vprn>interface>>spoke-sdp>pw-path-id |
| Description | This command configures the attachment group identifier for an MPLS-TP PW. |
| Parameters | <i>agi</i> — Specifies the attachment group identifier. Values 0 — 4294967295 |

saii-type2

| | |
|--------------------|--|
| Syntax | saii-type2 <i>global-id:node-id:ac-id</i> no saii-type2 |
| Context | config>service>epipe>spoke-sdp>pw-path-id config>service>cpipe>spoke-sdp>pw-path-id config>service>apipe>spoke-sdp>pw-path-id config>service>vpls>spoke-sdp>pw-path-id config>service>ies>interface>>spoke-sdp>pw-path-id config>service>vprn>interface>>spoke-sdp>pw-path-id |
| Description | This command configures the source individual attachment identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the taii-type2 of the mate spoke-sdp. |
| Parameters | <i>global-id</i> — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values 0 — 4294967295 <i>node-id</i> — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values a.b.c.d or 0 — 4294967295 <i>ac-id</i> — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value. Values 1 — 4294967295 |

taii-type2

| | |
|---------------|--|
| Syntax | taii-type2 <i>global-id:node-id:ac-id</i> no taii-type2 |
|---------------|--|

| | |
|--------------------|---|
| Context | <pre>config>service>epipe>spoke-sdp>pw-path-id config>service>cpipe>spoke-sdp>pw-path-id config>service>apipe>spoke-sdp>pw-path-id config>service>vpls>spoke-sdp>pw-path-id config>service>ies>interface>>spoke-sdp>pw-path-id config>service>vprn>interface>>spoke-sdp>pw-path-id</pre> |
| Description | This command configures the target individual attachment identifier (TAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the saii-type2 of the mate spoke-sdp. |
| Parameters | <p><i>global-id</i> — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p>Values 0 — 4294967295</p> <p><i>node-id</i> — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p>Values a.b.c.d or 0 — 4294967295</p> <p><i>ac-id</i> — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.</p> <p>Values 1 — 4294967295</p> |

control-channel-status

| | |
|--------------------|---|
| Syntax | [no] control-channel-status |
| Context | <pre>config>service>cpipe>spoke-sdp config>service>epipe>spoke-sdp config>service>apipe>spoke-sdp config>service>ies>interface>>spoke-sdp config>service>vpls>spoke-sdp config>service>vprn>interface>>spoke-sdp</pre> |
| Description | <p>This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF.</p> <p>A control-channel-status no shutdown is allowed only if all of the following are true:</p> <ul style="list-style-type: none"> • The system is using network chassis mode D • SDP signaling is off. • The control-word is enabled (the control-word is disabled by default) • The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN • Mate SDP signaling is off (in vc-switched services) • The pw-path-id is configured for this spoke-SDP. <p>The no form of this command removes control channel status signaling from a spoke-SDP. It can only be removed if control channel status is shut down.</p> |
| Default | no control-channel-status |

acknowledgment

| | |
|--------------------|--|
| Syntax | [no] acknowledgment |
| Context | config>service>cpipe>spoke-sdp>control-channel-status config>service>epipe>spoke-sdp>control-channel-status config>service>apipe>spoke-sdp>control-channel-status config>service>vpls>spoke-sdp>control-channel-status config>service>ies>interface>>spoke-sdp>control-channel-status config>service>vprn>interface>>spoke-sdp>control-channel-status |
| Description | This command enables the acknowledgement of control channel status messages. By default, no acknowledgement packets are sent. |

refresh-timer

| | |
|--------------------|--|
| Syntax | refresh-timer <i>value</i> no refresh-timer |
| Context | config>service>epipe>spoke-sdp>control-channel-status config>service>cpipe>spoke-sdp>control-channel-status config>service>apipe>spoke-sdp>control-channel-status config>service>vpls>spoke-sdp>control-channel-status config>service>ies>interface>>spoke-sdp>control-channel-status config>service>vprn>interface>>spoke-sdp>control-channel-status |
| Description | This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent. |
| Default | no refresh-timer |
| Parameters | <i>value</i> — Specifies the refresh timer value. |
| Values | 10 — 65535 seconds |
| Default | 0 (off) |

request-timer

| | |
|----------------|--|
| Syntax | request-timer <i>timer1</i> retry-timer <i>timer2</i> timeout-multiplier <i>multiplier</i> no request-timer |
| Context | config>service>cpipe>spoke-sdp>control-channel-status config>service>epipe>spoke-sdp>control-channel-status config>service>apipe>spoke-sdp>control-channel-status config>service>vpls>spoke-sdp>control-channel-status config>service>ies>interface>>spoke-sdp>control-channel-status config>service>vprn>interface>>spoke-sdp>control-channel-status |

| | |
|--------------------|--|
| Description | This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value. |
| Parameters | <p><i>timer1</i> — Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the “request” bit set, are sent.</p> <p>Values 10 — 65535 seconds</p> <p>retry-timer <i>timer2</i> — specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.</p> <p>Values 0, 3 — 60 seconds</p> <p>timeout-multiplier <i>multiplier</i> — If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it will assume the pseudowire is down. This parameter is optional.</p> <p>Values 3 — 20 seconds</p> |

control-word

| | |
|--------------------|--|
| Syntax | [no] control-word |
| Context | config>service>ies>interface>spoke-sdp config>service>vprn>interface>spoke-sdp |
| Description | <p>This command enables/disables the PW control word on spoke-sdps terminated on an IES or VPRN interface. The control word must be enabled to allow MPLS-TP OAM on the spoke-sdp</p> <p>It is only valid for MPLS-TP spoke-sdps when used with IES and VPRN services.</p> |
| Default | no control-word |

egress

| | |
|--------------------|--|
| Syntax | egress |
| Context | config>service>cpipe>spoke-sdp config>service>ipipe>spoke-sdp |
| Description | This command configures the egress SDP context. |

hash-label

| | |
|----------------|--|
| Syntax | hash-label [signal-capability] no hash label |
| Context | config>service>ipipe>spoke-sdp config>service>epipe>spoke-sdp |

Description This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to 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 or for a VPRN service using the autobind mode with the `gre` option..

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to 1 to indicate that.

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.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note however that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets that are generated in CPM and forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the Hash Label is set to a value of 0.

The **no** form of this command disables the use of the hash label.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh-sdp, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The ESS-Series 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 ESS-Series 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 <i>network-policy-id</i> port-redirect-group <i>queue-group-name</i> [instance <i>instance-id</i>] no qos |
| Context | config>service>apipe>spoke-sdp>egress config>service>cpipe>spoke-sdp>egress config>service>epipe>spoke-sdp>egress config>service>fpipes>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 | <p>This command is used to redirect PW packets to an egress port queue-group for the purpose of shaping.</p> <p>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.</p> <p>Operationally, the provisioning model consists of the following steps:</p> <ol style="list-style-type: none"> 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 | <i>network-policy-id</i> — Specifies the network policy identification. The value uniquely identifies the policy on the system. |
| Values | 1—65535 |
| port-redirect-group | <i>queue-group-name</i> — Specifies the name of the queue group template up to 32 characters in length. |
| instance | <i>instance-id</i> — Specifies the optional identification of a specific instance of the queue-group. |
| Values | 1—40960 |

qos

| | |
|--------------------|---|
| Syntax | qos <i>network-policy-id</i> fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i> 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 | <p>This command is used to redirect PW packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.</p> <p>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.</p> <p>Operationally, the provisioning model in the case of the ingress PW 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 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/VRPN 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 VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to the default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface the pseudowire packet is received on.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

| | |
|--------------------------|---|
| Parameters | <i>network-policy-id</i> — Specifies the network policy identification on the system. |
| Values | 1—65535 |
| fp-redirect-group | <i>queue-group-name</i> — Specifies the name of the queue group template up to 32 characters in length. |
| instance | <i>instance-id</i> — Specifies the identification of a specific instance of the queue-group. |
| Values | 1—16384 |

vc-label

| | |
|--------------------|--|
| Syntax | [no] vc-label <i>vc-label</i> |
| Context | config>service>cpipe>spoke-sdp>egress config>service>ipipe>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 | [no] vc-label <i>vc-label</i> |
| Context | config>service>cpipe>spoke-sdp>ingress config>service>ipipe>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 |

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>ipipe>spoke-sdp config>service>epipe>spoke-sdp |
| Description | This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic. The no form of the command returns the precedence value to the default. |
| Default | 4 |
| Parameters | <i>precedence-value</i> — 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 <i>vc-label</i> |
| Context | config>service>cpipe>spoke-sdp>egress config>service>epipe>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 | [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 | <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 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.</p> <p>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.</p> <p>The no form of this command disables the command</p> |
| 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 | <p>This command binds a service to an existing Service Distribution Point (SDP), using a dynamic MS-PW.</p> <p>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>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.</p> |

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>aii-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>his 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 TAI associated with that spoke-SDP. It allows a single-sided provisioning model where an incoming label mapping message with a TAI that matches the SAI 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> <p>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.</p> <p>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 SAI and TAI of the spoke-SDP. If the SAI has the greater prefix</p> |

value, then the 7x50 will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAIL 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 <i>name</i> 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 | <i>path-name</i> — The name of the explicit path to be used, as configured under config>service>pw-routing. |

precedence

| | |
|--------------------|--|
| Syntax | precedence <i>prec-value</i> 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 | <i>precedence-value</i> — 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”. 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

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 l2-ap SAP. This can only be changed when the l2-ap is shutdown. The no form of the command removes the template, the SAP will use default parameters. |
| Default | None |
| Parameters | <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 <i>filter-id</i> 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 <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters. Values 1 — 65535 |

ipv6

| | |
|--------------------|--|
| Syntax | ipv6 <i>filter-id</i> 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 <i>ipv6-filter-id</i> — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters. Values 1 — 65535 |

mac

| | |
|--------------------|---|
| Syntax | mac <i>filter-id</i> 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. |

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
    — egress-label start-label [end-label]
    — id service-id
      — all
      — authentication
      — base
      — bgp-vpws
      — endpoint [endpoint-name]
      — labels
      — retailers
      — 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]
      — wholesalers
    — ingress-label start-label [end-label]
    — 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] [vppls] [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
      — arp
      — host-tracking [sap sap-id] [host ip-address]

```

- **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
```

VLL Show Commands

egress-label

| | |
|--------------------|--|
| Syntax | egress-label <i>egress-label1</i> [<i>egress-label2</i>] |
| Context | show>service |
| Description | <p>This command displays services using the range of egress labels. 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> |
| Output | Show Service Egress Command Output — The following table describes show service egress label output fields. |

Table 11: Show Service Egress Label Output Fields

| Label | Description |
|--------------------------|--|
| Svc Id | The ID that identifies a service. |
| Sdp Id | The ID that identifies an SDP. |
| Type | Indicates whether the SDP binding is spoke or mesh. |
| I. Lbl | The VC label used by the far-end device to send packets to this device in this service by the SDP. |
| E. Lbl | The VC label used by this device to send packets to the far-end device in this service by the SDP. |
| Number of bindings found | The total number of SDP bindings that exist within the specified egress label range. |

Sample Output

```

*A:ALA-12# show service egress-label 0 10000
=====
Martini Service Labels
=====
Svc Id      Sdp Id      Type I.Lbl      E.Lbl
-----
1           10:1        Mesh 0          0
1           20:1        Mesh 0          0
1           30:1        Mesh 0          0
1           100:1       Mesh 0          0
...
1           107:1       Mesh 0          0
1           108:1       Mesh 0          0
1           300:1       Mesh 0          0
1           301:1       Mesh 0          0
1           302:1       Mesh 0          0
1           400:1       Mesh 0          0
1           500:2       Spok 131070     2001
1           501:1       Mesh 131069     2000
100         300:100     Spok 0          0
200         301:200     Spok 0          0
300         302:300     Spok 0          0
400         400:400     Spok 0          0
-----
Number of Bindings Found : 23
=====
*A:ALA-12#

```

ingress-label

| | |
|--------------------|---|
| Syntax | ingress-label <i>start-label</i> [<i>end-label</i>] |
| Context | show>service |
| Description | <p>This command displays services using the range of ingress labels. If only the mandatory <i>start-label</i> parameter is specified, only services using the specified label are displayed.</p> <p>If both <i>start-label</i> and <i>end-label</i> parameters are specified, the services using the range of labels X where <i>start-label</i> <= X <= <i>end-label</i> are displayed.</p> <p>Use the show router vprn-service-id ldp bindings command to display dynamic labels.</p> |
| Parameters | <p><i>start-label</i> — The starting ingress label value for which to display services using the label range. If only <i>start-label</i> is specified, services only using <i>start-label</i> are displayed.</p> <p>Values 0, 2048 — 131071</p> <p><i>end-label</i> — The ending ingress label value for which to display services using the label range.</p> <p>Default The <i>start-label</i> value.</p> <p>Values 2049 — 131071</p> |
| 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 a spoke or a mesh. |
| I.Lbl | The ingress label used by the far-end device to send packets to this device in this service by the SDP. |
| E.Lbl | The egress label used by this device to send packets to the far-end device in this service by the SDP. |
| Number of Bindings Found | The number of SDP bindings within the label range specified. |

Sample Output

```
*A:ALA-12# show service ingress-label 0
=====
Martini Service Labels
=====
Svc Id      Sdp Id      Type I.Lbl      E.Lbl
-----
1           10:1        Mesh 0         0
1           20:1        Mesh 0         0
1           30:1        Mesh 0         0
1           50:1        Mesh 0         0
1           100:1       Mesh 0         0
1           101:1       Mesh 0         0
1           102:1       Mesh 0         0
1           103:1       Mesh 0         0
1           104:1       Mesh 0         0
1           105:1       Mesh 0         0
1           106:1       Mesh 0         0
1           107:1       Mesh 0         0
1           108:1       Mesh 0         0
1           300:1       Mesh 0         0
1           301:1       Mesh 0         0
1           302:1       Mesh 0         0
1           400:1       Mesh 0         0
100         300:100     Spok 0         0
200         301:200     Spok 0         0
300         302:300     Spok 0         0
400         400:400     Spok 0         0
-----
Number of Bindings Found : 21
-----
*A:ALA-12#
```

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*
sap-using authentication-policy *policy-name*

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 or egress 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 1319 for command syntax.
interface *ip-address* — The IP address of the interface for which to display matching SAPs.
Values 1.0.0.0 — 223.255.255.255
interface *ip-int-name* — The IP interface name for which to display matching SAPs.
authentication-policy *policy name* — Specifies an existing authentication policy.
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. |
| Sap MTU | The SAP MTU value. |
| Ing. QoS | The SAP ingress QoS policy number specified on the ingress SAP. |
| Ing Fltr | The MAC or IP filter policy ID applied to the ingress SAP. |
| Egr. QoS | The SAP egress QoS policy number specified on the egress SAP. |
| Egr. Fltr | The MAC or IP filter policy ID applied to the egress SAP. |

| Label | Description (Continued) |
|-------|--------------------------------------|
| 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
=====
```

A:ALA-42# show service sap-using

```
=====
Service Access Points
=====
PortId                      SvcId      Ing.  Ing.  Egr.  Egr.  Adm  Opr
                        QoS    Fltr  QoS   Fltr
-----
1/1/2:100                   1          1    none  1     none  Up   Down
1/1/4:0                     2          1    none  1     none  Up   Down
1/2/9:0                     6          1    none  1     none  Up   Down
1/1/3:0                     88         1    none  1     none  Up   Down
1/1/5:0                     88         1    none  1     none  Up   Down
1/1/2:0                     218        1    none  1     none  Up   Down
1/1/9:0                     700        1    none  1     none  Up   Down
1/1/9:10                    1000       1    none  1     none  Up   Down
1/1/12:11                   1000       1    none  1     none  Up   Down
1/1/13:10.20                1000       1    none  1     none  Up   Down
-----
Number of SAPs : 10
=====
```

A:ALA-42#

*A:ALA-48# show service sap-using sap 1/1/21:0

```
=====
Service Access Points Using Port 1/1/21:0
=====
PortId                      SvcId      Ing.  Ing.  Egr.  Egr.  Anti  Adm  Opr
                        QoS    Fltr  QoS   Fltr  Spoof
-----
```

Show, Clear, Debug Commands

```
1/1/21:0          1          1      none      1      none      none      Up      Down
-----
Number of SAPs : 1
=====
*A:ALA-48#
```

sdp

| | |
|--------------------|--|
| Syntax | sdp <i>[[sdp-id[:vc-id] far-end ip-address] [detail]]</i> sdp <i>sdp-id:vc-id mrp</i> |
| 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 | <i>sdp-id</i> — Specifies the SDP ID. Values 1 — 17407 <i>vc-id</i> — 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. Octets.       : 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

```

```

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** **epipe** [**ies**] [**vpls**][**mirror**] [**ipipe**] [**sdp** *sdp-id*] [**customer** *customer-id*] [**b-vpls**] [**i-vpls**] [**m-vpls**]
- 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. |

| Label | Description |
|------------------|---|
| CustomerID | The ID of the customer who owns this service. |
| Last Mgmt Change | The date and time of the most recent management-initiated change to this service. |

Sample Output

```
*A:ALA-12# show service service-using customer 10
=====
Services
=====
ServiceId      Type      Adm      Opr      CustomerId      Last Mgmt Change
-----
1              VPLS      Up       Up       10              09/05/2006 13:24:15
100           IES       Up       Up       10              09/05/2006 13:24:15
300           Epipe     Up       Up       10              09/05/2006 13:24:15
-----
Matching Services : 3
=====
*A:ALA-12#

*A:ALA-12# show service service-using
=====
Services
=====
ServiceId      Type      Adm      Opr      CustomerId      Last Mgmt Change
-----
1              uVPLS     Up       Up       1              10/26/2006 15:44:57
2              Epipe     Up       Down    1              10/26/2006 15:44:57
10             mVPLS     Down    Down    1              10/26/2006 15:44:57
11             mVPLS     Down    Down    1              10/26/2006 15:44:57
100            mVPLS     Up       Up       1              10/26/2006 15:44:57
101            mVPLS     Up       Up       1              10/26/2006 15:44:57
102            mVPLS     Up       Up       1              10/26/2006 15:44:57
999            uVPLS     Down    Down    1              10/26/2006 16:14:33
-----
Matching Services : 8
-----
*A:ALA-12#
*A:ALA-12# show service service-using epipe
=====
Services [epipe]
=====
ServiceId      Type      Adm      Opr      CustomerId      Last Mgmt Change
-----
6              Epipe     Up       Up       6              06/22/2006 23:05:58
7              Epipe     Up       Up       6              06/22/2006 23:05:58
8              Epipe     Up       Up       3              06/22/2006 23:05:58
103            Epipe     Up       Up       6              06/22/2006 23:05:58
-----
Matching Services : 4
=====
*A:ALA-12#
```

spoke-sdp-fec-using

| | |
|--------------------|--|
| Syntax | spoke-sdp-fec-using [spoke-sdp-fec-id <i>spoke-sdp-fec-id</i>] [saii-type2 <i>global-id:prefix:ac-id</i>] [taii-type2 <i>global-id:prefix:ac-id</i>] [path <i>name</i>] |
| 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
=====
SvcId      SpokeSdpFec  Oper-SdpBind      SAII-Type2
Path                                     TAII-Type2
-----
1          1            17407:4294967245  3:10.20.1.3:1
n/a                                     6:10.20.1.6:1
2          2            17407:4294967247  3:10.20.1.3:2
n/a                                     6:10.20.1.6:2
3          3            17407:4294967248  3:10.20.1.3:3
n/a                                     6:10.20.1.6:3
4          4            17407:4294967249  3:10.20.1.3:4
n/a                                     6:10.20.1.6:4
5          5            17407:4294967250  3:10.20.1.3:5
n/a                                     6:10.20.1.6:5
6          6            17407:4294967251  3:10.20.1.3:6
n/a                                     6:10.20.1.6:6
7          7            17407:4294967252  3:10.20.1.3:7
n/a                                     6:10.20.1.6:7
8          8            17407:4294967253  3:10.20.1.3:8
n/a                                     6:10.20.1.6:8
9          9            17407:4294967254  3:10.20.1.3:9
n/a                                     6:10.20.1.6:9
10         10            17407:4294967255  3:10.20.1.3:10
n/a                                     6:10.20.1.6:10
-----
Entries found: 10
=====
```

vccv-bfd

| | |
|--------------------|--|
| Syntax | vccv-bfd session [detail] [sdp sdp-id[:vc-id]] vccv-bfd session [detail] |
| Context | show>service>id |
| Description | <p>This command shows whether VCCV BFD is configured for a particular service and information about the VCCV session state.</p> <p>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.</p> <p>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.</p> |
| Parameters | <p><i>service-id</i> — The identification number of the specific service.</p> <p>Values service-id: 1 — 214748364</p> <p><i>sdp-id</i> — Specifies the SDP ID.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP.</p> <p>Values 1 — 4294967295</p> |

Sample Output

```
*A:Dut-C# show service id 1000 vccv-bfd session
=====
BFD Session
=====
Interface/Lsp Name          State          Tx Intvl  Rx Intvl  Multipl
Remote Address/Info         Protocols      Tx Pkts   Rx Pkts   Type
LAG port/sdp-id:vc-id      LAG ID/SvcId
-----
N/A                          Up (3)         1000      1000      3
N/A                          vccv          152       151       central
100:100                      1000
-----
No. of BFD sessions: 1
=====

A:bksim1611>show>service# id 1000 sdp 25:105 vccv-bfd session detail
=====
BFD Session
=====
Svc-id : 1000
sdp-id:vc-id   : 25:105  sdpBindInstanceId : 100
Admin State    : Up      Oper State       : Up (3)
Up Time        : 0d 00:00:27  Up Transitions   : 1
Down Time      : None      Down Transitions  : 0
Version Mismatch : 0
```

Show, Clear, Debug Commands

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 <i>service-id</i> { 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 | <i>service-id</i> — 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:bksim1611>config>service>ipipe# show service id 1009 all
=====
Service Detailed Information
=====
Service Id       : 1009                Vpn Id           : 0
Service Type     : Ipipe
Name             : (Not Specified)
Description      : (Not Specified)
Customer Id      : 1
Last Status Change: 09/15/2010 13:06:46 Last Mgmt Change : 09/15/2010 13:06:02
Admin State      : Up                  Oper State       : Up
MTU              : 1500
Vc Switching     : False
SAP Count        : 1                  SDP Bind Count   : 1
CE IPv4 Discovery : Enabled            CE IPv6 Discovery : Enabled
-----
Service Destination Points(SDPs)
-----
Sdp Id 5:1009   -(5.5.5.5)
-----
Description     : (Not Specified)
SDP Id          : 5:1009                Type             : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp : (Not Specified)
```

Show, Clear, Debug Commands

```

VC Type           : Ipipe           VC Tag           : 0
Admin Path MTU    : 0               Oper Path MTU    : 1568
Far End          : 5.5.5.5         Delivery         : MPLS
Tunnel Far End   : n/a
Hash Label       : Disabled

Admin State       : Up              Oper State       : Up
Acct. Pol        : None            Collect Stats    : Disabled
Ingress Label    : 131048          Egress Label    : 131053
Ingr Mac Fltr-Id : n/a            Egr Mac Fltr-Id : n/a
Ingr IP Fltr-Id  : n/a            Egr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a           Egr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred  Oper ControlWord : False
Admin BW(Kbps)   : 0              Oper BW(Kbps)    : 0
Last Status Change : 09/15/2010 13:06:46
Last Mgmt Change  : 09/15/2010 13:06:02
Signaling         : TLDP
Endpoint         : N/A            Precedence       : 4
PW Status Sig     : Enabled
Class Fwding State : Down
Flags            : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel

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

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

```

----- RSVP/Static LSPs

----- Associated LSP LIST :

```

Lsp Name          : to-bksim180-1
Admin State       : Up              Oper State         : Up
Time Since Last Tr*: 16h07m44s

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

```

----- Class-based forwarding :

```

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

```

=====

FC Mapping Table

```

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

```

IPIPE Service Destination Point specifics

Configured CE IP Addr : n/a Peer CE IP Addr : 0.0.0.0

Peer IPv6 Capability : No
Peer IPv6 LL Addr : FE80::2009:2009:2
Peer IPv6 Global Addr : 3FFE:1200:2009:2009:9:9:9:8

Number of SDPs : 1

Service Access Points

SAP 1/7/3:1009

| | | | |
|---------------------|-----------------------|--------------------|------------|
| Service Id | : 1009 | | |
| SAP | : 1/7/3:1009 | Encap | : q-tag |
| Description | : (Not Specified) | | |
| Admin State | : Up | Oper State | : Up |
| Flags | : None | | |
| Multi Svc Site | : None | | |
| Last Status Change | : 09/15/2010 13:06:21 | | |
| Last Mgmt Change | : 09/15/2010 13:06:02 | | |
| Sub Type | : regular | | |
| Dot1Q Ethertype | : 0x8100 | QinQ Ethertype | : 0x8100 |
| Split Horizon Group | : (Not Specified) | | |
| Admin MTU | : 1518 | Oper MTU | : 1518 |
| Ingr IP Fltr-Id | : n/a | Egr IP Fltr-Id | : n/a |
| Ingr Mac Fltr-Id | : n/a | Egr Mac Fltr-Id | : n/a |
| Ingr IPv6 Fltr-Id | : n/a | Egr IPv6 Fltr-Id | : n/a |
| tod-suite | : None | qinq-pbit-marking | : both |
| Ing Agg Rate Limit | : max | Egr Agg Rate Limit | : max |
| Endpoint | : N/A | | |
| Q Frame-Based Acct | : Disabled | | |
| Acct. Pol | : None | Collect Stats | : Disabled |
| Oper Group | : (none) | Monitor Oper Grp | : (none) |

ETH-CFM SAP specifics

Tunnel Faults : n/a CFM Hold-Timer : n/a

Ipipe SAP Configuration Information

| | | | |
|------------------|---------------------|--------------------|-------------|
| Configured CE IP | : n/a | Discovered CE IP | : 209.1.1.1 |
| SAP MAC Address | : ac:55:01:07:00:03 | Mac Refresh Inter* | : 14400 |

Ipipe SAP IPv4 ARP Entry Info

| | |
|-----------|---------------------------|
| 209.1.1.1 | 00:11:22:33:44:55 dynamic |
|-----------|---------------------------|

Ipipe SAP IPv6 Neighbor Entry Info

FE80::2009:2009:1

00:11:22:33:44:55 dynamic

QOS

```
Ingress qos-policy : 1
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 : 1
Multipoint shared : Disabled
```

Sap Statistics

Last Cleared Time : N/A

| | Packets | Octets |
|-------------------------|---------|--------|
| Forwarding Engine Stats | | |
| Dropped | : 2 | 172 |
| Off. HiPrio | : 0 | 0 |
| Off. LowPrio | : 17 | 1978 |
| Off. Uncolor | : 0 | 0 |

Queueing Stats(Ingress QoS Policy 1)

| | | |
|--------------|------|------|
| Dro. HiPrio | : 0 | 0 |
| Dro. LowPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 17 | 1978 |

Queueing Stats(Egress QoS Policy 1)

| | | |
|--------------|------|------|
| Dro. InProf | : 0 | 0 |
| Dro. OutProf | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 15 | 1790 |

Sap per Queue stats

| | Packets | Octets |
|--------------------------------------|---------|--------|
| Ingress Queue 1 (Unicast) (Priority) | | |
| Off. HiPrio | : 0 | 0 |
| Off. LoPrio | : 17 | 1978 |
| Dro. HiPrio | : 0 | 0 |
| Dro. LoPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 17 | 1978 |

Egress Queue 1

| | | |
|--------------|------|------|
| For. InProf | : 0 | 0 |
| For. OutProf | : 15 | 1790 |
| Dro. InProf | : 0 | 0 |
| Dro. OutProf | : 0 | 0 |

Service Endpoints

No Endpoints found.


```

VPLS Sites
=====
Site              Site-Id  Dest              Mesh-SDP  Admin  Oper  Fwdr
-----
No Matching Entries
=====
show service id x all
-----
SAP 1/1/4:500
-----
Service Id       : 500
SAP              : 1/1/4:500          Encap              : q-tag
Description      : (Not Specified)
Admin State      : Up                Oper State         : Down
Flags            : PortOperDown
Multi Svc Site   : None
Last Status Change : 09/19/2013 11:43:04
Last Mgmt Change  : 09/19/2013 11:43:05
Sub Type         : regular
Dot1Q Ethertype  : 0x8100            QinQ Ethertype     : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU        : 1518              Oper MTU           : 1518
Ingr IP Fltr-Id  : n/a              Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id : n/a              Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a            Egr IPv6 Fltr-Id  : n/a
tod-suite        : None              qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint         : N/A
Q Frame-Based Acct : Disabled
Vlan-translation : None

Acct. Pol        : None              Collect Stats      : Disabled

Application Profile: None
Transit Policy    : None

Oper Group        : (none)           Monitor Oper Grp   : (none)
Host Lockout Plcy : n/a
Ignore Oper Down  : Disabled
Lag Link Map Prof : (none)
Cflowd           : Disabled
-----
ETH-CFM SAP specifics
-----
Tunnel Faults    : n/a              AIS                : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels    : 0 1 2 3 4 5 6 7
-----
QOS
-----
Ingress qos-policy : 1              Egress qos-policy : 1
.
.
.
-----
Service Destination Points(SDPs)
-----

```

Show, Clear, Debug Commands

```

-----
Sdp Id 1:2  -(1.1.1.1)
-----
Description      : (Not Specified)
SDP Id           : 1:2                               Type           : Spoke
Spoke Descr      : (Not Specified)
Split Horiz Grp  : (Not Specified)
VC Type          : Ether                               VC Tag          : n/a
Admin Path MTU   : 0                                   Oper Path MTU   : 0
Delivery         : GRE
Far End          : 1.1.1.1
Tunnel Far End   : n/a                               LSP Types       : n/a
Hash Label       : Disabled                           Hash Lbl Sig Cap : Disabled
Oper Hash Label  : Disabled

Admin State      : Up                                Oper State      : Down
Acct. Pol        : None                              Collect Stats   : Disabled
Ingress Label    : 0                                Egress Label    : 0
Ingr Mac Fltr-Id : n/a                              Egr Mac Fltr-Id : n/a
Ingr IP Fltr-Id  : n/a                              Egr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a                             Egr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred                    Oper ControlWord : False
Last Status Change : 09/11/2013 20:02:40             Signaling       : TLDP
Last Mgmt Change  : 09/15/2013 13:56:56             Force Vlan-Vc   : Disabled
Endpoint         : N/A                               Precedence      : 4
PW Status Sig     : Enabled
Class Fwding State : Down
Flags             : SdpOperDown
                   NoIngVCLabel NoEgrVCLabel
                   PathMTUTooSmall

Time to RetryReset : never                           Retries Left    : 3
Mac Move           : Blockable                        Blockable Level : Tertiary
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None
Transit Policy     : None
Max Nbr of MAC Addr: No Limit                         Total MAC Addr  : 0
Learned MAC Addr   : 0                               Static MAC Addr  : 0
OAM MAC Addr       : 0                               DHCP MAC Addr   : 0
Host MAC Addr      : 0                               Intf MAC Addr   : 0
SPB MAC Addr       : 0                               Cond MAC Addr   : 0

MAC Learning       : Enabled                           Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                             Block On Mesh Fail: False
Oper Group         : (none)                             Monitor Oper Grp  : (none)
Rest Prot Src Mac   : Disabled
Auto Learn Mac Prot: Disabled                         RestProtSrcMacAct : Disable

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

```

```

ETH-CFM SDP-Bind specifics
-----
V-MEP Filtering      : Disabled

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

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

Squelch Levels        : 0 1 2 3 4 5 6 7

```

authentication

| | |
|--------------------|--|
| Syntax | authentication |
| Context | show>service>id |
| Description | This command enables the context to display subscriber authentication information. |

statistics

| | |
|--------------------|---|
| Syntax | statistics [<i>policy name</i>] [<i>sap sap-id</i>] |
| Context | show>service>id>authentication |
| Description | This command displays session authentication statistics for this service. |
| Parameters | <p>policy name — Specifies the subscriber authentication policy statistics to display.</p> <p>sap sap-id — Specifies the SAP ID statistics to display. See Common CLI Command Descriptions on page 1319 for command syntax.</p> |

Sample Output

```

*A:ALA-1# show service id 11 authentication statistics
=====
Authentication statistics
=====
Interface / SAP                Authentication  Authentication
                               Successful         Failed
-----
vpls-11-90.1.0.254             1582           3
-----
Number of entries: 1
=====
*A:ALA-1#

```

all

| | |
|--------------------|--|
| Syntax | all |
| Context | show>service>id |
| Description | This command displays detailed information for all aspects of the service. |
| Output | 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. |
| SDP Id | The SDP identifier. |
| Description | Generic information about the service. |
| Customer Id | The customer identifier. |
| Last Mgmt Change | The date and time of the most recent management-initiated change. |
| Endpoint | Specifies the name of the service endpoint. |
| Flags | Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode. |
| SAP Count | The number of SAPs specified for this service. |
| SDP Bind Count | The number of SDPs bound to this service. |
| Split Horizon Group specifics | |
| Split Horizon Group | Name of the split horizon group for this VPLS. |
| Description | Description of the split horizon group. |
| Last Changed | The date and time of the most recent management-initiated change to this split horizon group. |
| Service Destination Points (SDPs) | |
| SDP Id | The SDP identifier. |

| Label | Description (Continued) |
|-----------------------------|--|
| Type | Indicates whether this Service SDP binding is a spoke or a mesh. |
| Admin Path MTU | The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |
| Oper Path MTU | The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |
| Delivery | Specifies the type of delivery used by the SDP: GRE or MPLS. |
| Admin State | The administrative state of this SDP. |
| Oper State | The operational state of this SDP. |
| Jitter Buffer (packets) | Indicates the jitter buffer length in number of packet buffers. |
| Playout Threshold (packets) | Indicates the playout buffer packets threshold in number of packet buffers. |
| Playout Threshold (packets) | Indicates the current packet depth of the jitter buffer. |
| 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:SR12# show service id 10 all
=====Service
Detailed Information
=====Service Id
: 10                Vpn Id                : 0
Service Type       : Apipe                VLL Type         : ATMCell
Name               : (Not Specified)
Description        : (Not Specified)
Customer Id        : 2
Last Status Change: 10/07/2010 05:03:47 Last Mgmt Change : 10/07/2010 05:03:51
Admin State        : Up                    Oper State        : Down
MTU                : 1508                  Signaling Override: ATMVCC
Vc Switching       : False
```

Show, Clear, Debug Commands

```

SAP Count          : 1                      SDP Bind Count      : 1

..... (No change to SDP description)

----- SAP 2/1/
4:cp.10
-----Service Id
: 10
SAP                : 2/1/4:cp.10            Encap                : atm
Description         : (Not Specified)
Admin State        : Up                      Oper State             : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 10/16/2010 06:58:41
Last Mgmt Change   : 10/16/2010 06:58:41
Sub Type           : regular
Split Horizon Group: (Not Specified)

Admin MTU          : 1524                    Oper MTU                : 1524
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id         : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id        : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id       : n/a
tod-suite          : None                     qinq-pbit-marking      : both
Ing Agg Rate Limit : max                     Egr Agg Rate Limit     : max
Endpoint           : N/A

Acct. Pol          : None                     Collect Stats           : Disabled

Oper Group         : (none)                   Monitor Oper Grp       : (none)

*B:ALA-Dut-H# show service id 100 all
=====
Service Detailed Information
=====
Service Id         : 100                      Vpn Id              : 100
Service Type       : Epipe
Description        : Default epipe description for service id 100
Customer Id        : 100
Last Status Change: 02/06/2007 10:03:11
Last Mgmt Change   : 02/06/2007 09:43:27
Admin State        : Up                      Oper State           : Up
MTU                : 1514
Vc Switching       : False
SAP Count          : 1                      SDP Bind Count      : 4
-----
Service Destination Points(SDPs)
-----
Sdp Id 1:10100    -(10.20.1.7)
-----
SDP Id            : 1:10100                    Type                : Spoke
VC Type           : Ether                      VC Tag              : n/a
Admin Path MTU    : 1560                      Oper Path MTU       : 1560
Far End           : 10.20.1.7                  Delivery             : MPLS

Admin State       : Up                      Oper State           : Up
Acct. Pol        : None                     Collect Stats        : Disabled
Ingress Label    : 130065                   Egress Label        : 130368
Ing mac Fltr     : n/a                      Egr mac Fltr        : n/a
Ing ip Fltr      : n/a                      Egr ip Fltr         : n/a
Ing ipv6 Fltr    : n/a                      Egr ipv6 Fltr       : n/a
Admin ControlWord: Not Preferred              Oper ControlWord     : False

```

```

Last Status Change : 02/06/2007 10:03:24      Signaling      : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint          : y                          Precedence     : 4
Flags             : SapOperDown
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel
MAC Pinning       : Disabled

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

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

Associated LSP LIST :
Lsp Name          : lsp1_G
Admin State       : Up                        Oper State     : Up
Time Since Last Tr*: 01h40m15s
-----
Sdp Id 2:100  -(10.20.1.4)
-----
SDP Id           : 2:100                      Type           : Spoke
VC Type          : Ether                      VC Tag         : n/a
Admin Path MTU   : 1560                      Oper Path MTU  : 1560
Far End          : 10.20.1.4                  Delivery       : MPLS
Admin State      : Up                        Oper State     : Up
Acct. Pol        : None                      Collect Stats   : Disabled
Ingress Label    : 130671                    Egress Label   : 130367
Ing mac Fltr     : n/a                      Egr mac Fltr   : n/a
Ing ip Fltr      : n/a                      Egr ip Fltr    : n/a
Ing ipv6 Fltr    : n/a                      Egr ipv6 Fltr  : n/a
Admin ControlWord : Not Preferred              Oper ControlWord : False
Last Status Change : 02/06/2007 10:03:11      Signaling      : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint         : y                          Precedence     : 0
Flags            : None
Peer Pw Bits      : pwFwdingStandby
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel
MAC Pinning       : Disabled

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

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

Associated LSP LIST :
Lsp Name          : lsp2_D
Admin State       : Up                        Oper State     : Up
Time Since Last Tr*: 01h40m16s
-----

```

Show, Clear, Debug Commands

```

Sdp Id 3:100  -(10.20.1.5)
-----
SDP Id           : 3:100                      Type           : Spoke
VC Type          : Ether                      VC Tag          : n/a
Admin Path MTU   : 1560                      Oper Path MTU   : 1560
Far End          : 10.20.1.5                  Delivery        : MPLS

Admin State      : Up                        Oper State      : Up
Acct. Pol       : None                      Collect Stats   : Disabled
Ingress Label    : 130971                   Egress Label    : 130368
Ing mac Fltr     : n/a                      Egr mac Fltr    : n/a
Ing ip Fltr      : n/a                      Egr ip Fltr     : n/a
Ing ipv6 Fltr    : n/a                      Egr ipv6 Fltr   : n/a
Admin ControlWord : Not Preferred            Oper ControlWord : False
Last Status Change : 02/06/2007 10:03:17    Signaling       : TLDP
Last Mgmt Change  : 02/06/2007 09:43:27
Endpoint         : y                        Precedence      : 4
Flags            : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel
MAC Pinning      : Disabled

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

Statistics         :
I. Fwd. Pkts.      : 0                     I. Dro. Pkts.    : 0
E. Fwd. Pkts.      : 0                     E. Fwd. Octets   : 0
Associated LSP LIST :
Lsp Name           : lsp3_E
Admin State        : Up                    Oper State       : Up
Time Since Last Tr*: 01h40m16s
-----
...
=====
*B:ALA-Dut-H#

*A:ALU-76# show service id 200 all
=====
Service Detailed Information
=====
Service Id       : 200                      Vpn Id          : 0
Service Type     : Cpipe                     VLL Type        : CESoPSN
Customer Id      : 1
Last Status Change: 09/11/2008 19:05:29
Last Mgmt Change  : 09/10/2008 19:51:06
Admin State      : Up                        Oper State       : Up
MTU              : 1400
Vc Switching     : False
SAP Count        : 1                        SDP Bind Count   : 1
-----
Service Destination Points(SDPs)
-----
Sdp Id 5:200  -(5.5.5.5)
-----
SDP Id           : 5:200                      Type           : Spoke

```



```

VC Type          : CESoPSN                      VC Tag          : 0
Admin Path MTU   : 0                          Oper Path MTU   : 1568
Far End          : 5.5.5.5                     Delivery        : MPLS

Admin State      : Up                          Oper State      : Up
Acct. Pol       : None                        Collect Stats   : Disabled
Ingress Label    : 131061                     Egress Label    : 131066
Ing mac Fltr     : n/a                        Egr mac Fltr    : n/a
Ing ip Fltr      : n/a                        Egr ip Fltr     : n/a
Admin ControlWord : Preferred                   Oper ControlWord : True
Admin BW(Kbps)   : 0                          Oper BW(Kbps)   : 0
Last Status Change : 09/11/2008 19:05:29       Signaling       : TLDP
Last Mgmt Change  : 09/10/2008 19:51:06
Endpoint         : N/A                        Precedence      : 4
Class Fwding State : Down
Flags            : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : pwe3ControlWord mplsRouterAlertLabel

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

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

Associated LSP LIST :
Lsp Name         : to-ALU-80-1
Admin State      : Up                          Oper State      : Up
Time Since Last Tr*: 03d21h52m

-----
Class-based forwarding :
-----
Class forwarding   : disabled                   EnforcedSTELspFc : disabled
Default LSP       : Uknwn                       Multicast LSP     : None
=====
FC Mapping Table
=====
FC Name           LSP Name
-----
No FC Mappings

-----
CPIPE Service Destination Point specifics
-----
Local Bit-rate    : 12                          Peer Bit-rate     : 12
Local Payload Size : 192                        Peer Payload Size : 192
Local Sig Pkts     : No Sig.                     Peer Sig Pkts     : No Sig.
Local CAS Framing   : No CAS                      Peer CAS Framing   : No CAS
Local RTP Header    : No                          Peer RTP Header    : No
Local Differential   : No                          Peer Differential   : No
Local Timestamp     : 0                          Peer Timestamp     : 0
-----
Number of SDPs : 1
-----
Service Access Points
-----

```

Show, Clear, Debug Commands

```

SAP 1/5/1.1.1.1
-----
Service Id      : 200
SAP             : 1/5/1.1.1.1          Encap           : cem
Admin State    : Up                   Oper State      : Up
Flags          : None
Multi Svc Site : None
Last Status Change : 09/10/2008 19:51:27
Last Mgmt Change  : 09/10/2008 19:51:06
Sub Type       : regular

Admin MTU       : 1578                 Oper MTU        : 1578
Ingr IP Fltr-Id : n/a                 Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a                 Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a                Egr IPv6 Fltr-Id : n/a
tod-suite      : None                 qinq-pbit-marking : both
Ing Agg Rate Limit : max               Egr Agg Rate Limit: max
Endpoint       : N/A

Acct. Pol      : None                  Collect Stats    : Disabled
-----
QOS
-----
Ingress qos-policy : 1                 Egress qos-policy : 1
Shared Q plcy     : n/a                Multipoint shared : Disabled
-----
Sap Statistics
-----
Last Cleared Time : N/A
                  Packets                Octets
Forwarding Engine Stats
Dropped           : 0                    0
Off. HiPrio       : 0                    0
Off. LowPrio      : 0                    0
Off. Uncolor      : 0                    0
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio       : 0                    0
Dro. LowPrio      : 0                    0
For. InProf       : 0                    0
For. OutProf      : 0                    0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf       : 0                    0
Dro. OutProf      : 0                    0
For. InProf       : 0                    0
For. OutProf      : 0                    0
-----
Sap per Queue stats
-----
                  Packets                Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio       : 0                    0
Off. LoPrio       : 0                    0
Dro. HiPrio       : 0                    0
Dro. LoPrio       : 0                    0
For. InProf       : 0                    0
For. OutProf      : 0                    0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio       : 0                    0
Off. LoPrio       : 0                    0

```

```

Dro. HiPrio      : 0          0
Dro. LoPrio      : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0

```

```

Egress Queue 1
For. InProf      : 0          0
For. OutProf     : 0          0
Dro. InProf      : 0          0
Dro. OutProf     : 0          0

```

CEM SAP Configuration Information

```

Endpoint Type      : NxDS0          Bit-rate           : 12
Payload Size       : 192             Jitter Buffer (ms)  : 8
Jitter Buffer (packets): 4           Playout Threshold (packets): 3
Use RTP Header     : No              Differential        : No
Timestamp Freq     : 0               CAS Framing         : No CAS

```

```

Cfg Alarm          : stray malformed pktloss overrun underrun
Alarm Status       :

```

CEM SAP Statistics

| | Packets | Seconds | Events |
|---------------------|---------|---------|--------|
| Egress Stats | | | |
| Forwarded | : 0 | | |
| Dropped | : 0 | | |
| Missing | : 0 | | |
| Reordered Forwarded | : 0 | | |
| Underrun | : 0 | | 0 |
| Overrun | : 0 | | 0 |
| Misordered Dropped | : 0 | | |
| Malformed Dropped | : 0 | | |
| LBit Dropped | : 0 | | |
| Multiple Dropped | : 0 | | |
| Error | : | 0 | |
| Severely Error | : | 0 | |
| Unavailable | : | 0 | |
| Failure Count | : | | 0 |
| Jitter Buffer Depth | : 0 | | |
| Ingress Stats | | | |
| Forwarded | : 0 | | |
| Dropped | : 0 | | |

Service Endpoints

No Endpoints found.

*A:ALU-76#

*A:bksim180# show service id 1000 all

Service Detailed Information

```

Service Id        : 1000          Vpn Id             : 0
Service Type      : Ipipe
Customer Id       : 1
Last Status Change: 03/11/1973 10:20:24
Last Mgmt Change  : 03/11/1973 10:20:23

```

Show, Clear, Debug Commands

```

Admin State      : Up                      Oper State      : Up
MTU              : 1400
Vc Switching    : False
SAP Count       : 1                      SDP Bind Count  : 1
CE Addr Discovery : enabled
-----
Service Destination Points(SDPs)
-----
Sdp Id 22:1000  -(2.2.2.2)
-----
SDP Id          : 22:1000                      Type           : Spoke
VC Type         : Ipipe                        VC Tag         : 0
Admin Path MTU  : 0                          Oper Path MTU   : 1568
Far End         : 2.2.2.2                      Delivery        : MPLS

Admin State     : Up                          Oper State     : Up
Acct. Pol      : None                        Collect Stats   : Disabled
Ingress Label   : 131070                      Egress Label    : 131062
Ing mac Fltr    : n/a                        Egr mac Fltr    : n/a
Ing ip Fltr     : n/a                        Egr ip Fltr     : n/a
Admin ControlWord : Not Preferred              Oper ControlWord : False
Admin BW(Kbps)  : 0                          Oper BW(Kbps)    : 0
Last Status Change : 03/11/1973 10:20:24      Signaling       : TLDP
Last Mgmt Change  : 03/11/1973 10:19:21
Endpoint        : N/A                        Precedence      : 4
Class Fwding State : Down
Flags           : None
Time to RetryReset : 1999616832 seconds        Retries Left    : 2984947
Mac Move        : Ukwn                       Blockable Level  : Unknown
Peer Pw Bits    : None
Peer Fault Ip   : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel

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

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

Associated LSP LIST :
Lsp Name           : to-bksim176-1
Admin State        : Up                      Oper State         : Up
Time Since Last Tr*: 00h01m28s
-----
Class-based forwarding :
-----
Class forwarding    : disabled
Default LSP         : Ukwn                      Multicast LSP      : None
=====
FC Mapping Table
=====
FC Name            LSP Name
-----
No FC Mappings
-----
IPIPE Service Destination Point specifics

```

```

-----
Configured CE IP Addr : n/a                      Peer CE IP Addr : 1.1.1.2
-----
Number of SDPs : 1
-----
Service Access Points
-----
SAP 1/7/1
-----
Service Id      : 1000
SAP             : 1/7/1                      Encap          : null
Admin State    : Up                        Oper State      : Up
Flags          : None
Multi Svc Site : None
Last Status Change : 03/11/1973 10:20:23
Last Mgmt Change  : 03/11/1973 10:19:21
Sub Type       : regular
Dot1Q Ethertype : 0x8100                      QinQ Ethertype  : 0x8100

Admin MTU      : 1514                      Oper MTU        : 1514
Ingr IP Fltr-Id : n/a                      Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a                      Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a                    Egr IPv6 Fltr-Id : n/a
tod-suite      : None                      qinq-pbit-marking : both
Ing Agg Rate Limit : max                    Egr Agg Rate Limit: max
Endpoint       : N/A

Q Frame-Based Acct : Disabled \

Acct. Pol      : None                      Collect Stats   : Disabled
-----
Ipipe SAP Info
-----
Configured CE IP : n/a                      Discovered CE IP : 1.1.1.1
SAP MAC Address  : 8c:c7:01:07:00:01          Mac Refresh Inter*: 14400
-----
Ipipe SAP ARP Entry Info
-----
1.1.1.1          8c:c7:01:07:00:03  dynamic  04h00m00s
-----
QOS
-----
Ingress qos-policy : 1                      Egress qos-policy : 1
Shared Q plcy      : n/a                    Multipoint shared : Disabled
-----
Sap Statistics
-----
Last Cleared Time : N/A
                  Packets                      Octets

Forwarding Engine Stats
Dropped           : 0                          0
Off. HiPrio       : 0                          0
Off. LowPrio      : 0                          0
Off. Uncolor      : 0                          0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio       : 0                          0
Dro. LowPrio      : 0                          0
For. InProf       : 0                          0
For. OutProf      : 0                          0

```

Show, Clear, Debug Commands

```

Queueing Stats(Egress QoS Policy 1)
Dro. InProf      : 0          0
Dro. OutProf     : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0
-----
Sap per Queue stats
-----
                Packets          Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0          0
Off. LoPrio      : 0          0
Dro. HiPrio      : 0          0
Dro. LoPrio      : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio      : 0          0
Off. LoPrio      : 0          0
Dro. HiPrio      : 0          0
Dro. LoPrio      : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0

Egress Queue 1
For. InProf      : 0          0
For. OutProf     : 0          0
Dro. InProf      : 0          0
Dro. OutProf     : 0          0
-----
Service Endpoints
-----
No Endpoints found.
=====
*A:bksiml80#

*A:ces-A# show service id 1 all
=====
Service Detailed Information
=====
Service Id       : 1          Vpn Id           : 0
Service Type     : Cpipe      VLL Type        : SAToPT1
Description      : (Not Specified)
Customer Id      : 1
Last Status Change: 07/06/2010 19:21:14
Last Mgmt Change : 07/06/2010 19:21:14
Admin State      : Up         Oper State       : Up
MTU              : 1514
Vc Switching     : False
SAP Count        : 1          SDP Bind Count   : 1
-----
Service Destination Points(SDPs)
-----
Sdp Id 12:1    -(2.2.2.2)
-----
Description      : (Not Specified)
SDP Id          : 12:1          Type             : Spoke
VC Type         : SAToPT1      VC Tag           : 0

```

| | | | |
|-------------------------|--|------------------|------------|
| Admin Path MTU | : 0 | Oper Path MTU | : 9190 |
| Far End | : 2.2.2.2 | Delivery | : MPLS |
| Admin State | : Up | Oper State | : Up |
| Acct. Pol | : None | Collect Stats | : Disabled |
| Ingress Label | : 131064 | Egress Label | : 131064 |
| Admin ControlWord | : Preferred | Oper ControlWord | : True |
| Admin BW(Kbps) | : 0 | Oper BW(Kbps) | : 0 |
| Last Status Change | : 07/06/2010 19:21:14 | Signaling | : TLDP |
| Last Mgmt Change | : 07/06/2010 19:21:14 | Precedence | : 4 |
| Endpoint | : N/A | | |
| Flags | : None | | |
| Peer Pw Bits | : None | | |
| Peer Fault Ip | : None | | |
| Peer Vccv CV Bits | : lspPing | | |
| Peer Vccv CC Bits | : pwe3ControlWord mplsRouterAlertLabel | | |
| KeepAlive Information : | | | |
| Admin State | : Enabled | Oper State | : Alive |
| Hello Time | : 10 | Hello Msg Len | : 0 |
| Max Drop Count | : 3 | Hold Down Time | : 10 |
| Statistics : | | | |
| I. Fwd. Pkts. | : 141578 | I. Fwd. Octs. | : 31430316 |
| E. Fwd. Pkts. | : 141583 | E. Fwd. Octets | : 31431426 |
| Associated LSP LIST : | | | |
| Lsp Name | : to_b_1_2 | | |
| Admin State | : Up | Oper State | : Up |
| Time Since Last Tr* | : 04h08m22s | | |

base

| | |
|--------------------|---|
| Syntax | base |
| Context | show>service>id |
| Description | 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 | 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
=====
```

```
Remote-Ve Information
-----
Remote VE Name      : PE2                      Remote VE Id      : 2
=====
*A:cses-E11>config>service>epipe>bgp-vpws#
```

endpoint

| | |
|--------------------|--|
| Syntax | endpoint [<i>endpoint-name</i>] |
| Context | show>service>id |
| Description | This command displays service endpoint information. |
| Parameters | <i>endpoint-name</i> — 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#
```

retailers

| | |
|--------------------|---|
| Syntax | retailers |
| Context | show>service>id |
| Description | This command displays the service ID of the retailer subscriber service to which this DHCP lease belongs. |

wholesalers

| | |
|--------------------|---|
| Syntax | wholesalers |
| Context | show>service>id |
| Description | This command displays service wholesaler information. |

sap

| | |
|--------------------|---|
| Syntax | sap <i>sap-id</i> [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 | <p><i>sap-id</i> — The ID that displays SAPs for the service in the form <i>slot/mda/port</i>. See Common CLI Command Descriptions on page 1319 for command syntax.</p> <p>interface <i>interface-name</i> — Displays information for the specified IP interface.</p> <p>ip-address <i>ip-address</i> — Displays information associated with the specified IP address.</p> <p>detail — Displays detailed information.</p> <p>wholesaler <i>service-id</i> — The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored.</p> <p>Values 1 — 2147483648</p> <p>detail — Displays detailed information for the SAP.</p> |
| 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, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode. |
| Last Status Change | The time of the most recent operating status change to this SAP. |
| Last Mgmt Change | The time of the most recent management-initiated change to this SAP. |

| Label | Description (Continued) |
|--------------------|--|
| Admin MTU | The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented. |
| Oper MTU | The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented. |
| Ingress qos-policy | The ingress QoS policy ID assigned to the SAP. |
| Egress qos-policy | The egress QoS policy ID assigned to the SAP. |
| Ingress Filter-Id | The ingress filter policy ID assigned to the SAP. |
| Egress Filter-Id | The egress filter policy ID assigned to the SAP. |
| Acct. Pol | The accounting policy ID assigned to the SAP. |
| Collect Stats | Specifies whether collect stats is enabled. |
| 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|
                                     vpil.vpi2|
...
                                     ima-grp     - <bundle-id>[:vpi/vci|vpi|vpil.vpi2|
```

Show, Clear, Debug Commands

```
A:ALA-48>config>service>epipe# show service id 8 sap 881/1/2:4094
=====
Service Access Points(SAP)
=====
Service Id      : 8
SAP             : 8/1/2:4094          Encap           : bcpDot1q
Admin State     : Up                  Oper State      : Down
Flags           : ServiceAdminDown   PortOperDown
Last Status Change : 02/06/2007 12:01:14
Last Mgmt Change  : 02/06/2007 12:01:17
Admin MTU        : 1522               Oper MTU        : 1522
Ingress qos-policy : 1                Egress qos-policy : 1
Shared Q plcy    : n/a                Multipoint shared : Disabled
Ingress Filter-Id : n/a               Egress Filter-Id  : n/a
tod-suite       : None

Multi Svc Site   : None
Acct. Pol        : None                Collect Stats    : Disabled
=====
A:ALA-48>config>service>epipe#
A:ALA-48>config>service>epipe# show service id 8 sap 881/1/2:4094 detail
=====
Service Access Points(SAP)
=====
Service Id      : 8
SAP             : 8/1/2:4094          Encap           : bcpDot1q
Admin State     : Up                  Oper State      : Down
Flags           : ServiceAdminDown   PortOperDown
Last Status Change : 02/06/2007 12:01:14
Last Mgmt Change  : 02/06/2007 12:01:17
Admin MTU        : 1522               Oper MTU        : 1522
Ingress qos-policy : 1                Egress qos-policy : 1
Shared Q plcy    : n/a                Multipoint shared : Disabled
Ingress Filter-Id : n/a               Egress Filter-Id  : n/a
tod-suite       : None

Multi Svc Site   : None
Acct. Pol        : None                Collect Stats    : Disabled
-----
Sap Statistics
-----
Packets          Octets
Forwarding Engine Stats
Dropped          : 0                0
Off. HiPrio      : 0                0
Off. LowPrio     : 0                0
Off. Uncolor     : 0                0
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio      : 0                0
Dro. LowPrio     : 0                0
For. InProf      : 0                0
For. OutProf     : 0                0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf      : 0                0
Dro. OutProf     : 0                0
For. InProf      : 0                0
For. OutProf     : 0                0
-----
```

```

Sap per Queue stats
-----
Packets
Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0      0
Off. LoPrio      : 0      0
Dro. HiPrio      : 0      0
Dro. LoPrio      : 0      0
For. InProf      : 0      0
For. OutProf     : 0      0
Egress Queue 1
For. InProf      : 0      0
For. OutProf     : 0      0
Dro. InProf      : 0      0
Dro. OutProf     : 0      0
=====
A:ALA-48>config>service>epipe#

```

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  Total MAC Addr  : 0
Learned MAC Addr  : 0          Static MAC Addr : 0
Admin MTU        : 1518        Oper MTU        : 1518
Ingress qos-policy : 1190      Egress qos-policy : 1190
Shared Q plcy    : n/a         Multipoint shared : Disabled
Ingr IP Fltr-Id  : n/a         Egr IP Fltr-Id   : n/a
Ingr Mac Fltr-Id : n/a         Egr Mac Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a        Egr IPv6 Fltr-Id : n/a
tod-suite       : suite_sixteen  qinq-pbit-marking : both
Egr Agg Rate Limit : max
ARP Reply Agent   : Unknown     Host Conn Verify : Disabled
Mac Learning      : Enabled      Discard Unkwn Srce: Disabled
Mac Aging         : Enabled      Mac Pinning      : Disabled
L2PT Termination  : Disabled     BPDU Translation : Disabled
Multi Svc Site    : None
I. Sched Pol      : SchedPolCust1_Night
E. Sched Pol      : SchedPolCust1Egress_Night
Acct. Pol         : None         Collect Stats     : Disabled
Anti Spoofing     : None         Nbr Static Hosts : 0
=====
A:ALA-48#

A:kerckhot_4# show service id 1 sap 1/1/1:6

```

Show, Clear, Debug Commands

```
=====
Service Access Points(SAP)
=====
Service Id      : 1
SAP             : 1/1/1:6
Dot1Q Ethertype : 0x8100
Admin State     : Up
Encap           : q-tag
QinQ Ethertype  : 0x8100
Oper State      : Down
Flags          : TodResourceUnavail
Last Status Change : 12/01/2006 09:59:42
Last Mgmt Change  : 12/01/2006 09:59:45
...
A:kerckhot_4#
```


sdp

| | |
|--------------------|---|
| Syntax | sdp <i>[[sdp-id[:vc-id] far-end ip-address] [detail]</i> sdp <i>sdp-id:vc-id mrp</i> |
| 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 | <i>sdp-id</i> — Displays only information for the specified SDP ID. Default All SDPs. Values 1 — 17407 <i>far-end ip-address</i> — 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
Fwd Transitions      : 0                               Designated Port Id : 0
Cfg BPDUs rcvd       : 0                               Bad BPDUs rcvd     : 0
TCN 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#

```

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
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                  Total MAC Addr   : 0
Learned MAC Addr   : 0                        Static MAC Addr  : 0
MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                BPDU Translation  : Disabled
MAC Pinning        : Disabled
KeepAlive Information :
Admin State        : Disabled                 Oper State       : Disabled
Hello Time         : 10                      Hello Msg Len    : 0
Max Drop Count     : 3                      Hold Down Time   : 10
Statistics         :
I. Fwd. Pkts.      : 0                      I. Dro. Pkts.    : 0
E. Fwd. Pkts.      : 0                      E. Fwd. Octets   : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

-----
Number of SDPs : 1
=====
*A:ALA-Dut-B>config>service>epipe#

*A:SetupCLI>config>service>epipe>spoke-sdp# show service id 2 sdp 2000:1 detail
=====
Service Destination Point (Sdp Id : 2000:1) Details
=====
-----
Sdp Id 2000:1  -(101.101.101.101)
-----
Description      : (Not Specified)
SDP Id           : 2000:1                      Type            : Spoke
Spoke Descr      : (Not Specified)
VC Type          : Ether
Admin Path MTU   : 1500                      Oper Path MTU    : 1500
Far End          : 101.101.101.101             Delivery        : MPLS
Hash Label       : Enabled
Admin State      : Up                        Oper State       : Down
Acct. Pol        : None                      Collect Stats    : Disabled
Ingress Label    : 0                        Egress Label     : 0
Ingr Mac Fltr-Id : n/a                      Egr Mac Fltr-Id  : n/a
Ingr IP Fltr-Id  : n/a                      Egr IP Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a                     Egr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred             Oper ControlWord  : False
Admin BW(Kbps)   : 0                        Oper BW(Kbps)    : 0
Last Status Change : 10/08/2009 06:55:54      Signaling        : TLDP
Last Mgmt Change   : 10/08/2009 07:04:27      Force Vlan-Vc    : Disabled
Endpoint          : N/A                      Precedence       : 4
Class Fwding State : Down
Flags             : SvcAdminDown SdpOperDown
                  : NoIngVCLabel NoEgrVCLabel
                  : PathMTUTooSmall
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
```



```

Application Profile: None

KeepAlive Information :
Admin State           : Enabled                Oper State           : No response
Hello Time            : 600                    Hello Msg Len        : 1500
Max Drop Count        : 3                      Hold Down Time       : 10

Statistics            :
I. Fwd. Pkts.         : 0                      I. Dro. Pkts.        : 0
E. Fwd. Pkts.         : 0                      E. Fwd. Octets       : 0
-----RSVP/Static
LSPs
-----Associated
LSP LIST :
No LSPs Associated
-----Class-based
forwarding :
-----
Class forwarding      : Disabled                EnforceDSTELspFc    : Disabled
Default LSP           : Uknwn                  Multicast LSP        : None
=====
FC Mapping Table
=====
FC Name               LSP Name
-----
No FC Mappings
-----
Number of SDPs : 1
=====
*A:SetupCLI>config>service>epipe>spoke-sdp#

```

Sample Output for L2TPv3 SDP binding

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
=====

```

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)
=====
```

```
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                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> <i>service-name</i> neighbor |
| Context | clear>service clear>service>statistics |
| Description | <p>This command clears commands for a specific service.</p> <p>It clears the discovered IPv6 address of the neighboring CE associated with an iPipe SAP. Note that when IPv6CP comes back up following the execution of this command on an IPv6CP SAP, the 7x50 will check to see if an IPv6 address has been learned for the remote CE attached to the ipipe service. If one has been learned, then this is used to bring up IPv6CP.</p> |
| 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. |
| | <i>service-name</i> — Neighboring IPv6 address. |

arp

| | |
|--------------------|---|
| Syntax | arp |
| Context | clear>service>id |
| Description | This command clears all ARP entries. This command is only valid for Ipipe services. |

neighbor

| | |
|--------------------|--|
| Syntax | neighbor |
| Context | clear>service>id |
| Description | Clears the discovered IPv6 address of the neighboring CE associated with an iPipe SAP. |

host-tracking

| | |
|--------------------|---|
| Syntax | host-tracking [sap <i>sap-id</i>] [host <i>ip-address</i>] |
| Context | clear>service>id |
| Description | This command clears host tracking data. |

mesh-sdp

| | |
|--------------------|--|
| Syntax | mesh-sdp <i>sdp-id</i> [: <i>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</i> : <i>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. lt2pv3 — 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 | <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1319 for command syntax. all — Clears all SAP queue statistics and STP statistics. counters — Clears all queue statistics associated with the SAP. stp — Clears all STP statistics associated with the SAP. |

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 | <i>sdp-id</i> — The SDP ID for which to clear keepalive statistics. Values 1 — 17407 |

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. |
| 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. |

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 service-id: 1 — 214748364 svc-name: A string up to 64 characters in length. |

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 neighbor-discovery 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. neighbor-discovery — displays the status of IPv6 neighbour discovery for the sap or the spoke-sdp. |

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)
```

Show, Clear, Debug Commands

```
hwType      : 0x0001
prType      : 0x0800
hwLength    : 0x06
prLength    : 0x04
srcMac      : 8c:c7:01:07:00:03
destMac     : 00:00:00:00:00:00
srcIp       : 200.1.1.2
destIp      : 200.1.1.1
"

4 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000 SAP 1/7/1 "Service
1000 SAP 1/7/1:
TX: ARP_RESPONSE (0x0002)
hwType      : 0x0001
prType      : 0x0800
hwLength    : 0x06
prLength    : 0x04
srcMac      : 00:03:0a:0a:0a:0a
destMac     : 8c:c7:01:07:00:03
srcIp       : 200.1.1.1
destIp      : 200.1.1.2
"
```

sdp

| | |
|--------------------|--|
| Syntax | [no] sdp <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. |

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 411](#)
- [VPLS Features on page 415](#)
 - [VPLS Packet Walkthrough on page 412](#)
 - [VPLS Enhancements on page 415](#)
 - [VPLS over MPLS on page 416](#)
 - [VPLS Service Pseudowire VLAN Tag Processing on page 417](#)
 - [VPLS MAC Learning and Packet Forwarding on page 421](#)
 - [VPLS Using G.8031 Protected Ethernet Tunnels on page 424](#)
 - [Pseudowire Control Word on page 425](#)
 - [Table Management on page 426](#)
 - [VPLS and Spanning Tree Protocol on page 434](#)
 - [Multiple Spanning Tree on page 436](#)
 - [Egress Multicast Groups on page 443](#)
 - [VPLS Redundancy on page 454](#)
 - [Object Grouping and State Monitoring on page 472](#)
 - [MAC Flush Message Processing on page 474](#)
 - [ACL Next-Hop for VPLS on page 478](#)
 - [SDP Statistics for VPLS and VLL Services on page 479](#)
 - [BGP Auto-Discovery for LDP VPLS on page 480](#)

- [Multicast-Aware VPLS on page 499](#)
- [RSVP and LDP P2MP LSP for Forwarding VPLS/B-VPLS BUM and IP Multicast Packets on page 503](#)
- [VPLS Service Considerations on page 521](#)
 - [SAP Encapsulations on page 521](#)
 - [VLAN Processing on page 521](#)
 - [Ingress VLAN Swapping on page 522](#)
 - [Service Auto-Discovery using Multiple VLAN Registration Protocol \(MVRP\) on page 523](#)
 - [VPLS E-Tree Services on page 533](#)

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 55) from site-A, which is connected to PE-Router-A, to site-B, which is connected to PE-Router-C (Figure 56).

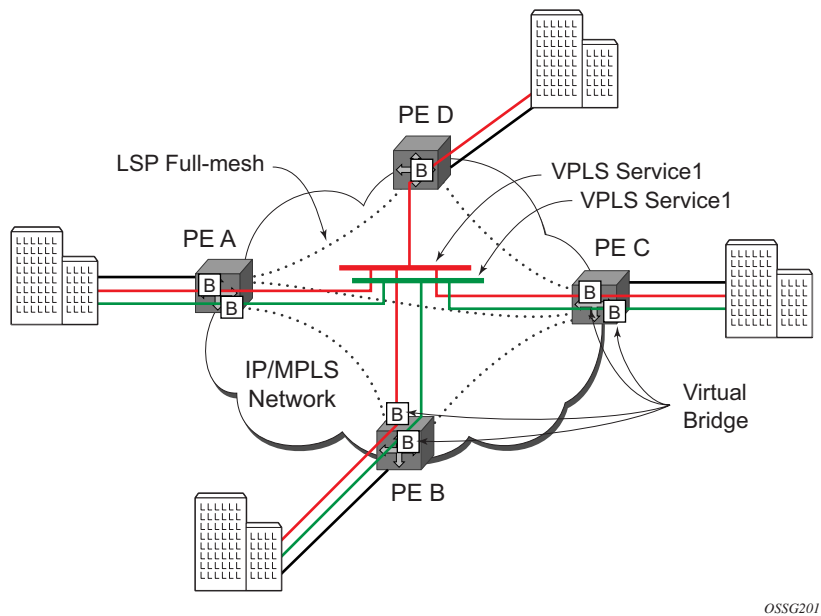


Figure 55: VPLS Service Architecture

- 1. PE-Router-A (Figure 56)
 - a. Service packets arriving at PE-Router-A are associated with a VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet

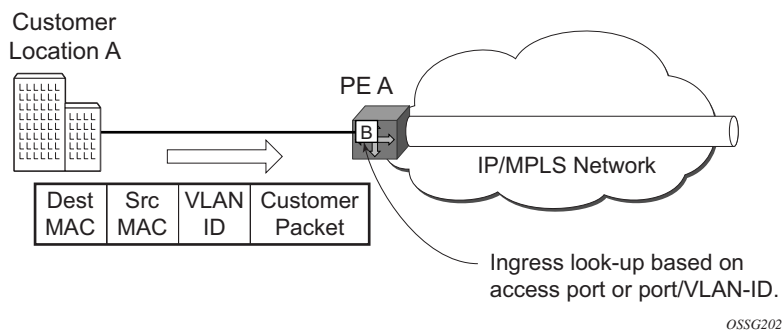


Figure 56: Access Port Ingress Packet Format and Lookup

- 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 57):

- 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 57):

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.

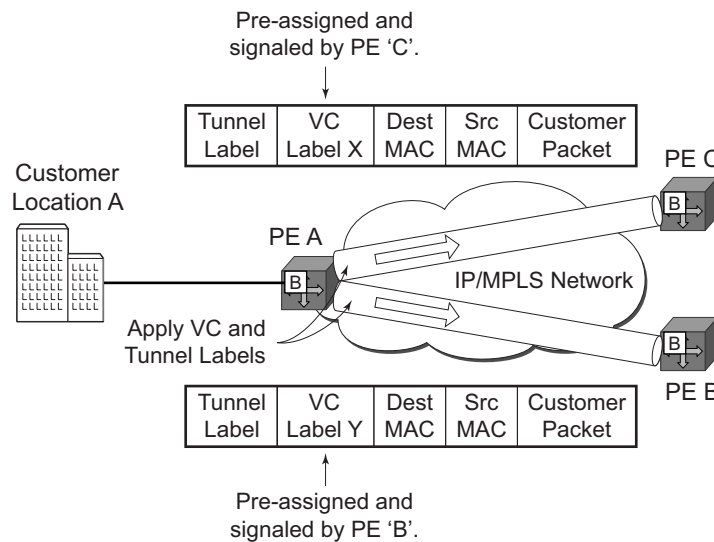


Figure 57: Network Port Egress Packet Format and Flooding

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 58)

- 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.

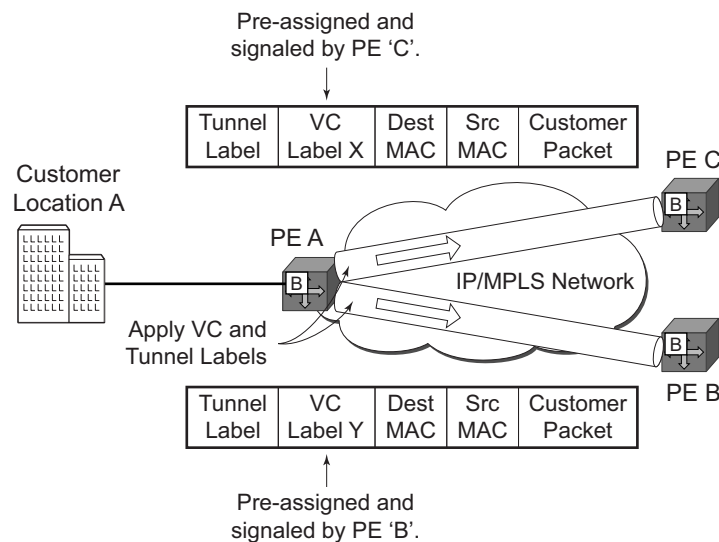


Figure 58: Access Port Egress Packet Format and Lookup

VPLS Features

This section features:

- [VPLS Enhancements on page 415](#)
 - [Pseudowire Control Word on page 425](#)
 - [Split Horizon SAP Groups and Split Horizon Spoke SDP Groups on page 433](#)
 - [VPLS and Spanning Tree Protocol on page 434](#)
 - [Egress Multicast Groups on page 443](#)
 - [VPLS Redundancy on page 454](#)
 - [VPLS Access Redundancy on page 470](#)
 - [VCCV BFD Support for VPLS Services on page 493](#)
-

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 on a per VPLS, per SAP and per spoke SDP basis
 - FIB size alarms on a per VPLS basis
 - MAC learning disable on a per VPLS, per SAP and per spoke SDP basis
 - Discard unknown on a per VPLS basis
 - Separate aging timers for locally and remotely learned MAC addresses.
- Ingress rate limiting for broadcast, multicast, and destination unknown flooding on a per SAP basis.
- Implementation of 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.
- DHCP snooping and anti-spoofing on a per-SAP and per-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**:

- 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 12 and Table 13 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 12: 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 13: 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 | <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 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 13: 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 7450 ESS edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the 7450 ESS to reduce the amount of unknown destination MAC address flooding.

7450 ESS routers learn the source MAC addresses of the traffic arriving on their access and network ports.

Each 7450 ESS 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, subscribers connected to SAPs A, B, C can create a denial of service attack by sending packets sourcing the gateway MAC address. This will move the learned gateway MAC from the uplink SDP/SAP to the subscriber's SAP causing all communication to the gateway to be disrupted. If local content is attached to the same VPLS (D), a similar attack can be launched against it. Communication between subscribers is also disallowed but split-horizon will not be sufficient in the topology depicted in [Figure 59](#).

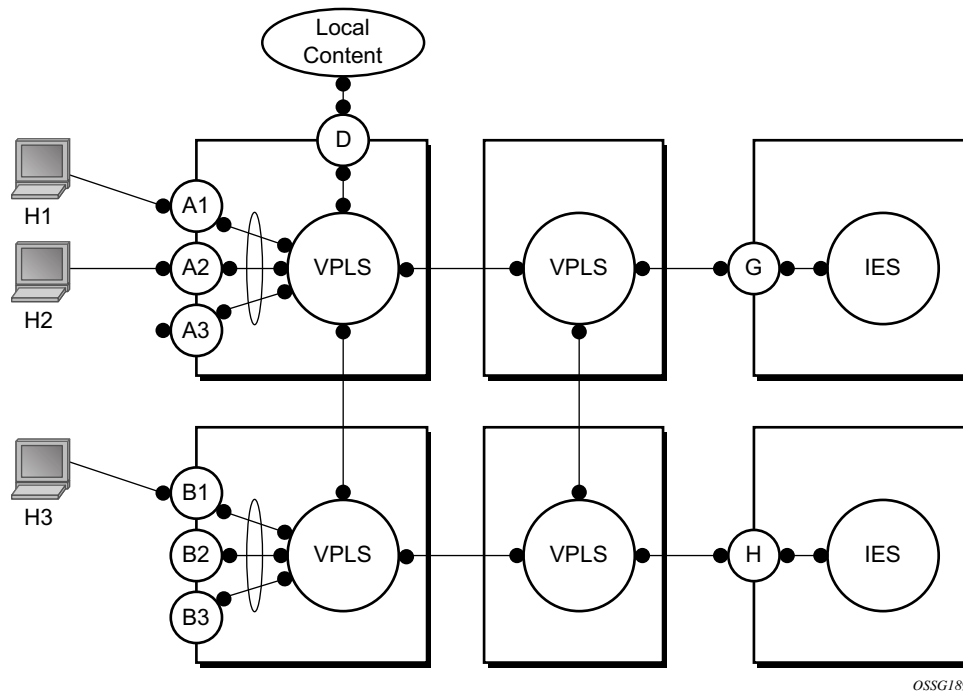


Figure 59: MAC Learning Protection

7450 ESSs 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 59) 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 subscriber 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 60](#)).



Figure 60: 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. Refer to the *IEEE 802.1ah PBB Guide* for examples.

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 ingressed on a SAP. A remote MAC address is a MAC address received by an SDP from another router for the VPLS instance. The local-age timer for the VPLS instance specifies the aging time for locally learned MAC addresses, and the remote-age timer specifies the aging time for remotely learned MAC addresses.

In general, the remote-age timer is set to a longer period than the local-age timer to reduce the amount of flooding required for destination unknown MAC addresses. The aging mechanism is considered a low priority process. In most situations, the aging out of MAC addresses 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 7450 ESS 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.

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

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.

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 7450 ESS 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 539](#).

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 7450 ESS 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 7450 ESS supports 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 61](#) 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 61: 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 7450 ESS 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 7450 ESS 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 UDLD. This allows running these protocols between customer CPEs without involving backbone infrastructure.

7450 ESS 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.

7450 ESS 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 7450 ESS 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 444](#)
 - [SAP Port Dot1Q EtherType on page 444](#)
 - [Egress Multicast Groups on page 445](#)
 - [SAP Egress Filter on page 445](#)
-

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.
 - All SAP members of the multicast-group must have the same inner-vlan-tag (the default setting) or must have the same value of outer-vlan-tag as defined by the **qinq-fixed-tag-value** command.
-

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 subscriber 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 IOM (Input Output Module) automatically creates the SAP chains on each egress forwarding plane (typically all ports on an MDA are part of a single forwarding plane except in the case of the 10 Gigabit IOM which has two MDAs on a single 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 IOM 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 448](#)). 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 IOM 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 449](#) and [MVR IGMP Snooping \(s,g\) Flooding List on page 449](#)). 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 IOM 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 IOM 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 IOM 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 IOM 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 IOM 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 449](#). 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 IOM 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 IOM 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 IOM.

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.

IOM Chain Management

As previously stated, the IOM automatically creates the chain lists from the available egress multicast group SAPs. The IOM 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 IOM 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 14](#).

Table 14: 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 IOM 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 IOM 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.

IOM Mode B Capability

Efficient multicast replication uses an egress forwarding plane that supports chassis mode b due to the expanded memory requirements to store the replication chain information. The system does not need to be placed into mode b for efficient multicast replication to be performed. Any IOM that is capable of mode “b” operation automatically performs efficient multicast replication when a flooding list contains SAPs that are members of an egress multicast group.

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.

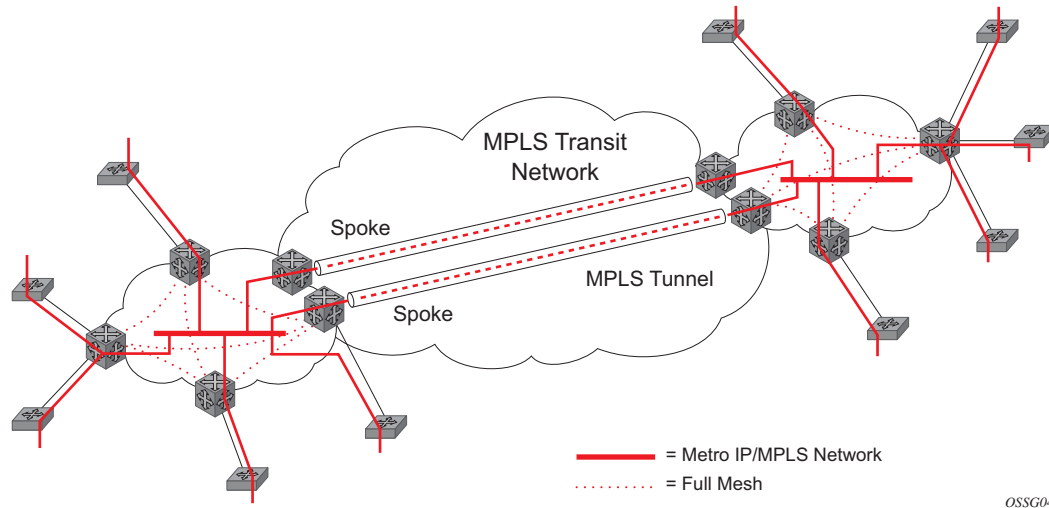
Alcatel-Lucent routers have implemented special features for improving the resilience of hierarchical VPLS instances, in both MTU and inter-metro applications.

Spoke SDP Redundancy for Metro Interconnection

When two or more meshed VPLS instances are interconnected by redundant spoke SDPs (as shown in [Figure 62](#)), 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.

**Figure 62: 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 62](#), 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 63](#).




Figure 63: 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 64](#) can be used.

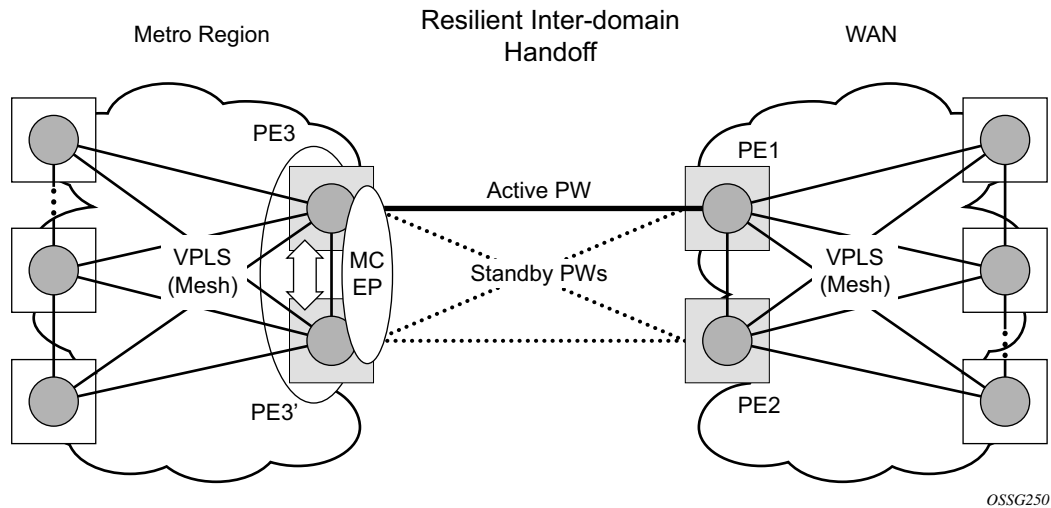


Figure 64: Multi-Chassis Pseudowire Endpoint for VPLS

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.

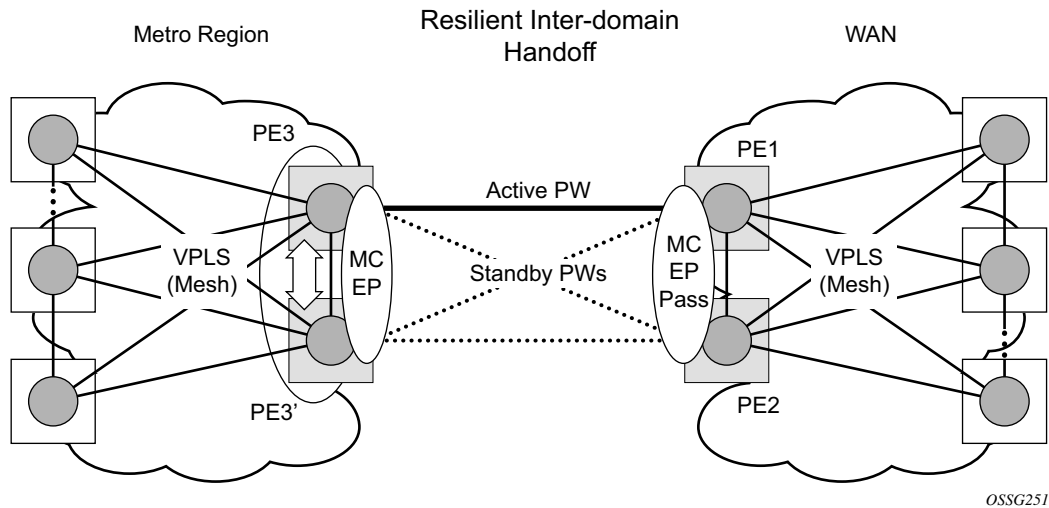


Figure 65: MC-EP in Passive Mode

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 66](#) 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.

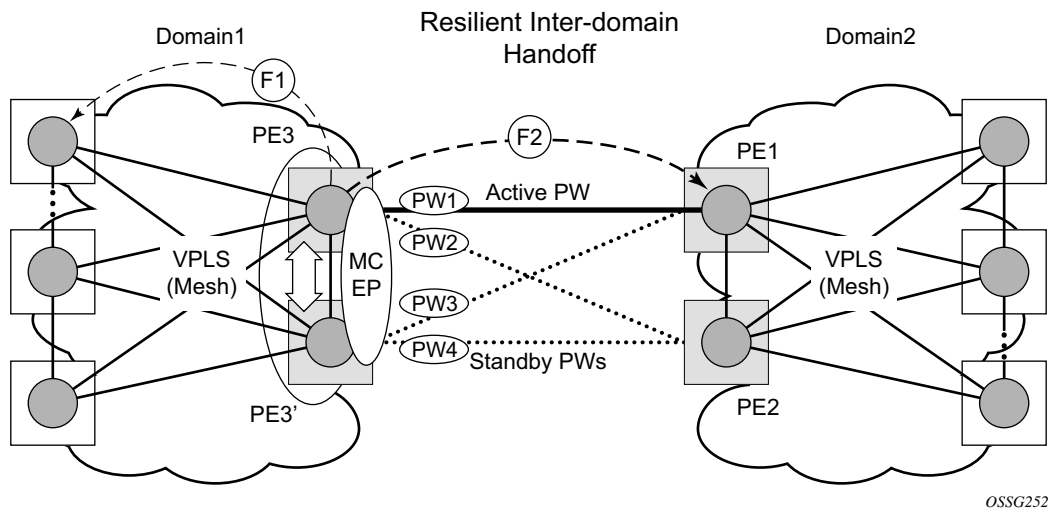


Figure 66: MAC Flush in the MC-EP Solution

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 66](#)):

- 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 619](#) 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.

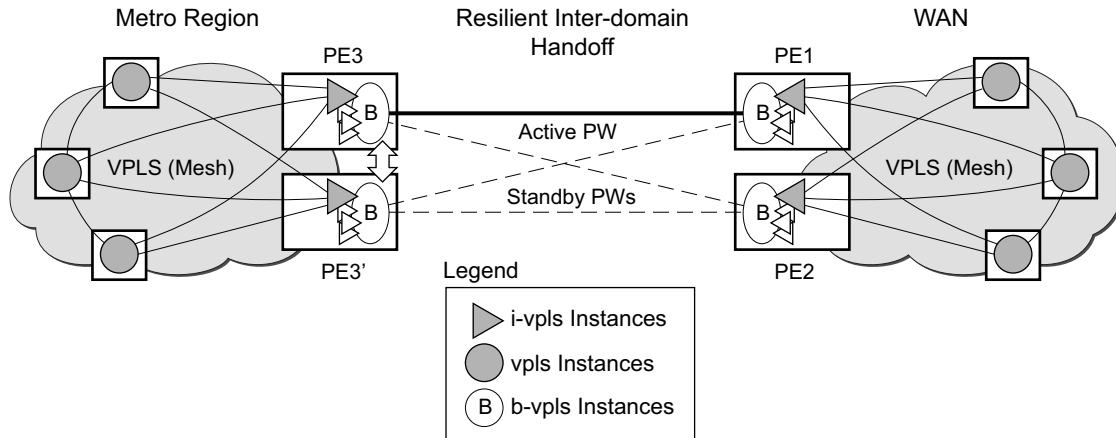
In the MC-EP case the revertive behavior is supported for pseudowire defined as primary (precedence 0). The following rules apply:

- 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 67](#) below.



OSSG487

Figure 67: 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 67](#), 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 68](#) is used to define the blackholing problem in PBB-VPLS using MC-EP.

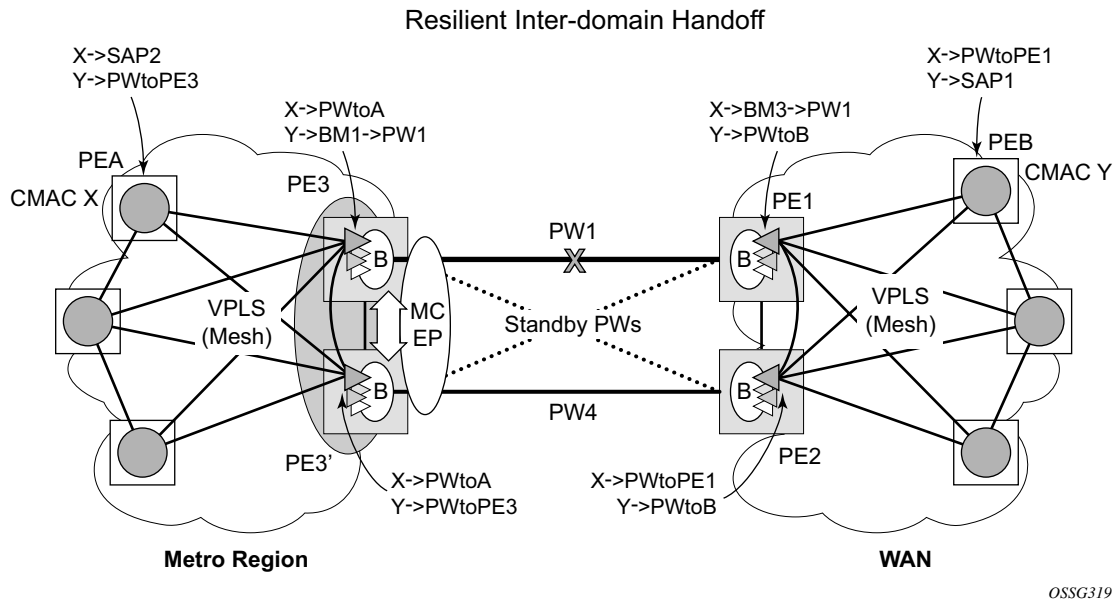


Figure 68: MC-EP with B-VPLS Failure Scenario

In topology displayed in [Figure 68](#), 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 69](#).

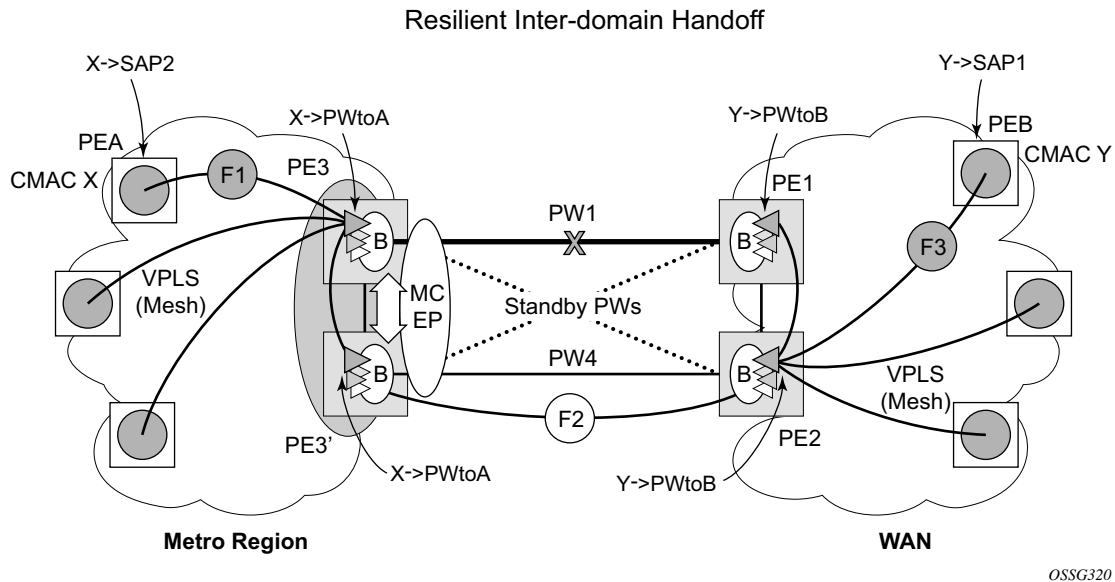


Figure 69: MC-EP with B-VPLS Mac Flush Solution

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: PEA, 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

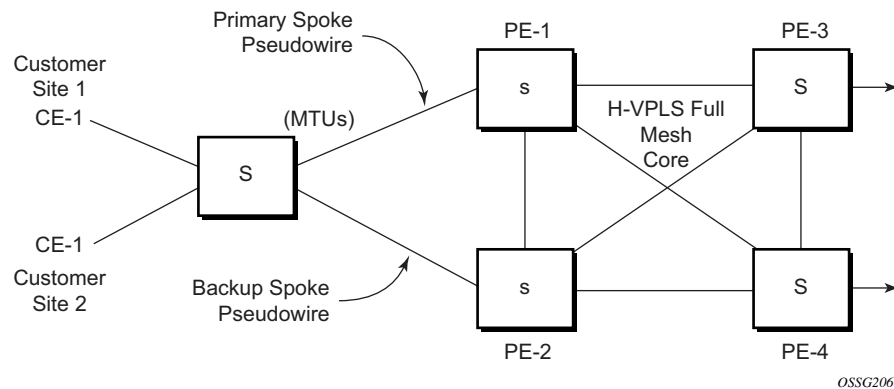


Figure 70: Dual Homed MTU-s in Two-Tier Hierarchy H-VPLS

In configuration shown in [Figure 70](#), 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 an loop resolution mechanism, every TCN (Topology Change Notification) received in a context of STP instance is translated into LDP- MAC address withdrawal message (also referred to as MAC-flush message) requesting to clear all FDB entries, 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 62](#) 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 7450 ESS supports 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 than 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 7750 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

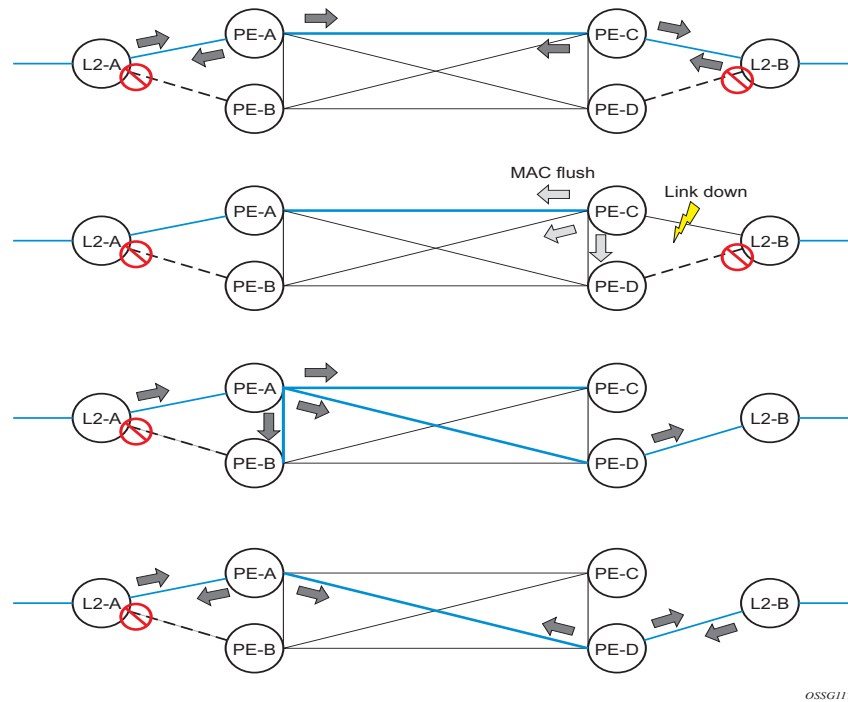


Figure 71: Dual Homed CE Connection to VPLS

Figure 71 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 71](#)) 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.

ACL Next-Hop for VPLS

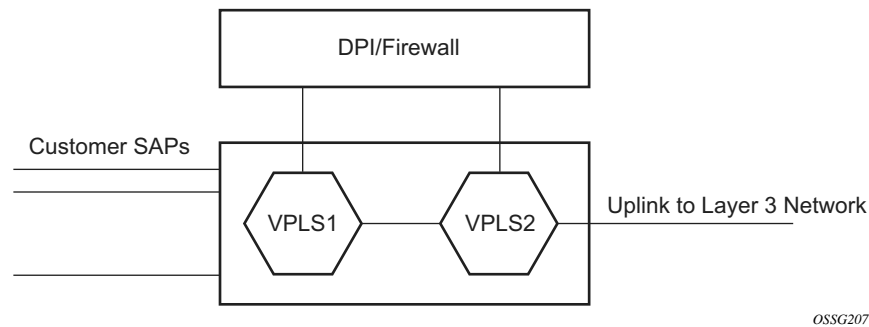


Figure 72: Application 1 Diagram

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 73: SDP Statistics for VPLS and VLL Services

The simple three-node network described in [Figure 73](#) 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 74](#). The system IP may be used for the last 4 bytes of the VSI ID further simplifying the addressing and the provisioning process.




Figure 74: 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 75](#).

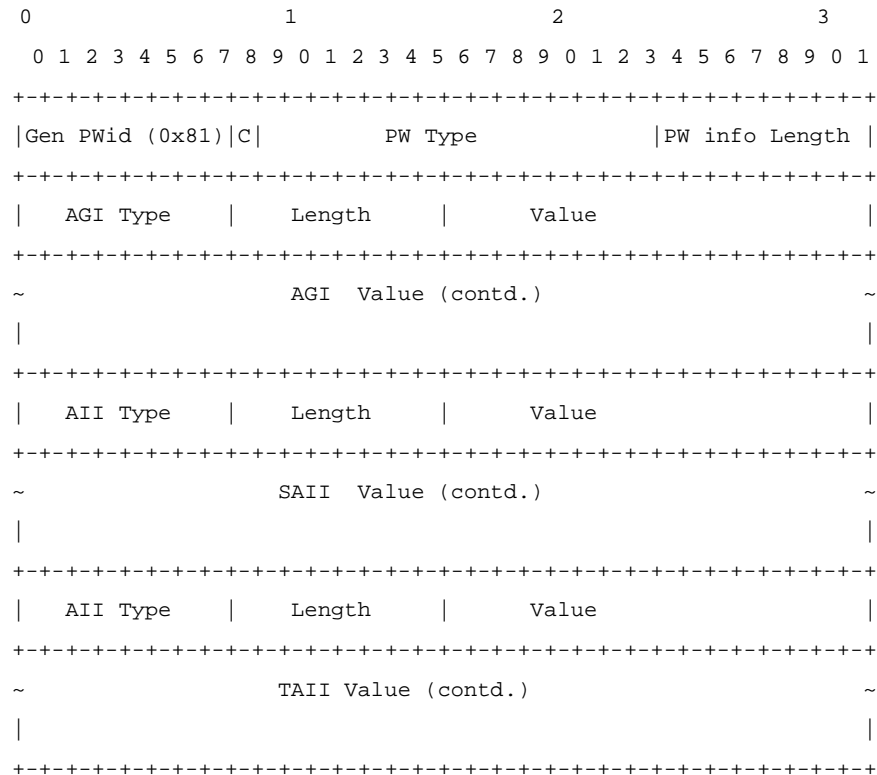


Figure 75: Generalized Pseudowire-ID FEC Element

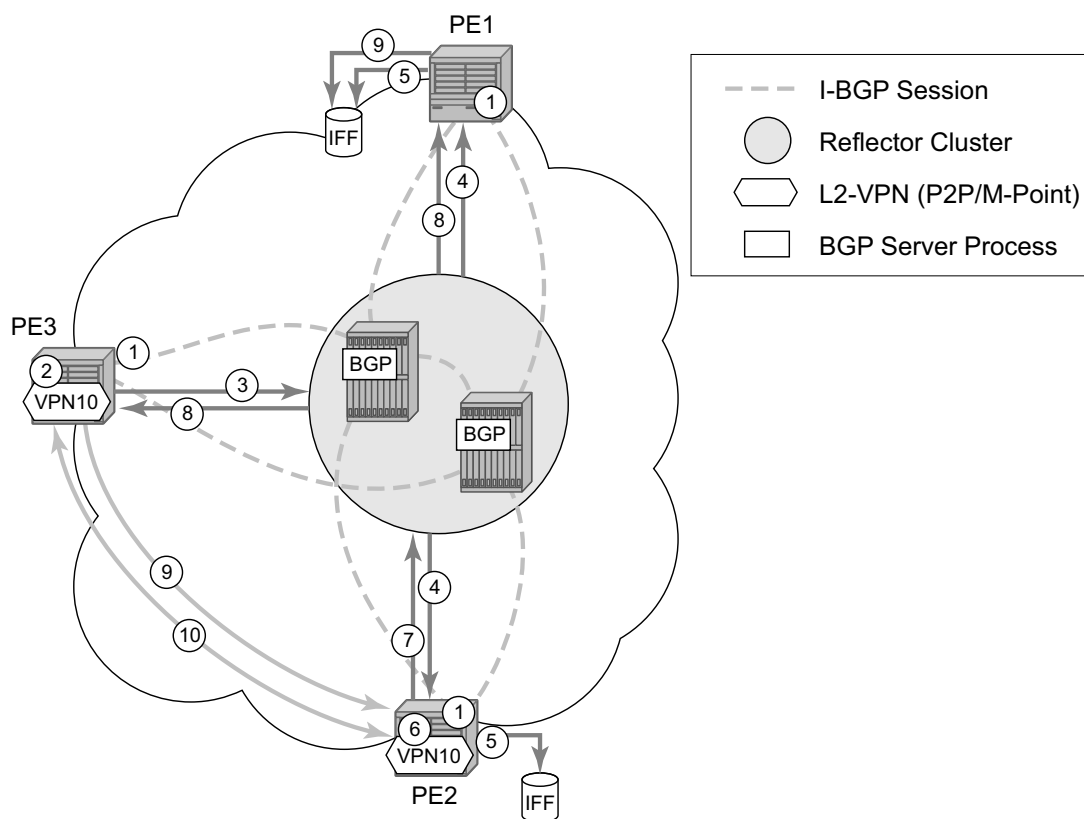
Each of the FEC fields are designed as a sub-TLV equipped with its own type and length providing support for new applications. To accommodate the BGP AD information model the following FEC formats are used:

- 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 76](#).



OSSG210

Figure 76: 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.

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.

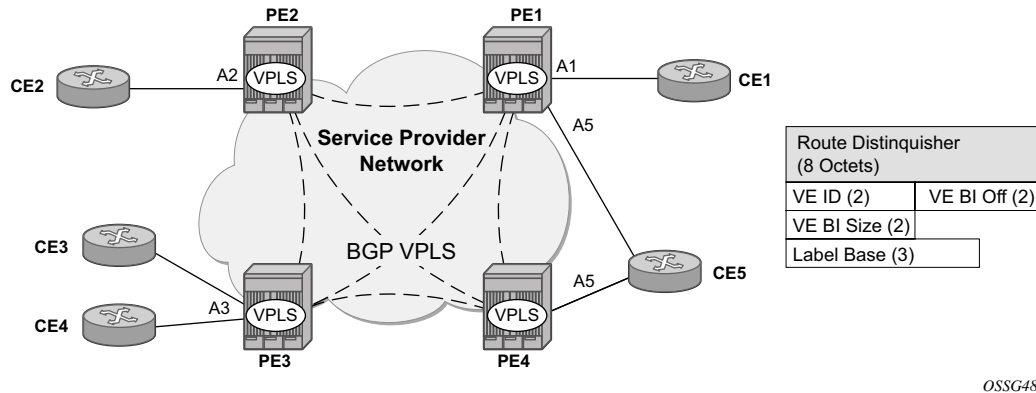


Figure 77: BGP VPLS Solution

Figure 77 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

- 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**.

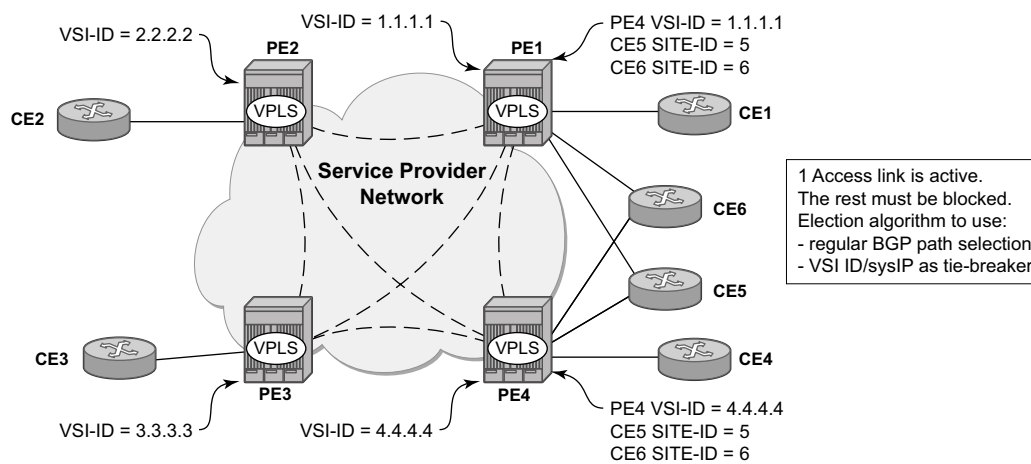
For BGP-AD VPLS, a BFD template is referenced from the pseudowire template context. To configure VCCV BFD for BGP-AD, use the command **config>service>vpls>bgp-ad>pw-template-binding>bfd-template** *name* and enable it using the command **config>service>vpls>bgp-ad>pw-template-binding>bfd-enable**.

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 78: BGP Multi-Homing for VPLS

Figure 78 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 79](#), 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 NLRI in [Figure 79](#).

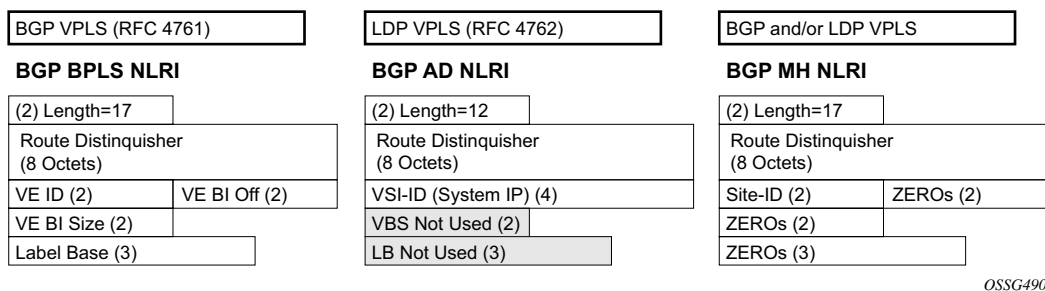


Figure 79: BGP MH-NLRI for VPLS Multi-Homing

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 80](#).




Figure 80: 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.
- When LDP multicast state distribution is enabled, generate PIM messages for LDP.

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 IOM2, 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

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. A

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.

Routed VPLS and I-VPLS

IES or VPRN IP Interface Binding

A standard IP interface within an existing IES or VPRN service context may be bound to a service name. Subscriber and group IP interfaces are not allowed to bind to a VPLS or I-VPLS service context or I-VPLS. **For the remainder of this section Routed VPLS and Routed I-VPLS will both be described as a VPLS service and differences will be pointed out where applicable.** A VPLS service only supports binding for a single IP interface.

While an IP interface may only be bound to a single VPLS service, the routing context containing the IP interface (IES or VPRN) may have other IP interfaces bound to other VPLS service contexts of the same type (all VPLS or all I-VPLS). In other words, Routed VPLS allows the binding of IP interfaces in IES or VPRN services to be bound to VPLS services and Routed I-VPLS allows of IP interfaces in IES or VPRN services to be bound to I-VPLS services.

Assigning a Service Name to a VPLS Service

When a service name is applied to any service context, the name and service ID association is registered with the system. A service name cannot be assigned to more than one service ID.

Special consideration is given to a service name that is assigned to a VPLS service that has the **configure>service>vpls>allow-ip-int-binding** command is enabled. If a name is applied to the VPLS service while the flag is set, the system will scan the existing IES and VPRN services for an IP interface that is bound to the specified service name. If an IP interface is found, the IP interface will be attached to the VPLS service associated with the name. Only one interface can be bound to the specified name.

If the **allow-ip-int-binding** command is not enabled on the VPLS service, the system will not attempt to resolve the VPLS service name to an IP interface. As soon as the **allow-ip-int-binding** flag is configured on the VPLS, the corresponding IP interface will be bound and become operational up. There is no need to toggle the **shutdown/no shutdown** command.

If an IP interface is not currently bound to the service name used by the VPLS service, no action is taken at the time of the service name assignment.

Service Binding Requirements

In the event that the defined service ID is created on the system, the system will check to ensure that the service type is VPLS. If the service type is not VPLS or I-VPLS, service creation will not be allowed and the service ID will remain undefined within the system.

If the created service type is VPLS, the IP interface will be eligible to enter the operationally up state.

Bound Service Name Assignment

In the event that a bound service name is assigned to a service within the system, the system will first check to ensure the service type is VPLS or I-VPLS. Secondly the system will ensure that the service is not already bound to another IP interface via the service ID. If the service type is not VPLS or I-VPLS or the service is already bound to another IP interface via the service ID, the service name assignment will fail.

In the event that a single VPLS Service ID and service name is assigned to two separate IP interfaces, the VPLS service will not be allowed to enter and be operational/up state.

Binding a Service Name to an IP Interface

An IP interface within an IES or VPRN service context may be bound to a service name at anytime. Only one interface can be bound to a service.

When an IP interface is bound to a service name and the IP interface is administratively up, the system will scan for a VPLS service context using the name and take the following actions:

- If the name is not currently in use by a service, the IP interface will be placed in an operationally down: Non-existent service name or inappropriate service type state.
- If the name is currently in use by a non-VPLS service or the wrong type of VPLS service, the IP interface will be placed in the operationally down: Non-existent service name or inappropriate service type state.
- If the name is currently in use by a VPLS service without the **allow-ip-int-binding** flag set, the IP interface will be placed in the operationally down: VPLS service **allow-ip-int-binding** flag not set state. There is no need to toggle the **shutdown/no shutdown** command.
- If the name is currently in use by a valid VPLS service and the **allow-ip-int-binding** flag is set, the IP interface will be eligible to be placed in the operationally up state depending on other operational criteria being met.

Bound Service Deletion or Service Name Removal

In the event that a VPLS service is deleted while bound to an IP interface, the IP interface will enter the 'Down: Non-existent svc-ID' operational state. If the IP interface was bound to the VPLS service name, the IP interface will enter the 'Down: Non-existent svc-name' operational state. No console warning is generated.

If the created service type is VPLS, the IP interface will be eligible to enter the operationally up state.

IP Interface Attached VPLS Service Constraints

Once a VPLS service has been bound to an IP interface through its service name, the service name assigned to the service cannot be removed or changed unless the IP interface is first unbound from the VPLS service name.

A VPLS service that is currently attached to an IP interface cannot be deleted from the system unless the IP interface is unbound from the VPLS service name.

The **allow-ip-int-binding** flag within an IP interface attached VPLS service cannot be reset. The IP interface must first be unbound from the VPLS service name to reset the flag.

IP Interface and VPLS Operational State Coordination

When the IP interface is successfully attached to a VPLS service, the operational state of the IP interface will be dependent upon the operational state of the VPLS service.

The VPLS service itself remains down until at least one virtual port (SAP, spoke SDP or mesh SDP) is operational.

IP Interface MTU and Fragmentation

The VPLS service is affected by two MTU values; port MTUs and the VPLS service MTU. The MTU on each physical port defines the largest Layer 2 packet (including all DLC headers) that may be transmitted out a port. The VPLS itself has a service level MTU that defines the largest packet supported by the service. This MTU does not include the local encapsulation overhead for each port (QinQ, Dot1Q, TopQ or SDP service delineation fields and headers) but does include the remainder of the packet. As virtual ports are created in the system, the virtual port cannot become operational unless the configured port MTU minus the virtual port service delineation overhead is greater than or equal to the configured VPLS service MTU. Thus, an operational virtual port is ensured to support the largest packet traversing the VPLS service. The service delineation overhead on each Layer 2 packet is removed before forwarding into a VPLS service. VPLS services do not support fragmentation and must discard any Layer 2 packet larger than the service MTU after the service delineation overhead is removed.

When an IP interface is associated with a VPLS service, the IP-MTU is based on either the administrative value configured for the IP interface or an operational value derived from VPLS service MTU. The operational IP-MTU cannot be greater than the VPLS service MTU minus 14 bytes.

- If the configured (administrative) IP-MTU is configured for a value greater than the normalized IP-MTU, based on the VPLS service-MTU, then the operational IP-MTU is reset to equal the normalized IP-MTU value (VPLS service MTU – 14 bytes).
- If the configured (administrative) IP-MTU is configured for a value less than or equal to the normalized IP-MTU, based on the VPLS service-MTU, then the operational IP-MTU is set to equal the configured (administrative) IP-MTU value.

Unicast IP Routing into a VPLS Service

The VPLS service MTU and the IP interface MTU parameters may be changed at anytime.

ARP and VPLS FIB Interactions

Two address-oriented table entries are used when routing into a VPLS service. On the routing side, an ARP entry is used to determine the destination MAC address used by an IP next-hop. In the case where the destination IP address in the routed packet is a host on the local subnet represented by the VPLS instance, the destination IP address itself is used as the next-hop IP address in the ARP cache lookup. If the destination IP address is in a remote subnet that is reached by another router attached to the VPLS service, the routing lookup will return the local IP address on the VPLS service of the remote router will be returned. If the next-hop is not currently in the ARP cache, the system will generate an ARP request to determine the destination MAC address associated with the next-hop IP address. IP routing to all destination hosts associated with the next-hop IP address stops until the ARP cache is populated with an entry for the next-hop. The ARP cache may be populated with a static ARP entry for the next-hop IP address. While dynamically populated ARP entries will age out according to the ARP aging timer, static ARP entries never age out.

The second address table entry that affects VPLS routed packets is the MAC destination lookup in the VPLS service context. The MAC associated with the ARP table entry for the IP next-hop may or may not currently be populated in the VPLS Layer 2 FIB table. While the destination MAC is unknown (not populated in the VPLS FIB), the system will flood all packets destined to that MAC (routed or bridged) to all virtual ports within the VPLS service context. Once the MAC is known (populated in the VPLS FIB), all packets destined to the MAC (routed or bridged) will be targeted to the specific virtual port where the MAC has been learned. As with ARP entries, static MAC entries may be created in the VPLS FIB. Dynamically learned MAC addresses are allowed to age out or be flushed from the VPLS FIB while static MAC entries always remain associated with a specific virtual port. Dynamic MACs may also be relearned on another VPLS virtual port than the current virtual port in the FIB. In this case, the system will automatically move the MAC FIB entry to the new VPLS virtual port.

Routed VPLS Specific ARP Cache Behavior

In typical routing behavior, the system uses the IP route table to select the egress interface and then at the egress forwarding engine, an ARP entry is used forward the packet to the appropriate Ethernet MAC. With routed VPLS, the egress IP interface may be represented by multiple egress forwarding engine (wherever the VPLS service virtual ports exists).

In order to optimize routing performance, the ingress forwarding engine processing has been augmented to perform an ingress ARP lookup in order to resolve which VPLS MAC address the IP frame must be routed towards. This MAC address may be currently known or unknown within the VPLS FIB. If the MAC is unknown, the packet is flooded by the ingress forwarding engine to all egress forwarding engines where the VPLS service exists. When the MAC is known on a virtual port, the ingress forwarding engine forwards the packet to the proper egress forwarding engine. [Table 15](#) describes how the ARP cache and MAC FIB entry states interact at ingress and [Table 16](#) describes the corresponding egress behavior.

Table 15: Ingress Routed to VPLS Next-Hop Behavior

| Next-Hop ARP Cache Entry | Next-Hop MAC FIB Entry | Ingress Behavior |
|---------------------------|------------------------|---|
| ARP Cache Miss (No Entry) | Known or Unknown | Flood to all egress forwarding engines associated with the VPLS/I-VPLS context. |
| | Unknown | Flood to all egress forwarding engines associated with the VPLS/I-VPLS context |
| | Unknown | Flood to all egress forwarding engines associated with the VPLS for forwarding out all VPLS /I-VPLS virtual ports |

Table 16: Egress Routed VPLS Next-Hop Behavior

| Next-Hop ARP Cache Entry | Next-Hop MAC FIB Entry | Egress Behavior |
|--|------------------------|---|
| ARP Cache Miss (No Entry) ² | Known | No ARP entry. The MAC address is unknown and the ARP request is flooded out of all virtual ports of the VPLS/I-VPLS instance |
| | Unknown | Request control engine ARP processing ARP request transmitted out all virtual port associated with the VPLS/I-VPLS service. Only the first egress forwarding engine ARP processing request triggers egress ARP request. |

Table 16: Egress Routed VPLS Next-Hop Behavior (Continued)

| Next-Hop ARP Cache Entry | Next-Hop MAC FIB Entry | Egress Behavior |
|-----------------------------|---------------------------|--|
| ARP Cache Hit | Known | Forward out specific egress VPLS/I-VPLS virtual port where MAC has been learned. |
| | Unknown | Flood to all egress VPLS/I-VPLS virtual ports on forwarding engine. |

The allow-ip-int-binding VPLS Flag

The **allow-ip-int-binding** 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-binding** flag is set (routing support enabled) on a VPLS/I-VPLS service. SAPs within the service can be created on standard Ethernet, HSMDA.

Routed VPLS SAPs Only Supported on FP2 (or later) Based Systems or IOM/IMM

The Ethernet ports must be populated on a FP2 or FP3 system IOMs in order for the routing enabled VPLS SAPs to be created.

Network Ports Restricted to FP2-Based Systems or IOMs

When at least one VPLS context is configured with the **allow-ip-int-binding** flag set, all ports within the system defined as mode network must be on an FP2 or greater forwarding plane. If one or more network ports are on an FP1 based forwarding plane, the **allow-ip-int-binding** flag cannot be set in a VPLS service context. Once the **allow-ip-int-binding** flag is set on a VPLS service context, a port on an FP1 based forwarding plane cannot be placed in mode network.

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-binding** 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.
- SDPs used in spoke or mesh SDP bindings cannot be configured as GRE.
- The VPLS service type cannot be B-VPLS or M-VPLS.
- MVR from Routed VPLS and to another SAP is not supported.
- Enhanced and Basic Subscriber Management (BSM) features.
- Network domain on SDP bindings.
- Per Service Hashing not supported
- No BGP-AD
- No BGP-VPLS
- IOM3+ cards only

Note: IES/VP RN 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 when the source of the multicast stream is on the IP side of the routed VPLS service, and the multicast traffic is either flooded in the VPLS service or sent to IGMP clients. The IP interface supports the configuration of PIM and IGMP. When IGMP is configured on the IP interface, enabling IGMP snooping in the VPLS service is optional. If IGMP snooping is enabled, it is mandatory 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.




Figure 81: IPv4 Multicast with a Router VPLS service

An example scenario is shown in [Figure 81](#). The IP multicast traffic entering the system is replicated to IP interfaces, two of which are part of an IES routed VPLS service. The IP multicast traffic is sent only to those SAPs/SDPs from which a corresponding IGMP join has been received as igmp-snooping is enabled in the VPLS services.

It is possible to configure PIM on the IP interface and have neighboring downstream PIM routers connecting via the VPLS, however, any multicast traffic will be flooded in the VPLS service.

If a multicast source was connected by a SAP/SDP into the VPLS service, any multicast traffic would be replicated within the VPLS service but would be dropped at the routed VPLS IP interface.

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.

7450 Mixed Mode Chassis

The mixed mode on the 7450 that allows 7750 based IOM3s to be populated and operational in a 7450 chassis supports routed VPLS as long as all the forwarding plane and port type restrictions are observed.

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.
-

Routed VPLS Supported Routing Related Protocols

The following protocols are supported on IP interfaces bound to a VPLS service:

- BGP
- OSPF
- ISIS
- PIM
- IGMP
- BFD
- VRRP

- ARP
 - DHCP Relay
-

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 various 7450 ESS 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 7450 ESS VPLS service:

- Ethernet null
 - Ethernet Dot1q
 - Ethernet QinQ
 - SONET/SDH BCP-null
 - SONET/SDH BCP-dot1q
-

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:

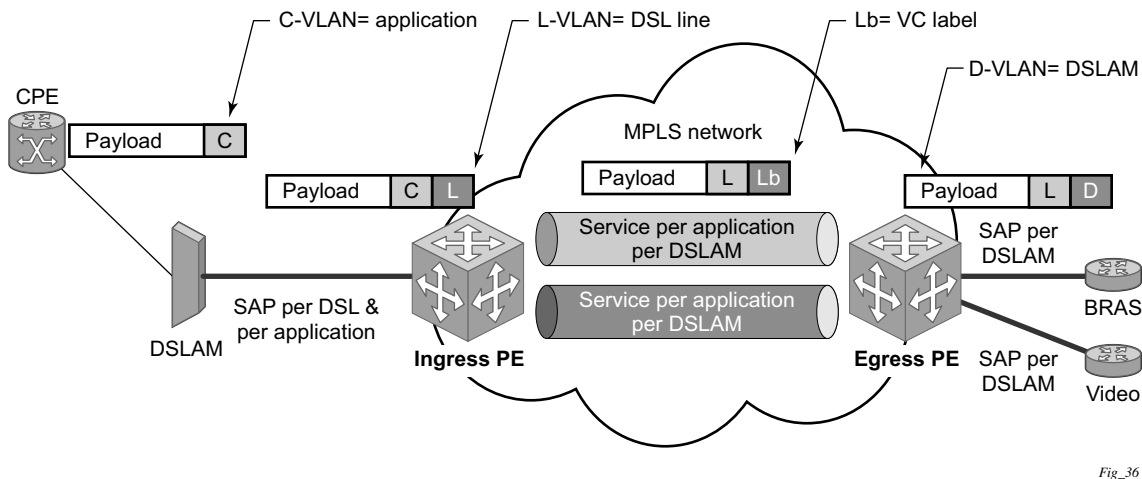
1. Null encapsulation defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP.
2. Dot1q encapsulation defined on ingress — Only first label is considered.
3. QinQ encapsulation defined on ingress— Both labels are considered.
Note that the SAP can be defined 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 82: 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 83](#).



OSSG492

Figure 83: 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:

Service Auto-Discovery using Multiple VLAN Registration Protocol (MVRP)

```
mrp
    no shutdown
    mmrp
        shutdown
    mvrp
        no shutdown
sap 1/1/1:0
    mrp mvrp
        no shutdown
sap 2/1/2:0
    mrp mvrp
        no shutdown
sap 3/1/10:0
    mrp mvrp
        no shutdown
vpls-group 1
    service-range 100-2000
    vpls-template-binding Autovpls1
    vpls-sap-template-binding Autosap1
    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 84](#).




Figure 84: 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.

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 17: 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.
- Temporary flooding behavior is supported only on SAPs located on IOM3s for all chassis modes. If IOM1 or IOM2 are involved, the flooding will not work on related SAP or PW endpoints.

VPLS E-Tree Services

This section describes the following topics:

- [VPLS E-Tree Services Overview on page 533](#)
 - [Leaf-ac and Root-ac SAPs on page 534](#)
 - [Leaf-ac and Root-ac SDP Binds on page 535](#)
 - [Root-leaf-tag SAPs on page 535](#)
 - [Root-leaf-tag SDP Binds on page 536](#)
 - [Interaction between VPLS E-Tree Services and Other Features on page 537](#)
-

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 85](#) 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.

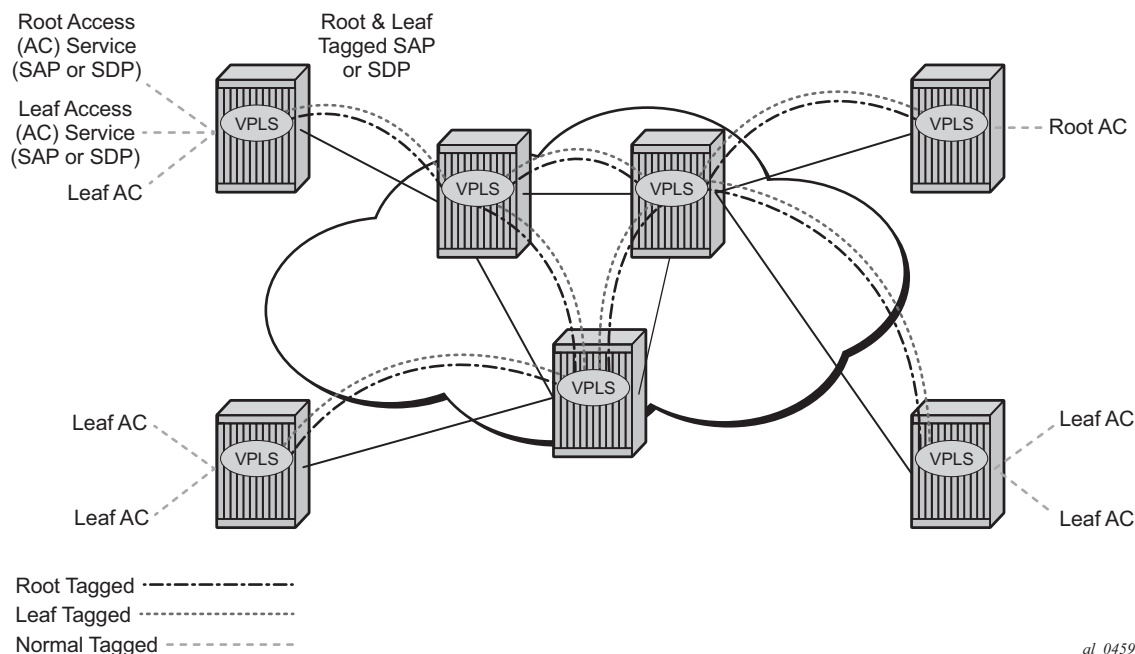


Figure 85: E-Tree Service

Leaf-ac and Root-ac SAPs

Figure 86 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.

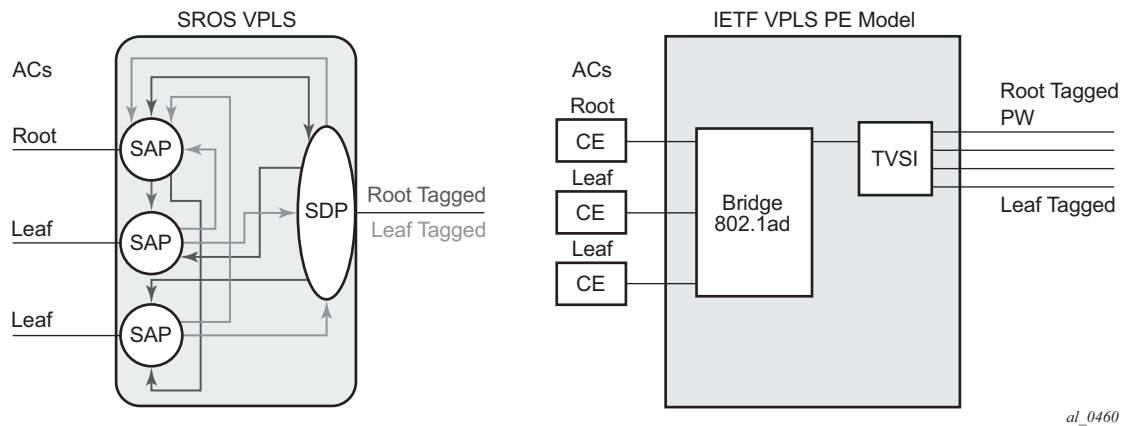


Figure 86: Mapping PE Model to 7x50 VPLS Service

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 87](#) 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
```

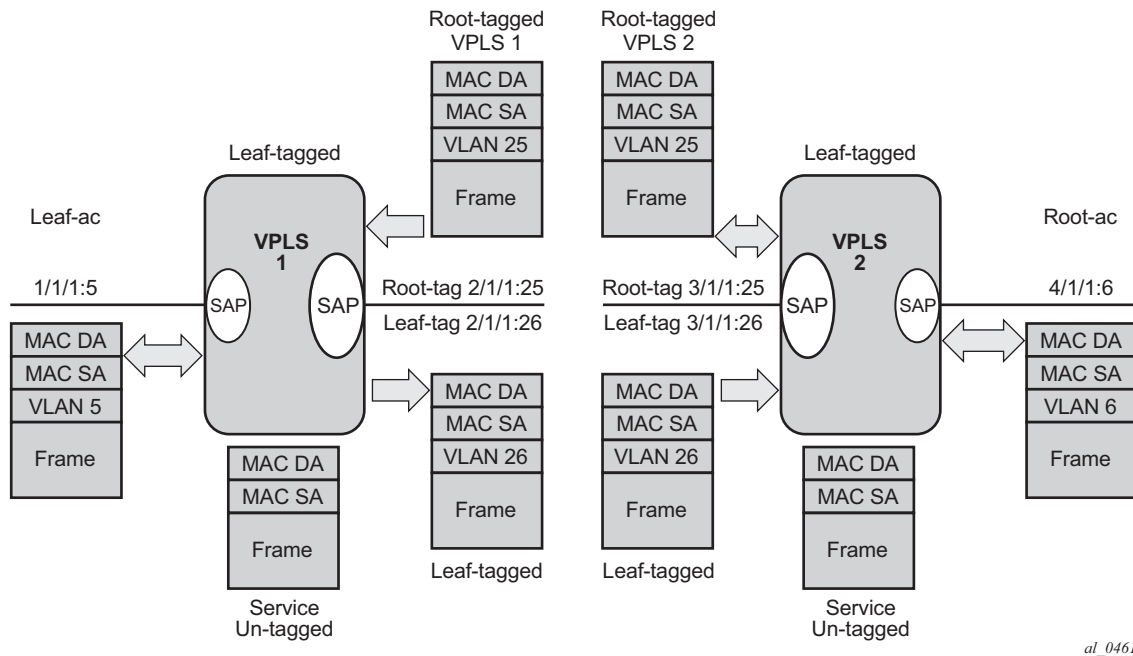


Figure 87: Leaf and Root Tagging Dot1q

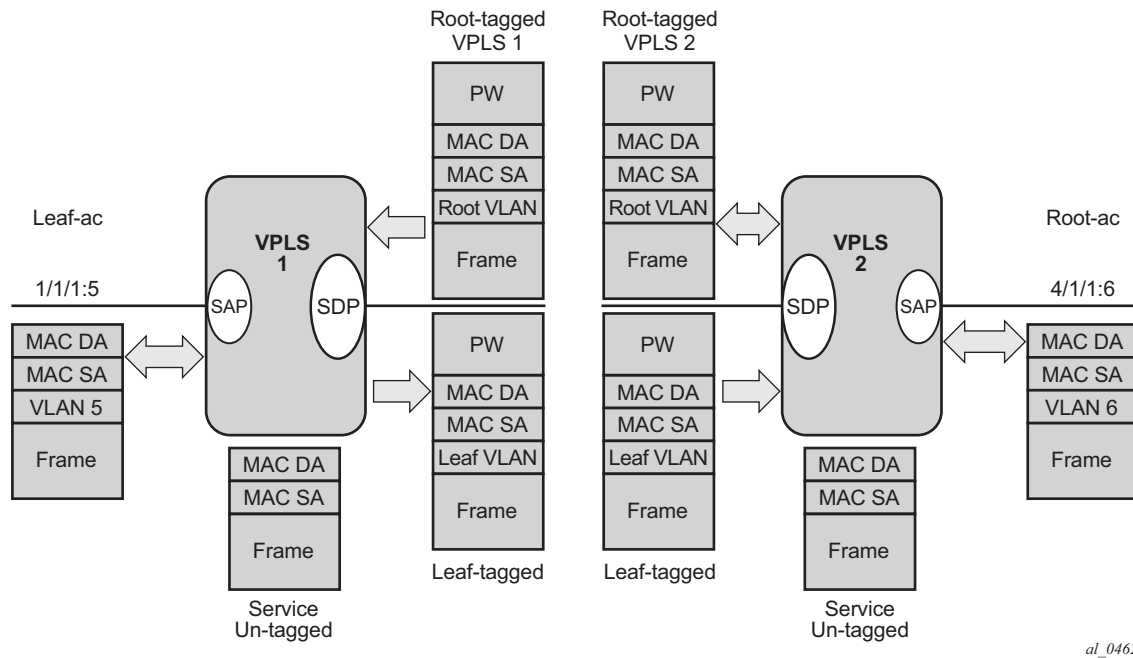
The root-leaf-tag SAP performs all of the operations for egress and ingress traffic for both tags (root and leaf):

- 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 88 illustrates the behavior when leaf-ac or root-ac interfaces exchange traffic over a root-leaf-tag SDP-binding.



al_0462

Figure 88: Leaf and Root Tagging PW

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 540](#)
- [Common Configuration Tasks on page 542](#)
 - [Configuring VPLS Components on page 543](#)
 - [Creating a VPLS Service on page 545](#)
 - [Configuring a VPLS SAP on page 556](#)
 - [Local VPLS SAPs on page 556](#)
 - [Distributed VPLS SAPs on page 557](#)
 - [Configuring SAP Subscriber Management Parameters on page 568](#)
 - [MSTP Control over Ethernet Tunnels on page 569](#)
 - [Configuring SDP Bindings on page 570](#)
- [Configuring VPLS Redundancy on page 583](#)
 - [Creating a Management VPLS for SAP Protection on page 583](#)
 - [Creating a Management VPLS for Spoke SDP Protection on page 586](#)
 - [Configuring BGP VPLS on page 606](#)
 - [Configuring Selective MAC Flush on page 594](#)
 - [Configuring Multi-Chassis Endpoints on page 595](#)
- [Configuring Policy-Based Forwarding for Deep Packet Inspection \(DPI\) in VPLS on page 609](#)
- [Configuring VPLS E-Tree Services on page 612](#)
- [Service Management Tasks on page 613](#)
 - [Modifying VPLS Service Parameters on page 613](#)
 - [Modifying Management VPLS Parameters on page 614](#)
 - [Deleting a VPLS Service on page 616](#)
 - [Disabling a VPLS Service on page 616](#)
 - [Re-Enabling a VPLS Service on page 617](#)

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 434](#))
5. Enable service

Configuring VPLS Components

Use the CLI syntax displayed below to configure the following entities:

- [Configuring Egress Multicast Groups on page 544](#)
- [Creating a VPLS Service on page 545](#)
 - [Enabling MAC Move on page 548](#)
- [Configuring a VPLS SAP on page 556](#)
 - [Local VPLS SAPs on page 556](#)
 - [Distributed VPLS SAPs on page 557](#)
 - [Configuring SAP-Specific STP Parameters on page 559](#)
 - [STP SAP Operational States on page 563](#)
 - [Configuring VPLS SAPs with Split Horizon on page 565](#)
 - [Configuring SAP Subscriber Management Parameters on page 568](#)
 - [Configuring Overrides on Service SAPs on page 571](#)
- [Configuring SDP Bindings on page 570](#)
 - [Mesh SDP on page 572](#)
 - [Spoke SDP on page 573](#)
- [Configuring VPLS Redundancy on page 583](#)

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 mmrp mac

| SAP/SDP | MAC Address | Registered | Declared |
|---------------|-------------------|------------|----------|
| sap:1/5/1:100 | 01:1e:83:00:00:0b | No | Yes |
| sdp:3101:100 | 01:1e:83:00:00:0b | No | Yes |
| sdp:3201:100 | 01:1e:83:00:00:0b | Yes | Yes |

*A:PE-C# show service id 100 sdp 3201:100 mrp

 Sdp Id 3201:100 MRP Information

| | | | |
|------------------|-------------|-------------------|------------|
| Join Time | : 0.2 secs | Leave Time | : 3.0 secs |
| Leave All Time | : 10.0 secs | Periodic Time | : 1.0 secs |
| Periodic Enabled | : false | | |
| Rx Pdus | : 7 | Tx Pdus | : 23 |
| Dropped Pdus | : 0 | | |
| Rx New Event | : 0 | Rx Join-In Event | : 6 |
| Rx In Event | : 0 | Rx Join Empty Evt | : 1 |
| Rx Empty Event | : 0 | Rx Leave Event | : 0 |
| Tx New Event | : 0 | Tx Join-In Event | : 4 |
| Tx In Event | : 0 | Tx Join Empty Evt | : 19 |
| Tx Empty Event | : 0 | Tx Leave Event | : 0 |

SDP MMRP Information

| MAC Address | Registered | Declared |
|-------------------|------------|----------|
| 01:1e:83:00:00:0b | Yes | Yes |

Number of MACs=1 Registered=1 Declared=1

 *A:PE-C#

*A:PE-C# show service id 100 mfib

=====

Multicast FIB, Service 100

| Source Address | Group Address | Sap/Sdp Id | Svc Id | Fwd/Blk |
|----------------|-------------------|--------------|--------|---------|
| * | 01:1E:83:00:00:0B | sdp:3201:100 | Local | Fwd |

Number of entries: 1

=====

*A:PE-C#

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 550](#)
- [Mode on page 551](#)
- [Bridge Priority on page 551](#)
- [Max Age on page 552](#)
- [Forward Delay on page 552](#)
- [Hello Time on page 553](#)
- [MST Instances on page 554](#)
- [MST Max Hops on page 554](#)
- [MST Name on page 554](#)
- [MST Revision on page 554](#)

STP always uses the locally configured values for the first three parameters (Admin State, Mode and Priority).

For the parameters Max Age, Forward Delay, Hello Time and Hold Count, the locally configured values are only used when this bridge has been elected root bridge in the STP domain, otherwise the values received from the root bridge are used. The exception to this rule is: when STP is running in RSTP mode, the Hello Time is always taken from the locally configured parameter. The other parameters are only used when running mode MSTP.

Bridge STP Admin State

The administrative state of STP at the VPLS level is controlled by the shutdown command.

When STP on the VPLS is administratively disabled, any BPDUs are forwarded transparently through the 7450 ESS. When STP on the VPLS is administratively enabled, but the administrative state of a SAP or spoke SDP is down, BPDUs received on such a SAP or spoke SDP are discarded.

CLI Syntax: `config>service>vpls service-id# stp
no shutdown`

Mode

To be compatible with the different iterations of the IEEE 802.1D standard, the 7450 ESS supports 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 434](#) 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 562](#)).

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 552](#).

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

Use the following CLI syntax to create:

- [Local VPLS SAPs on page 556](#)
- [Distributed VPLS SAPs on page 557](#)

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 570](#).

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 559](#)
 - [SAP Virtual Port Number on page 560](#)
 - [SAP Priority on page 560](#)
 - [SAP Path Cost on page 561](#)
 - [SAP Edge Port on page 561](#)
 - [SAP Auto Edge on page 562](#)
 - [SAP Link Type on page 562](#)
-

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. These are the values used for CIST when running MSTP.

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 560](#) for details on the virtual port number.

STP computes the actual SAP priority by taking the configured priority value and masking out the lower four bits. The result is the value that is stored in the SAP priority parameter. For example, if a value of 0 was entered, masking out the lower 4 bits would result in a parameter value of 0. If a value of 255 was entered, the result would be 240.

The default value for SAP priority is 128. This parameter can be modified within a range of 0 to 255, 0 being the highest priority. Masking causes the values actually stored and displayed to be 0 to 240, in increments of 16.

CLI Syntax: `config>service>vpls>sap>stp#
priority stp-priority
Range: 0 to 255 (240 largest value, in increments of 16)
Default: 128
Restore Default: no priority`

SAP Path Cost

The SAP path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP. When BPDUs are sent out other egress SAPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs are controlled by complex queuing dynamics, in the 7450 ESS, 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 552](#)). 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 561](#)).

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 563](#)
 - [Operationally Discarding on page 563](#)
 - [Operationally Learning on page 563](#)
 - [Operationally Forwarding on page 564](#)
-

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 552](#).

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 BPDU Encapsulation State

IEEE 802.1d (referred as Dot1d) and Cisco's per VLAN Spanning Tree (PVST) BPDU encapsulations are supported on a per SAP basis. STP is associated with a VPLS service like PVST is associated per VLAN. The main difference resides in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDU.

The following table shows differences between Dot1d and PVST Ethernet BPDU encapsulations based on the interface encap-type field:

Each SAP has a Read-Only operational state that shows which BPDU encapsulation is currently active on the SAP. The states are:

- **Dot1d** — This state specifies that the switch is currently sending IEEE 802.1d standard BPDUs. The BPDUs are tagged or non-tagged based on the encapsulation type of the egress interface and the encapsulation value defined in the SAP. A SAP defined on an interface with encapsulation type Dot1q continues in the dot1d BPDU encapsulation state until a PVST encapsulated BPDU is received in which case, the SAP will convert to the PVST encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged if the interface encapsulation type is defined as Dot1q. PVST BPDUs will be silently discarded if received when the SAP is on an interface defined with encapsulation type null.
- **PVST** — This state specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The SAP continues in the PVST BPDU encapsulation state until a dot1d encapsulated BPDU is received, in which case, the SAP reverts to the dot1d encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged with the encapsulation value defined for the SAP. PVST BPDUs are silently discarded if received when the SAP is on an interface defined with a null encapsulation type.

Dot1d is the initial and only SAP BPDU encapsulation state for SAPs defined on Ethernet interface with encapsulation type set to null.

Each transition between encapsulation types optionally generates an alarm that can be logged and optionally transmitted as an SNMP trap.

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 MAC Learning Protection

To configure MAC learning protection, configure split horizon, MAC protection, and SAP parameters.

The following example displays a VPLS configuration with split horizon enabled:

```
A:ALA-48>config>service>vpls# info
-----
description "IMA VPLS"
split-horizon-group "DSL-group1" create
    restrict-protected-src
    restrict-unprotected-dst
exit
mac-protect
    mac ff:ff:ff:ff:ff:ff
exit
sap 1/1/9:0 create
    ingress
        scheduler-policy "SLA1"
        qos 100 shared-queuing
    exit
    egress
        scheduler-policy "SLA1"
        filter ip 10
    exit
    restrict-protected-src
    arp-reply-agent
    host-connectivity-verify source-ip 143.144.145.1
exit
...
-----
A:ALA-48>config>service>vpls#
```

Applying an Egress Multicast Group to a VPLS Service SAP

Use the following CLI syntax to apply an egress multicast group to a VPLS service SAP:

CLI Syntax: `config>service>vpls service-id [customer customer-id] [vpn vpn-id] [mvpls]
 sap sap-id [split-horizon-group group-name]
 egress
 multicast-group group-name`

The following example displays a VPLS configuration with egress multicast group:

```
A:ALA-48>config>service>vpls# info
-----
description "VPLS with split horizon for DSL"
split-horizon-group "DSL-group1" create
    description "Split horizon group for DSL"
exit
stp
    shutdown
exit
sap 1/1/4:200 split-horizon-group "DSL-group1" create
    description "SAP for residential bridging"
exit
sap 1/1/3:100 split-horizon-group "DSL-group1" create
    description "SAP for residential bridging"
    egress
        multicast-group "vpls-emg-1"
exit
no shutdown
-----
A:ALA-48>config>service>vpls#
```

Configuring SAP Subscriber Management Parameters

Use the following CLI syntax to configure subscriber management parameters on a VPLS service SAP. The policies and profiles that are referenced in the **def-sla-profile**, **def-sub-profile**, **non-sub-traffic**, and **sub-ident-policy** commands must already be configured in the **config>subscriber** context.

CLI Syntax:

```
config>service>vpls service-id
  sap sap-id [split-horizon-group group-name]
  sub-sla-mgmt
    def-sla-profile default-sla-profile-name
    def-sub-profile default-subscriber-profile-name
    mac-da-hashing
    multi-sub-sap [number-of-sub]
    no shutdown
    single-sub-parameters
      non-sub-traffic sub-profile sub-profile-name sla-
        profile sla-profile-name [subscriber sub-ident-
          string]
      profiled-traffic-only
      sub-ident-policy sub-ident-policy-name
```

The following example displays a subscriber management configuration:

```
A:ALA-48>config>service>vpls#
-----
      description "Local VPLS"
      stp
        shutdown
      exit
      sap 1/2/2:0 create
        description "SAP for local service"
        sub-sla-mgmt
          def-sla-profile "sla-profile1"
          sub-ident-policy "SubIdent1"
        exit
      exit
      sap 1/1/5:0 create
        description "SAP for local service"
      exit
      no shutdown
-----
A:ALA-48>config>service>vpls#
```

MSTP Control over Ethernet Tunnels

When MSTP is used to control VLANs, a range of VLAN IDs is normally used to specify the VLANs to be controlled.

If an Ethernet tunnel SAP is to be controlled by MSTP, the Ethernet Tunnel SAP ID needs to be within the VLAN range specified under the mst-instance.

```
vpls 400 customer 1 m-vpls create
    stp
        mode mstp
        mst-instance 111 create
            vlan-range 1-100
        exit
        mst-name "abc"
        mst-revision 1
        no shutdown
    exit
    sap 1/1/1:0 create // untagged
    exit
    sap eth-tunnel-1 create
    exit
    no shutdown
exit
vpls 401 customer 1 create
    stp
        shutdown
    exit
    sap 1/1/1:12 create
    exit
    sap eth-tunnel-1:12 create
        // Ethernet tunnel SAP ID 12 falls within the VLAN
        // range for mst-instance 111
    eth-tunnel
        path 1 tag 1000
        path 8 tag 2000
    exit
    exit
    no shutdown
exit
```

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 582](#)).

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.

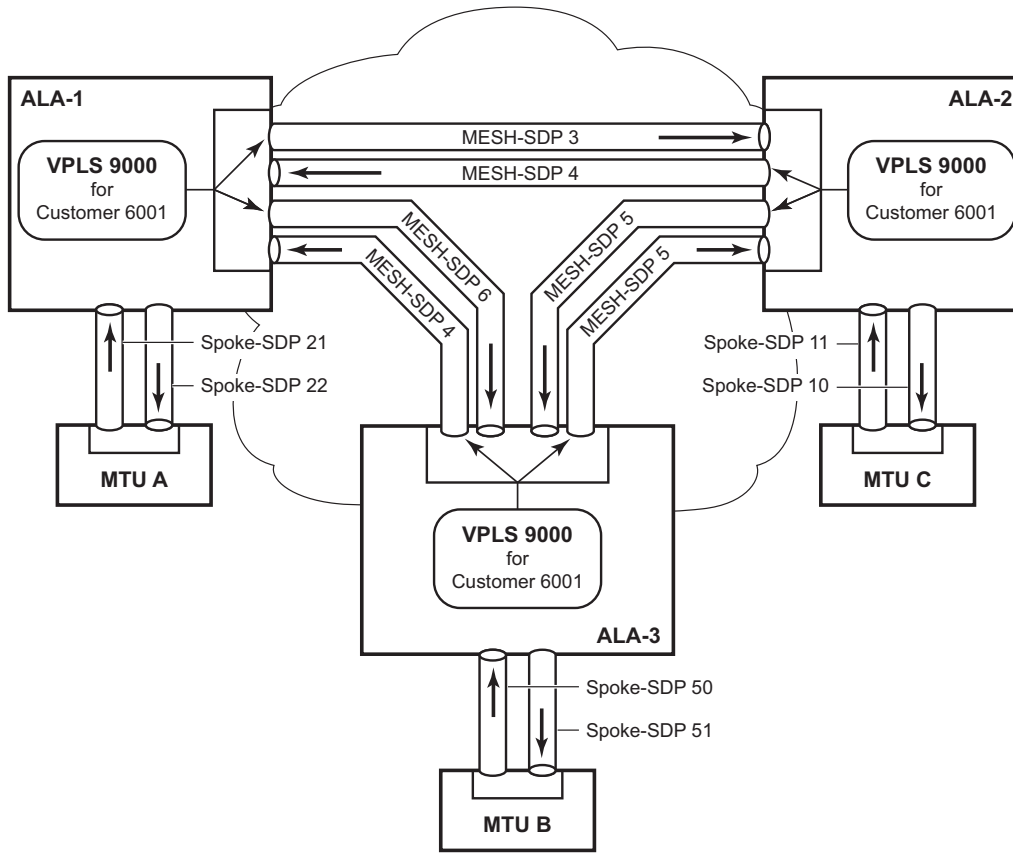
A VC-ID can be specified with the SDP-ID. The VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer SRs on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

[Figure 89](#) 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 89: 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-horizon-group group-name]
    egress
        filter {ip ip-filter-id|mac mac-filter-id}
        vc-label egress-vc-label
    ingress
        filter {ip ip-filter-id|mac mac-filter-id}
        vc-label ingress-vc-label
    limit-mac-move[non-blockable]
    vlan-vc-tag 0..4094
    no shutdown
    static-mac ieee-address
    stp
        path-cost stp-path-cost
        priority stp-priority
        no shutdown
    vlan-vc-tag [0..4094]
```

The following 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 575](#)
 - [Spoke SDP Virtual Port Number on page 576](#)
 - [Spoke SDP Priority on page 576](#)
 - [Spoke SDP Path Cost on page 577](#)
 - [Spoke SDP Edge Port on page 577](#)
 - [Spoke SDP Auto Edge on page 578](#)
 - [Spoke SDP Link Type on page 578](#)
-

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 576](#) 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 552](#)). 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 false (see [Spoke SDP Edge Port on page 577](#)).

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 579](#)
 - [Operationally Discarding on page 579](#)
 - [Operationally Learning on page 579](#)
 - [Operationally Forwarding on page 580](#)
-

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 552](#).

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 18](#) shows differences between dot1D and PVST Ethernet BPDUs encapsulations based on the interface encap-type field:

Table 18: 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 BPDU | 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 583](#)
 - [Creating a Management VPLS for Spoke SDP Protection on page 586](#)
-

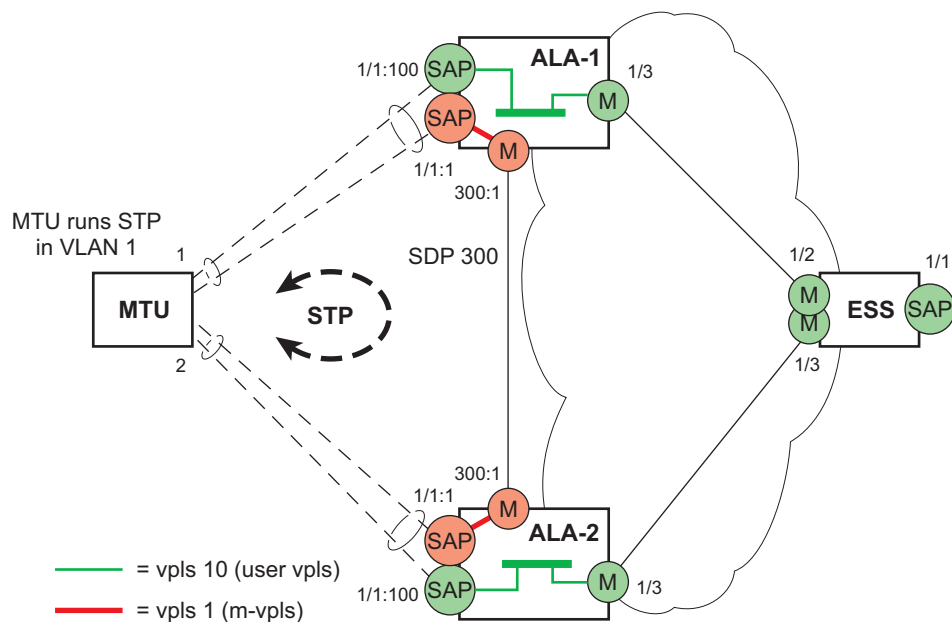
Creating a Management VPLS for SAP Protection

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 90](#). 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 454](#) 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.



OSSG047

Figure 90: Example Configuration for Protected VPLS SAP

Use the following CLI syntax to create a management VPLS:

CLI Syntax: `config>service# sdp sdp-id mpls create`
`far-end ip-address`
`lsp lsp-name`
`no shutdown`

CLI Syntax: `vpls service-id customer customer-id [m-vpls] create`
`description description-string`
`sap sap-id create`
`managed-vlan-list`
`range vlan-range`
`mesh-sdp sdp-id:vc-id create`
`stp`
`no shutdown`

The following example displays a VPLS configuration:

```
*A:ALA-1>config>service# info
-----
...
sdp 300 mpls create
  far-end 10.0.0.20
  lsp "toALA-A2"
  no shutdown
exit
vpls 1 customer 1 m-vpls create
  sap 1/1/1:1 create
```

```
        managed-vlan-list
        range 100-1000
    exit
exit
mesh-sdp 300:1 create
exit
stp
exit
no shutdown
exit
...
-----
*A:ALA-1>config>service#
```

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 91](#). 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 556](#) 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 589](#)).
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.

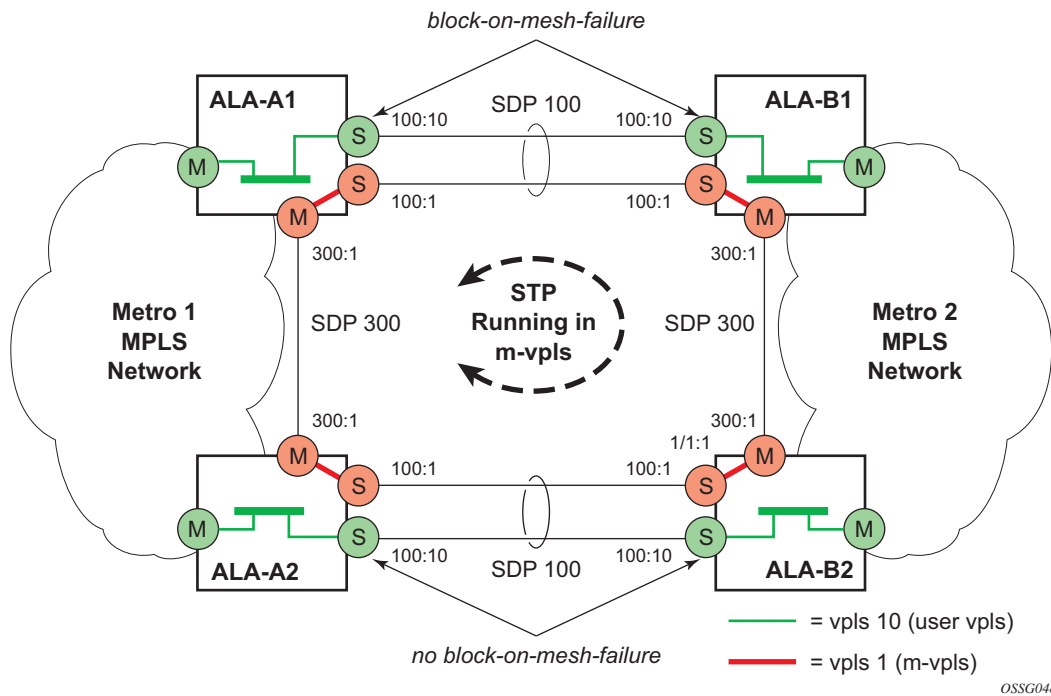


Figure 91: Example Configuration for Protected VPLS Spoke SDP

Use the following CLI syntax to create a management VPLS for spoke SDP protection:

CLI Syntax: `config>service# sdp sdp-id mpls create
far-end ip-address
lsp lsp-name
no shutdown`

CLI Syntax: `vpls service-id customer customer-id [m-vpls] create
description description-string
mesh-sdp sdp-id:vc-id create
spoke-sdp sdp-id:vc-id create
stp
no shutdown`

The following example 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.

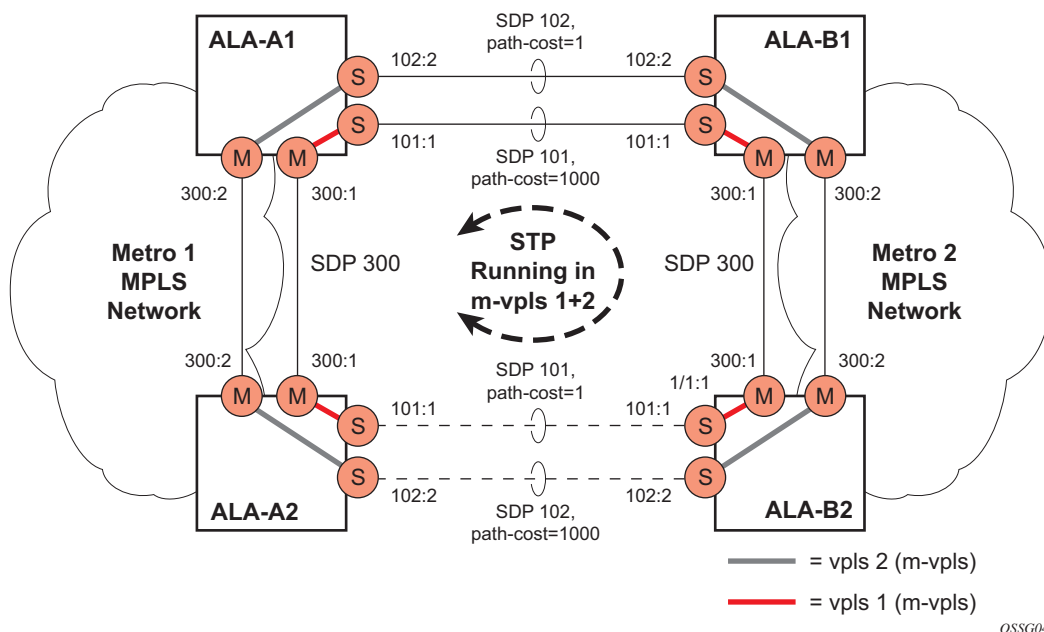


Figure 92: Example Configuration for Load Balancing Across Two Protected VPLS Spoke SDPs

Use the following CLI syntax to create a load balancing across two management VPLS instances:

CLI Syntax: `config>service# sdp sdp-id mpls create
far-end ip-address
lsp lsp-name
no shutdown`

CLI Syntax: `vpls service-id customer customer-id [m-vpls] create
description description-string
mesh-sdp sdp-id:vc-id create
spoke-sdp sdp-id:vc-id create
stp
path-cost
stp
no shutdown`

Note: the STP path costs in each peer node should be reversed.

The following example 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#
```

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
        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-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 457](#) 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 BGP Auto-Discovery

This section provides important information to explain the different configuration options used to populate the required BGP AD and generate the LDP generalized pseudowire-ID FEC fields. There are a large number of configuration options that are available with this feature. Not all these configurations option are required to start using BGP AD. At the end of this section, it will be apparent that a very simple configuration will automatically generate the required values used by BGP and LDP. In most cases, deployments will provide full mesh connectivity between all nodes across a VPLS instance. However, capabilities are available to influence the topology and build hierarchies or hub and spoke models.

Configuration Steps

Using [Figure 93](#), assume PE6 was previously configured with VPLS 100 as indicated by the configurations lines in the upper right. The BGP AD process will commence after PE134 is configured with the VPLS 100 instance as shown in the upper left. This shows a very basic and simple BGP AD configuration. The minimum requirement for enabling BGP AD on a VPLS instance is configuring the VPLS-ID and point to a pseudowire template.

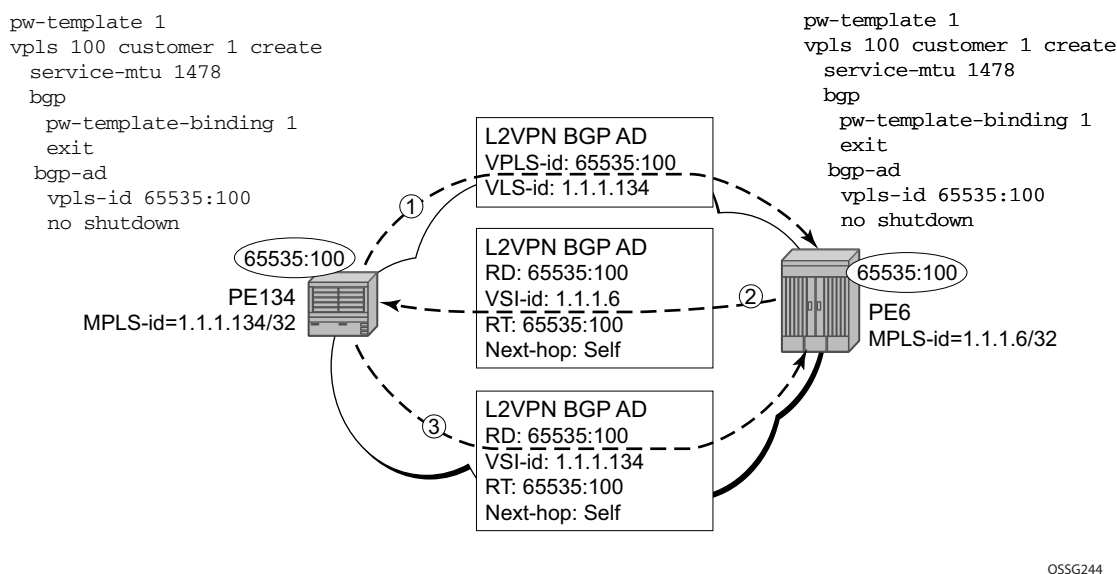


Figure 93: BGP AD Configuration Example

In many cases, VPLS connectivity is based on a pseudowire mesh. To reduce the configuration requirement, the BGP values can be automatically generated using the VPLS-ID and the MPLS router-ID. By default, the lower six bytes of the VPLS-ID are used to generate the RD and the RT

values. The VSI-ID value is generated from the MPLS router-ID. All of these parameters are configurable and can be coded to suit requirements and build different topologies.

```
PE134>config>service>vpls>bgp-ad#
[no] shutdown - Administratively enable/disable BGP auto-discovery
vpls-id - Configure VPLS-ID
vsi-id + Configure VSI-id
```

Figure 94: BGP-AD CLI Command Tree

A helpful command displays the service information, the BGP parameters and the SDP bindings in use. When the discovery process is completed successfully each endpoint will have an entry for the service.

```
PE134># show service l2-route-table
=====
Services: L2 Route Information - Summary Service
=====
Svc Id      L2-Routes (RD-Prefix)          Next Hop      Origin
              Sdp Bind Id
-----
100          65535:100-1.1.1.6              1.1.1.6       BGP-L2
              17406:4294967295
-----
No. of L2 Route Entries: 1
=====
PERs6>#

PERs6># show service l2-route-table
=====
Services: L2 Route Information - Summary Service
=====
Svc Id      L2-Routes (RD-Prefix)          Next Hop      Origin
              Sdp Bind Id
-----
100          65535:100-1.1.1.134           1.1.1.134     BGP-L2
              17406:4294967295
-----
No. of L2 Route Entries: 1
=====
PERs6>#
```

When only one of the endpoints has an entry for the service in the l2-routing-table, it is most likely a problem with the RT values used for import and export. This would most likely happen when different import and export RT values are configured using a router policy or the route-target command.

Service specific commands continue to be available to display service specific information, including status.

```
PERs6# show service sdp-using
=====
SDP Using
=====
```

| SvcId | SdpId | Type | Far End | Opr S* | I.Label | E.Label |
|-------|------------------|-------|-----------|--------|---------|---------|
| 100 | 17406:4294967295 | BgpAd | 1.1.1.134 | Up | 131063 | 131067 |

Number of SDPs : 1

=====

* indicates that the corresponding row element may have been truncated.

BGP AD will advertise the VPLS-ID in the extended community attribute, VSI-ID in the NLRI and the local PE id in the BGP next hop. At the receiving PE, the VPLS-ID is compared against locally provisioned information to determine whether the two PEs share a common VPLS. If it is found that they do, the BGP information is used in the signaling phase (see [Configuring BGP VPLS on page 606](#)).

LDP Signaling

T-LDP is triggered once the VPN endpoints have been discovered using BGP. The T-LDP session between the PEs is established when one does not exist. The far-end IP address required for the T-LDP identification is gleaned from the BGP AD next hop information. The pw-template and pw-template-binding configuration statements are used to establish the automatic SDP or to map to the appropriate SDP. The FEC129 content is built using the following values:

- AGI from the locally configured VPLS-ID.
- The SAII from the locally configured VSI-ID.
- The TAI from the VSI-ID contained in the last 4 bytes of the received BGP NLRI.

Figure 95 below shows the different detailed phases of the LDP signaling path, post BGP AD completion. It also indicates how some fields can be auto generated when they are not specified in the configuration.




Figure 95: BGP AD Triggering LDP Functions

The first command will display the LDP peering relationships that have been established (Figure 96). The type of adjacency is displayed in the “Adj Type” column. In this case the type is “Both” meaning link and targeted sessions have been successfully established.

```
PERs6# show router ldp session
```

```
LDP Sessions
```

| Peer LDP Id | Adj Type | State | Msg Sent | Msg Recv | Up Time |
|--------------------|----------|-------------|----------|----------|-------------|
| 1.1.1.134:0 | Both | Established | 21482 | 21482 | 0d 15:38:44 |
| No. of Sessions: 1 | | | | | |

Figure 96: Show Router LDP Session Output

The second command shows the specific LDP service label information broken up per FEC element type, 128 or 129, basis (Figure 97). The information for FEC element 129 includes the AGI, SAI, and the TAIL.

```
PERs6# show router ldp bindings fec-type services
```

```
LDP LSR ID: 1.1.1.6
```

```
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
         S - Status Signaled Up, D - Status Signaled Down
         E - Epipe Service, V - VPLS Service, M - Mirror Service
         A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
         P - Ipipe Service, C - Cpipe Service
         TLV - (Type, Length: Value)
```

```
LDP Service FEC 128 Bindings
```

| Type | VCId | SvcId | SDPID | Peer | IngLbl | EgrLbl | LMTU | RMTU |
|---------------------------|------|-------|-------|------|--------|--------|------|------|
| No Matching Entries Found | | | | | | | | |

```
LDP Service FEC 129 Bindings
```

| AGI | | | SAI | | TAIL | | | |
|--------------------|-------|-------|-----------|---------|-----------|------|------|--|
| Type | SvcId | SDPID | Peer | IngLbl | EgrLbl | LMTU | RMTU | |
| 65535:100 | | | 1.1.1.6 | | 1.1.1.134 | | | |
| V-Eth | 100 | 17406 | 1.1.1.134 | 131063U | 131067S | 1464 | 1464 | |
| No. of FEC 129s: 1 | | | | | | | | |

Figure 97: Show Router LDP Bindings FEC-Type Services

Pseudowire Template

The pw-template is defined under the top level service command (**config>service# pw-template**) and specifies whether to use an automatically generated SDP or manually configured SDP. It also provides the set of parameters required for establishing the pseudowire (SDP binding) as displayed in [Figure 98](#).

```
PERs6>config>service# pw-template 1 create
- [no] pw-template <policy-id> [use-provisioned-sdp]

<policy-id>          : [1..2147483647]
<use-provisioned-s*> : keyword

[no] accounting-pol* - Configure accounting-policy to be used
[no] auto-learn-mac* - Enable/disable automatic update of MAC protect list
[no] block-on-peer-* - Enable/Disable block traffic on peer fault
[no] collect-stats   - Enable/disable statistics collection
[no] controlword     - Enable/Disable the use of ControlWord
[no] disable-aging   - Enable/disable aging of MAC addresses
[no] disable-learn*  - Enable/disable learning of new MAC addresses
[no] discard-unknow* - Enable/disable discarding of frames with unknown source
                        MAC address
                        egress + Spoke SDP binding egress configuration
[no] force-qinq-vc-* - Forces qinq-vc-type forwarding in the data-path
[no] force-vlan-vc-* - Forces vlan-vc-type forwarding in the data-path
[no] hash-label      - Enable/disable use of hash-label
                        igmp-snooping + Configure IGMP snooping parameters
                        ingress + Spoke SDP binding ingress configuration
[no] l2pt-terminati* - Configure L2PT termination on this spoke SDP
[no] limit-mac-move  - Configure mac move
[no] mac-pinning     - Enable/disable MAC address pinning on this spoke SDP
[no] max-nbr-mac-ad* - Configure the maximum number of MAC entries in the FDB
                        from this SDP
[no] restrict-prote* - Enable/disable protected src MAC restriction
[no] sdp-exclude     - Configure excluded SDP group
[no] sdp-include     - Configure included SDP group
[no] split-horizon-* + Configure a split horizon group
                        stp + Configure STP parameters
                        vc-type - Configure VC type
[no] vlan-vc-tag     - Configure VLAN VC tag
```

Figure 98: PW-Template CLI Tree

A **pw-template-binding** command configured within the VPLS service under the **bgp-ad** sub-command is a pointer to the pw-template that should be used. If a VPLS service does not specify an import-rt list, then that binding applies to all route targets accepted by that VPLS. The **pw-template-bind** command can select a different template on a per import-rt basis. It is also possible to specify specific pw-templates for some route targets with a VPLS service and use the single **pw-template-binding** command to address all unspecified but accepted imported targets.


```

PERs6>config>service>vpls>bgp-ad# pw-template-binding
- pw-template-binding <policy-id> [split-horizon-group <group-name>] [import-
rt
{ext-community, ...(upto 5 max)}}]
- no pw-template-binding <policy-id>

<policy-id>          : [1..2147483647]
<group-name>         : [32 chars max]
<ext-community>      : target:{<ip-addr:comm-val>|<as-number:ext-comm-val>}
                        ip-addr      - a.b.c.d
                        comm-val     - [0..65535]
                        as-number    - [1..65535]
                        ext-comm-val - [0..4294967295]

```

Figure 99: PW-Template-Binding CLI Syntax

It is important understand the significance of the split-horizon-group used by the pw-template. Traditionally, when a VPLS instance was manually created using mesh-sdp bindings, these were automatically placed in a common split-horizon-group to prevent forwarding between the pseudowire in the VPLS instances. This prevents loops that would have otherwise occurred in the Layer 2 service. When automatically discovering VPLS service using BGP AD the service provider has the option of associating the auto-discovered pseudowire with a split-horizon group to control the forwarding between pseudowires.

Configuring BGP VPLS

This section gives a configuration example required to bring up BGP VPLS in the VPLS PEs depicted in [Figure 100](#):




Figure 100: BGP VPLS Example

The red BGP VPLS is configured in the PE24, PE25 and PE26 using the commands shown in the following CLI examples.

```
*A:PE24>config>service>vpls# info
-----
      bgp
        route-distinguisher 65024:600
        route-target export target:65019:600 import target:65019:600
        pw-template-binding 1
      exit
    bgp-vpls
      max-ve-id 100
      ve-name 24
      ve-id 24
      exit
      no shutdown
    exit
    sap 1/1/20:600.* create
    exit
    no shutdown
  -----

*A:PE24>config>service>vpls#

*A:PE25>config>service>vpls# info
-----
      bgp
        route-distinguisher 65025:600
        route-target export target:65019:600 import target:65019:600
```

```

        pw-template-binding 1
    exit
    bgp-vpls
        max-ve-id 100
        ve-name 25
        ve-id 25
    exit
    no shutdown
exit
sap 1/1/19:600.* create
exit
no shutdown
-----
*A:PE25>config>service>vpls#

*A:PE26>config>service>vpls# info
-----
    bgp
        route-distinguisher 65026:600
        route-target export target:65019:600 import target:65019:600
        pw-template-binding 1
    exit
    bgp-vpls
        max-ve-id 100
        ve-name 26
        ve-id 26
    exit
    no shutdown
exit
sap 5/2/20:600.* create
exit
no shutdown
-----
*A:PE26>config>service>vpls#

```

Configuring a VPLS Management Interface

Use the following CLI syntax to create a VPLS management interface.

CLI Syntax: `config>service>vpls# interface ip-int-name
address ip-address[/mask] [netmask]
arp-timeout seconds
description description-string
mac ieee-address
no shutdown
static-arp ip-address ieee-address`

The following displays the configuration.

```
A:ALA-49>config>service>vpls>interface# info detail
-----
no description
mac 14:31:ff:00:00:00
address 123.231.10.10/24
no arp-timeout
no shutdown
-----
A:ALA-49>config>service>vpls>interface#
```

Configuring Policy-Based Forwarding for Deep Packet Inspection (DPI) in VPLS

The purpose of policy-based forwarding is to capture traffic from a customer and perform a deep packet inspection (DPI) and forward traffic, if allowed, by the DPI.

In the following example, the split horizon groups are used to prevent flooding of traffic. Traffic from customers enter at SAP 1/1/5:5. Due to the mac-filter 100 that is applied on ingress, all traffic with dot1p 07 marking will be forwarded to SAP 1/1/22:1, which is the DPI.

DPI performs packet inspection/modification and either drops the traffic or forwards the traffic back into the box through SAP 1/1/21:1. Traffic will then be sent to spoke-sdp 3:5.

SAP 1/1/23:5 is configured to see if the VPLS service is flooding all the traffic. If flooding is performed by the router then traffic would also be sent to SAP 1/1/23:5 (which it should not).

Figure 101 shows an example to configure policy-based forwarding for deep packet inspection on a VPLS service. For information about configuring filter policies, refer to the 7450 ESS OS Router Configuration Guide.

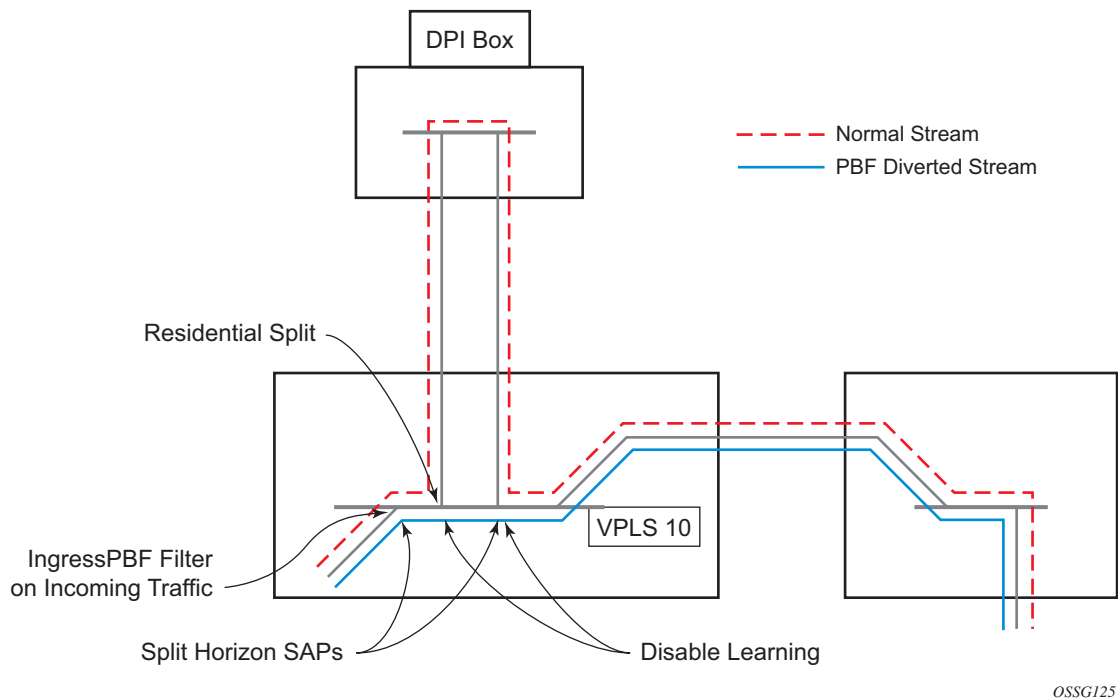


Figure 101: Policy-Based Forwarding For Deep Packet Inspection

The following example displays the service configuration:

```
*A:ALA-48>config>service# info
-----
...
    vpls 10 customer 1 create
        service-mtu 1400
        split-horizon-group "dpi" residential-group create
        exit
        split-horizon-group "split" create
        exit
        stp
            shutdown
        exit
        igmp-host-tracking
            expiry-time 65535
            no shutdown
        exit
        sap 1/1/21:1 split-horizon-group "split" create
            disable-learning
            static-mac 00:00:00:31:11:01 create
        exit
        sap 1/1/22:1 split-horizon-group "dpi" create
            disable-learning
            static-mac 00:00:00:31:12:01 create
        exit
        sap 1/1/23:5 create
            static-mac 00:00:00:31:13:05 create
        exit
        no shutdown
    exit
...
-----
*A:ALA-48>config>service#
```

The following example displays the MAC filter configuration:

```
*A:ALA-48>config>filter# info
-----
...
    mac-filter 100 create
        default-action forward
        entry 10 create
            match
                dot1p 7 7
            exit
            log 101
            action forward sap 1/1/22:1
        exit
    exit
...
-----
*A:ALA-48>config>filter#
```

The following example displays the service configuration with a MAC filter:

```
*A:ALA-48>config>service# info
-----
...
    vpls 10 customer 1 create
        service-mtu 1400
        split-horizon-group "dpi" residential-group create
        exit
        split-horizon-group "split" create
        exit
        stp
            shutdown
        exit
        igmp-host-tracking
            expiry-time 65535
            no shutdown
        exit
        sap 1/1/5:5 split-horizon-group "split" create
            ingress
                filter mac 100
            exit
            static-mac 00:00:00:31:15:05 create
        exit
        sap 1/1/21:1 split-horizon-group "split" create
            disable-learning
            static-mac 00:00:00:31:11:01 create
        exit
        sap 1/1/22:1 split-horizon-group "dpi" create
            disable-learning
            static-mac 00:00:00:31:12:01 create
        exit
        sap 1/1/23:5 create
            static-mac 00:00:00:31:13:05 create
        exit
        spoke-sdp 3:5 create
        exit
        no shutdown
    exit
....
-----
*A:ALA-48>config>service#
```

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 613](#)
 - [Modifying Management VPLS Parameters on page 614](#)
 - [Deleting a Management VPLS on page 614](#)
 - [Disabling a Management VPLS on page 615](#)
 - [Deleting a VPLS Service on page 616](#)
-

Modifying VPLS Service Parameters

You can change existing service parameters. The changes are applied immediately.

To display a list of services, use the **show service service-using vpls** command. Enter the parameter such as description, SAP, SDP, and/or service-MTU command syntax, and then enter the new information.

The following 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
remote-age 1000
stp
    shutdown
exit
sap 1/1/5:22 create
    description "VPLS SAP"
exit
spoke-sdp 2:22 create
exit
no shutdown
-----
*A:ALA-1>config>service>vpls#
```

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 620](#)
- [Oper Group Commands on page 627](#)
- [SAP Commands on page 628](#)
- [Mesh SDP Commands on page 640](#)
- [Spoke SDP Commands on page 643](#)
- [Provider Tunnel Commands on page 647](#)
- [Egress Multicast Group Commands on page 648](#)
- [Show Commands on page 650](#)
- [Clear Commands on page 654](#)
- [Debug Commands on page 656](#)

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
      — [no] allow-ip-int-binding
      — 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] 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]
— gsmp
    — [no] group name [create]
        — ancp
            — [no] dynamic-topology-discover
            — [no] line-configuration
            — [no] oam
            — description description-string
            — no description
            — hold-multiplier multiplier
            — no hold-multiplier
            — keepalive seconds
            — no keepalive
            — [no] neighbor ip-address
                — description description-string
                — no description
                — local-address ip-address
                — no local-address
                — priority-marking dscp dscp-name
                — priority-marking prec ip-prec-value
                — no priority-marking
                — [no] shutdown
            — [no] shutdown
        — [no] idle-filter
        — persistence-database
        — no persistence-database
        — [no] shutdown
— host-connectivity-verify source-ip ip-address [source-mac ieee-address] [interval
interval] [action { remove | alarm }]
— igmp-host-tracking
    — expiry-time expiry-time
    — no expiry-time
    — [no] shutdown
— igmp-snooping
    — mvr
        — description description-string
        — no description
        — group-policy policy-name
        — no group-policy
        — [no] shutdown
    — query-interval seconds
    — no query-interval
    — query-src-ip ip-address
    — no query-src-ip
    — report-src-ip ip-address
    — no report-src-ip
    — robust-count robust-count
    — no robust-count
    — [no] shutdown
— [no] interface ip-int-name
    — address ip-address[/mask] [netmask]
    — no address
    — arp-timeout seconds
    — no arp-timeout
    — description 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
— isis-policy
  — entry
    — [no] advertise-local
    — range isis [to isis]
    — 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
— mcast-info-policy policy-name
— no mcast-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]
        — 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] [capture-sap] [root-leaf-tag
leaf-tag-vid | leaf-ac]
— [no] sap eth-tunnel-tunnel-id [:eth-tunnel-sap-id]
— no sap sap-id
— accounting-policy acct-policy-id
— no accounting-policy
— [no] auto-learn-mac-protect
— anti-spoof {ip | mac | ip-mac}
— no anti-spoof
— app-profile app-profile-name
— no app-profile
— arp-host
— host-limit max-num-hosts
— no host-limit
— min-auth-interval min-auth-interval
— no min-auth-interval
— [no] shutdown
— arp-reply-agent [sub-ident]
— no arp-reply-agent
— authentication-policy name
— no authentication-policy
— bpdu-translation {auto | pvst | stp}
— no bpdu-translation
— calling-station-id {mac | remote-id | sap-id | sap-string}
— no calling-station-id
— [no] cflowd
— [no] collect-stats
— cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring[aggregate][car]]
— no cpu-protection
— default-msap-policy policy-name
— no default-msap-policy
— description description-string
— no description
— dhcp
— description description-string
— no description
— lease-populate [nbr-of-entries]
— no lease-populate
— [no] option
— action [dhcp-action]
— no action
— circuit-id [ascii-tuple | vlan-ascii-tuple]
— [no] remote-id [mac | string string]
— [no] vendor-specific-option
— [no] client-mac-address
— [no] sap-id
— [no] service-id

```



```

— string text
— no string
— [no] system-id
— proxy-server
— emulated-server ip-address
— no emulated-server
— lease-time [days days] [hrs hours] [min minutes] [sec
seconds] [radius-override]
— no lease-time
— [no] shutdown
— [no] shutdown
— [no] snoop
— dhcp-python-policy policy-name
— no dhcp-python-policy
— dhcp6-user-db local-user-db-name
— no dhcp6-user-db
— dhcp6
— description description-string
— no description
— [no] option
— interface-id
— interface-id ascii-tuple
— interface-id vlan-ascii-tuple
— no interface-id
— remote-id
— remote-id mac
— remote-id string [32 chars max]
— no remote-id
— [no] shutdown
— [no] snoop
— [no] disable-aging
— [no] disable-learning
— [no] discard-unknown
— dist-cpu-protection policy-name
— no dist-cpu-protection
— egress
— [no] agg-rate
— [no] limit-unused-bandwidth
— [no] queue-frame-based-accounting
— rate {max | rate}
— no rate
— 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] 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 subtract-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
— msap-defaults
    — [no] service service-id
    — policy msap-policy-name
    — no policy
— monitor-oper-group group-name
— no monitor-oper-group
— oper-group group-name
— no oper-group
— multi-service-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
— pppoe-policy pppoe-policy-name
— no pppoe-policy
— restrict-protected-src alarm-only
— restrict-protected-src [discard-frame]
— no restrict-protected-src
— restrict-unprotected-dst
— no restrict-unprotected-dst
— [no] shutdown
— static-host ip ip-address [mac ieee-address ] [create]
— static-host mac ieee-address [create]
— no static-host [ip ip-address ] mac ieee-address
— no static-host all [force]
— no static-host ip ip-address
    — ancp-string ancp-string
    — no ancp-string
    — app-profile app-profile-name
    — no app-profile
    — inter-dest-id intermediate-destination-id

```

```

— no inter-dest-id
— [no] shutdown
— sla-profile sla-profile-name
— no sla-profile
— sub-profile sub-profile-name
— no sub-profile
— subscriber sub-ident
— no subscriber
— [no] subscriber-sap-id
— [no] static-isid range entry-id isid [to isid] [create]
— [no] static-mac ieee-address
— stp
    — [no] auto-edge
    — [no] edge-port
    — link-type {pt-pt | shared}
    — no link-type {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
— [no] sub-sla-mgmt
    — def-sla-profile default-sla-profile-name
    — no def-sla-profile
    — def-sub-profile default-subscriber-profile-name
    — no def-sub-profile
    — [no] mac-da-hashing
    — multi-sub-sap [subscriber-limit]
    — [no] shutdown
    — single-sub-parameters
        — non-sub-traffic sub-profile sub-profile-name sla-
          profile sla-profile-name [subscriber sub-ident-string]
        — no non-sub-traffic
        — [no] profiled-traffic-only
    — sub-ident-policy sub-ident-policy-name
    — no sub-ident-policy
— tod-suite tod-suite-name
— no tod-suite
— trigger-packet [dhcp] [pppoe] [arp] [dhcp6] [ppp]
— no trigger-packet

```


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
— cpu-protection policy-id [mac-monitoring]
— no cpu-protection
— eth-cfm
    — [no] mip primary-vlan-enable [vlan vlan-id]
    — [no] squelch-ingress-levels [md-level [md-level...]]
— [no] disable-aging
— [no] disable-learning
— [no] discard-unknown
— 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 [shared-queuing | multipoint-shared]
    — 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
        — dhcp
          — description description-string
          — no description
          — snoop [l2-header]
          — no snoop
        — egress
          — filter ip ip-filter-id
          — filter ipv6 ipv6-filter-id
          — filter mac mac-filter-id
          — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
          — qos network-policy-id port-redirect-group queue-group-name [instance instance-id]
          — no qos
          — mfib-allowed-mda-destinations
            — [no] mda mda-id
          — vc-label egress-vc-label
          — no vc-label [egress-vc-label]
        — 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] squelch-ingress-levels [md-level [md-level...]]
— [no] vmep-filter
— fault-propagation-bmac [mac-name | ieee-address] [create]
— no fault-propagation-bmac [mac-name | ieee-address]
— [no] force-qinq-vc-forwarding
— [no] force-vlan-vc-forwarding
— [no] hash-label
— igmp-snooping
    — [no] disable-router-alert-check
    — [no] fast-leave
    — import policy-name
    — no import
    — last-member-query-interval interval
    — no last-member-query-interval
    — max-num-groups max-num-groups
    — no max-num-groups
    — max-num-grp-sources [1..32000]
    — no max-num-grp-sources
    — mcac
        — policy policy-name
        — no policy
        — unconstrained-bw bandwidth mandatory-bw mandatory-bw
        — no unconstrained-bw
    — [no] mrrouter-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]
— qos network-policy-id fp-redirect-group queue-group-name
   instance-id
— no qos
— mfib-allowed-mda-destinations
   — [no] mda mda-id
— vc-label ingress-vc-label
— no vc-label [ingress-vc-label]
— [no] mac-pinning
— mld-snooping
   — [no] disable-router-alert-check
   — [no] fast-leave
   — import policy-name
   — no import
   — last-member-query-interval interval
   — no last-member-query-interval
   — max-num-groups max-num-groups
   — no max-num-groups
   — 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
        — [no] control-channel-status
          — [no] acknowledgment
          — refresh-timer value
          — no refresh-timer
          — request-timer timer1 retry-timer timer2 [timeout-multiplier multiplier]
          — no request-timer
        — [no] control-word
        — description description-string
        — no description
        — dhcp
          — description description-string
          — no description
          — [no] snoop
        — [no] disable-aging
        — [no] disable-learning
        — [no] discard-unknown-source
        — 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
— no ccm-padding-size ccm-padding
— [no] csf-enable
    — multiplier multiplier-value
    — no multiplier
— [no] description
— [no] eth-test-enable
    — test-pattern {all-zeros | all-ones} [crc-enable]
    — no test-pattern
— fault-propagation-enable {use-if-tlv | suspend-ccm}
— no fault-propagation-enable
— 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] squelch-ingress-levels [md-level [md-level...]]
— vmep-filter
— egress
    — filter ip ip-filter-id
    — filter ipv6 ipv6-filter-id
    — filter mac mac-filter-id
    — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
    — qos network-policy-id port-redirect-group queue-group-name [instance-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]
— qos network-policy-id fp-redirect-group queue-group-name
— instance instance-id
— no qos
— mfib-allowed-mda-destinations
— [no] mda mda-id
— vc-label egress-vc-label
— no vc-label [egress-vc-label]
— [no] l2pt-termination
— limit-mac-move [blockable | non-blockable]
— no limit-mac-move
— [no] mac-pinning
— max-nbr-mac-addr table-size
— no max-nbr-mac-addr
— mld-snooping
— [no] disable-router-alert-check
— [no] fast-leave
— import policy-name
— no import
— last-member-query-interval interval
— no last-member-query-interval
— max-num-groups max-num-groups
— no max-num-groups
— 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-path-id
— agi agi
— no agi
— saii-type2 global-id:node-id:ac-id
— no saii-type2
— taii-type2 global-id:node-id:ac-id
— no taii-type2
— [no] pw-status-signaling
— propagate-mac-flush [precedence-value / primary]
— no propagate-mac-flush
— [no] shutdown
— [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

```


Show Commands

```

show
  — service
    — active-subscribers summary
    — active-subscribers [subscriber sub-ident-string] [sap-id sap-id] [sla-profile sla-profile-name] [detail | mirror]
    — active-subscribers hierarchy [subscriber sub-ident-string]
    — egress-label egress-label1 [egress-label2]
    — fdb-info
    — fdb-mac ieee-address [expiry]
    — id service-id
      — all
      — arp-host [wholesaler service-id] [sap sap-id | interface interface-name | ip-address ip-address[/mask] | mac ieee-address | {[port port-id] [no-inter-dest-id | inter-dest-id inter-dest-id]}] [detail]
      — arp-host statistics [sap sap-id | interface interface-name]
      — arp-host summary [interface interface-name]
      — authentication
        — statistics [policy name] [sap sap-id]
      — base [msap]
      — bgp-evpn
      — dhcp
        — lease-state [[sap sap-id] | [sdp sdp-id:vc-id] | [interface interface-name] | [ip-address ip-address]] [detail] [[mac ieee-address] [[wholesaler service-id] [detail]
        — lease-state [sap sap-id]
        — statistics [sdp sdp-id:vc-id]
        — statistics [interface interface-name]
        — summary
      — endpoint [endpoint-name]
      — etree
      — fdb [sap sap-id] [expiry]] | [sdp sdp-id [expiry]] | [mac ieee-address [expiry]] | endpoint endpoint | [detail] [expiry] [pbb]
      — gsmp
        — neighbors group [name] [ip-address]
        — sessions [group name] neighbor ip-address] [port port-number] [association] [statistics]
      — host [sap sap-id] [detail]
      — host summary
      — i-vpls
      — igmp-snooping
        — all
        — base
        — mroouters [detail]
        — port-db sap sap-id [detail]
        — port-db sap sap-id group grp-address
        — port-db sdp sdp-id:vc-id [detail]
        — port-db sdp sdp-id:vc-id group grp-address
        — port-db vxlan vtep ip-address vni [0..4294967295]
        — proxy-db [detail]
        — proxy-db [group grp-ip-address]
        — querier
        — static [sap sap-id | sdp sdp-id:vc-id]

```

- **isid-policy**
- **labels**
- **l2pt** disabled
- **l2pt** [detail]
- **mac-move**
- **mac-protect**
- **mfib** [brief | statistics] [ip | mac] brief
- **mfib** [group grp-address | *] [statistics]
- **mld-snooping**
 - **all**
 - **base**
 - **mrouters** [detail]
 - **mvr**
 - **port-db** sap sap-id
 - **port-db** sap sap-id detail
 - **port-db** sap sap-id group grp-ipv6-address
 - **port-db** sdp sdp-id:vc-id
 - **port-db** sdp \sdp-id:vc-id detail
 - **port-db** sdp \sdp-id:vc-id group grp-ipv6-address
 - **proxy-db** [detail]
 - **proxy-db** group grp-ipv6-address
 - **querier**
 - **static** [sap sap-id | sdp sdp-id:vc-id]
 - **static** [sap sap-id | sdp sdp-id:vc-id]
- **mmrp** mac [ieee-address]
- **mrp-policy** [mrp-policy]
- **mrp-policy** mrp-policy [association]
- **mrp-policy** mrp-policy [entry entry-id]
- **mstp-configuration**
- **provider-tunnel**
- **proxy-arp** [ip-address ip-address] [detail]
- **proxy-nd** [ip-address ip-address] [detail]
- **retailers**
- **sap** [sap-id [detail]]
- **sdp** [sdp-id | far-end ip-addr] [detail]
- **sdp** sdp-id:vc-id {mrp | mmrp}
- **site** [detail]
- **site** name
- **split-horizon-group** [group-name]
- **stp** mst-instance mst-inst-number
- **stp** [detail]
- **subscriber-hosts** [sap sap-id] [ip ip-address[/mask]] [mac ieee-address] [sub-profile sub-profile-name] [sla-profile sla-profile-name] [detail]
- **wholesalers**
- **vxlan**
- **ingress-label** start-label [end-label]
- **isid-using** [ISID]
- **sap-using** [msap] [dyn-script] [description]
- **sap-using** [sap sap-id] [vlan-translation | anti-spoof]
- **sap-using** app-profile app-profile-name
- **sap-using** authentication-policy policy-name [msap]
- **sap-using** encap-type encap-type
- **sap-using** eth-cfm collect-lmm-stats [sap sap-id]
- **sap-using** eth-ring [ring-id eth-ring-id]

```

— sap-using eth-tunnel [tunnel-id eth-tunnel-id]
— sap-using interface [ip-address | ip-int-name]
— sap-using [ingress | egress] 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
— subscriber-using [service-id service-id] [sap-id sap-id] [interface ip-int-name] [ip ip-address[/mask]] [mac ieee-address] [sub-profile sub-profile-name] [sla-profile sla-profile-name]
— vxlan
— egress-replication
— vpls vpls-service-id mda slot/mda
— vpls vpls-service-id mda slot/mda [igmp-record grp-address {source source-ip-address | starg}] | [mRouter]

show
— igmp
— group [grp-ip-address]
— ssm-translate
— interface [ip-int-name | ip-address] [group grp-address] [detail]
— static [ip-int-name | ip-addr]
— statistics [ip-int-name | ip-address]
— status

```


Show Multi-Chassis Endpoint Commands

```
show
  — service
    — id service-id
      — endpoint [endpoint-name]
  — redundancy
    — multi-chassis
      — mc-endpoint statistics
      — mc-endpoint peer [ip-address] statistics
      — mc-endpoint endpoint [mcep-id] statistics
      — mc-endpoint peer [ip-address]
```

Clear Commands

```
clear
  — service
    — id service-id
      — arp-host { mac ieee-address | sap sap-id | ip-address ip-address[/mask] }
      — arp-host [port port-id] [inter-dest-id intermediate-destination-id | no-inter-dest-id]
      — arp-host statistics [sap sap-id | interface interface-name]
      — authentication
        — statistics
      — dhcp
        — lease-state [no-dhcp-release]
        — lease-state ip-address [ip-address[/mask]] [no-dhcp-release]
        — lease-state mac ieee-address [no-dhcp-release]
        — lease-state sap sap-id [no-dhcp-release]
        — lease-state sdp sdp-id:vc-id [no-dhcp-release]
        — statistics [sap sap-id | sdp sdp-id:vc-id | interface ip-int-name | ip-address]
      — fdb { all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-sdp sdp-id:vc-id }
      — igmp-snooping
        — port-db sap sap-id [group grp-address [source ip-address]]
        — port-db sdp sdp-id:vc-id [group grp-address [source ip-address]]
        — querier
      — mfib
        — statistics { all | ip | mac }
        — statistics group grp-address
      — mld-snooping
        — port-db sap sap-id [group grp-ipv6-address]
        — port-db sap sap-id group grp-ipv6-address source src-ipv6-address
        — port-db sdp sdp-id:vc-id [group grp-ipv6-address]
        — port-db sdp sdp-id:vc-id group grp-ipv6-address source src-ipv6-address
        — querier
        — statistics all
        — statistics sap sap-id
        — statistics sdp sdp-id:vc-id
      — mesh-sdp sdp-id[:vc-id] ingress-vc-label
      — msap msap-id
      — spoke-sdp sdp-id:vc-id ingress-vc-label
      — proxy-arp { all | ip-address } [{dynamic|dup}]
      — proxy-nd { all | ipv6-address } [{dynamic|dup}]
      — stp
        — detected-protocols [all | sap sap-id / spoke-sdp [sdp-id[:vc-id]]]
    — statistics
      — id service-id
        — capture-sap sap-id [trigger]
        — cem
        — counters
        — l2pt
        — mesh-sdp sdp-id[:vc-id] { all | counters | stp | mrp }
        — mrp
        — pip
        — spoke-sdp sdp-id[:vc-id] { all | counters | stp | l2pt | mrp }
```

```
— spoke-sdp
— stp
— sap sap-id {all | cem | counters | l2pt | stp | mrp}
— sdp sap-id {keep-alive}

clear
— router
— dhcp
— statistics [interface ip-int-name | ip-address]
```

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 | sdpbinding-oper-status-change }
      — [no] host-connectivity-verify
        — [no] ip ip-address
        — [no] mac ieee-address
        — [no] sap sap-id
      — [no] proxy-arp [mac ieee-address] [ip [<ipaddr>] [all]]
      — [no] proxy-nd [mac ieee-address] [ip [<ipaddr>] [all]]
      — [no] sap sap-id
      — stp
        — all-events
        — [no] bpdu
        — [no] core-connectivity
        — [no] exception
        — [no] fsm-state-changes
        — [no] fsm-timers
        — [no] port-role
        — [no] port-state
        — [no] sap sap-id
        — [no] sdp sdp-id:vc-id
debug
  — igmp

```

- **router**
 - [no] **interface** *[ip-int-name / ip-address]*
 - [no] **mcs** *[ip-int-name]*
 - [no] **mcs**
 - [no] **packet** *[query/v1-report/v2-report/v3-report/v2-leave] [ip-int-name / ip-address]*

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 7450 ESS 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 | [no] shutdown |
| 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>sap>dhcp>proxy 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.

description

| | |
|--------------------|--|
| Syntax | description <i>description-string</i> no description |
| Context | config>service>vpls config>service>vpls>sap>dhcp6 config>service>vpls>gsmp>group config>service>vpls>gsmp>group>neighbor config>service>vpls>igmp-snooping>mvr config>service>vpls>interface config>service>vpls>split-horizon-group config>service>vpls>sap config>service>vpls>mesh-sdp config>service>vpls>spoke-sdp config>service>vpls>sap>dhcp config>service>vpls>mld-snooping>mvr |
| Description | This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The no form of this command removes the string from the configuration. |
| 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. |

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>To create a management VPLS, the m-vpls keyword must be specified. See section Hierarchical VPLS Redundancy for an introduction to the concept of management VPLS.</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> <table><tr><td>Values</td><td><i>service-id:</i></td><td>1 — 2147483648</td></tr><tr><td></td><td><i>svc-name:</i></td><td>64 characters maximum</td></tr></table> <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> <table><tr><td>Values</td><td>1 — 2147483647</td></tr></table> <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> | Values | <i>service-id:</i> | 1 — 2147483648 | | <i>svc-name:</i> | 64 characters maximum | Values | 1 — 2147483647 |
| Values | <i>service-id:</i> | 1 — 2147483648 | | | | | | | |
| | <i>svc-name:</i> | 64 characters maximum | | | | | | | |
| Values | 1 — 2147483647 | | | | | | | | |

Values 1 — 2147483647

Default null (0)

m-vpls — Specifies a management VPLS.

e-tree — Specifies a VPLS service as an E-Tree VPLS. E-Tree VPLS services have root and leaf-ac SAPs/SDP bindings and root-leaf-tag SAPs/SDP bindings for E-Tree interconnection. The access (AC) root SAP behaves as a normal VPLS SAP. The AC leaf SAP is restricted to communication only with root-connected services. AC leaf and root SAPs are externally normal SAPs. The AC root SDP bind behaves as a normal VPLS SDP bind. The AC leaf SDP bind is restricted to communication only with root-connected services. AC leaf and root SDP bindings are externally normal SDPs bindings.

In the E-Tree VPLS, the root-ac SAP/SDP bindings can communicate with other root-ac and leaf-ac SAP/SDP bind services locally and remotely. Root originated traffic is marked internally with a root indication and root tagged externally on tag SAP/SDP binds. The leaf-ac SAP/SDP bindings can communicate with other root SAP/SDP bindings locally and remotely. Leaf originated traffic is marked internally with a leaf indication and tagged externally on leaf tag SAP/SDP bindings.

There may be any number of AC SAPs of root or leaf up to typical SAP limits. Network Side tag SAPs for root-leaf use additional resources. These tag SAPs used two tags one for Root and one for Leaf. Network side tag SDPs use a hard coded tag of 1 for root and 2 for leaf. AC SDP bindings are designated as root or leaf SDP bindings but carry no tags marking traffic on the egress frames.

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. |
| Parameters | <i>vpls-id</i> — This value represents the VPLS ID value associated with the B-VPLS. <i>isid</i> — 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.

calling-station-id

| | |
|--------------------|--|
| Syntax | calling-station-id { mac remote-id sap-id sap-string } no calling-station-id |
| Context | config>service>vpls>sap |
| Description | This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages. |
| Default | no calling-station-id |
| Parameters | mac — Specifies that the mac-address will be sent. remote-id — Specifies that the remote-id will be sent. sap-id — Specifies that the sap-id will be sent. sap-string — Specifies that the value is the inserted value set at the SAP level. If no calling-station-id value is set at the SAP level, the calling-station-id attribute will not be sent. |

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. |

l2pt-termination

| | |
|--------------------|--|
| Syntax | l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp] no l2pt-termination |
| Context | config>service>vpls>spoke-sdp config>service>vpls>sap |
| Description | <p>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.</p> <p>This feature can be enabled only if STP is disabled in the context of the given VPLS service.</p> |
| Default | no l2pt-termination |
| Parameters | <p>cdp — Specifies the Cisco discovery protocol.</p> <p>dtp — Specifies the dynamic trunking protocol.</p> <p>pagp — Specifies the port aggregation protocol.</p> <p>stp — Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default).</p> <p>udld — Specifies unidirectional link detection.</p> <p>vtp — Specifies the virtual trunk protocol.</p> |

def-mesh-vc-id

| | |
|--------------------|---|
| Syntax | [no] def-mesh-vc-id vc-id |
| Context | config>service>vpls |
| Description | <p>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.</p> <p>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.</p> <p>The no form of this command disables the VC-ID.</p> |

VPLS Service Commands

| | |
|-------------------|--|
| 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. |

dhcp-python-policy

| | |
|--------------------|--|
| Syntax | dhcp-python-policy <i>policy-name</i> no dhcp-python-policy |
| Context | config>service>vpls>sap |
| Description | This command specifies the name of the Python policy. The Python policy is created in the config>python>python-policy <i>name</i> context. The no form of the command reverts to the default. |
| Default | none |
| Parameters | <i>policy-name</i> — Specifies a Python policy name up to 32 characters in length. |

dhcp6-user-db

| | |
|--------------------|--|
| Syntax | dhcp6-user-db <i>local-user-db-name</i> no dhcp6-user-db |
| Context | config>service>vpls>sap |
| Description | This command assigns a local user database for DHCP6 clients (capture SAP only)/ The no form of the command removes the name from the configuration. |
| Parameters | <i>local-user-db-name</i> — Specifies a local user database name up to 32 characters in length. |

dhcp6

| | |
|--------------------|--|
| Syntax | dhcp6 |
| Context | config>service>vpls>sap |
| Description | This command configures DHCP6 parameters for this SAP. |

interface-id

| | |
|--------------------|---|
| Syntax | interface-id interface-id ascii-tuple interface-id vlan-ascii-tuple no interface-id |
| Context | config>service>ies>if>ipv6>dhcp6>option |
| Description | This command configure the interface-id suboption of the DHCP6 Relay packet The no form of the command disables the sending of interface ID options in the DHCPv6 relay packet |
| Parameters | ascii-tuple — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q-encapsulated ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet. mac — This keyword specifies the MAC address of the remote end is encoded in the sub-option. |

remote-id

| | |
|--------------------|--|
| Syntax | remote-id remote-id mac remote-id string [32 chars max] no remote-id |
| Context | config>service>ies>if>ipv6>dhcp6>option |
| Description | This command enables the sending of remote ID option in the DHCPv6 relay packet. The client DHCP Unique Identifier (DUID) is used as the remote ID. The no form of the command disables the sending of remote ID option in the DHCPv6 relay packet. |

interface-id

| | |
|--------------------|---|
| Syntax | interface-id interface-id ascii-tuple interface-id vlan-ascii-tuple no interface-id |
| Context | config>service>vpls>sap>dhcp6>option |
| Description | This command configures the interface-id suboption of the DHCP6 relay packet. The no form of the command reverts to the default. |
| Default | none |
| Parameters | ascii-tuple — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used. vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet. |

disable-aging

| | |
|--------------------|---|
| Syntax | [no] disable-aging |
| Context | config>service>vpls config>service>vpls>spoke-sdp config>service>vpls>sap config>template>vpls-template |
| Description | This command disables MAC address aging across a VPLS service or on a VPLS service SAP or spoke SDP. Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB). The disable-aging command turns off aging for local and remote learned MAC addresses. |

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specific SAPs and/or spoke SDPs by entering the **disable-aging** command at the appropriate level.

When the **disable-aging** command is entered at the VPLS level, the **disable-aging** state of individual SAPs or SDPs will be ignored.

The **no** form of this command enables aging on the VPLS service.

Default no disable-aging

disable-learning

Syntax [no] disable-learning

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>template>vpls-template

Description This command disables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance, SAP instance or spoke SDP instance.

When **disable-learning** is enabled, new source MAC addresses will not be entered in the VPLS service forwarding database. This is true for both local and remote MAC addresses.

When **disable-learning** is disabled, new source MAC addresses will be learned and entered into the VPLS forwarding database.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses.

Default no disable-learning (Normal MAC learning is enabled)

discard-unknown

Syntax [no] discard-unknown

Context config>service>vpls
config>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.

dist-cpu-protection

| | |
|--------------------|--|
| Syntax | dist-cpu-protection <i>policy-name</i> no dist-cpu-protection |
| Context | config>service>vpls>sap |
| Description | This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid created DCP policy can be assigned to a SAP or a network interface. Note that this rule does not apply to templates such as msap-policy. |
| Default | no dist-cpu-protection |

endpoint

| | |
|--------------------|--|
| Syntax | endpoint <i>endpoint-name</i> [create] no endpoint |
| Context | config>service>vpls |
| Description | This command configures a service endpoint. |
| Parameters | <i>endpoint-name</i> — Specifies an endpoint name up to 32 characters in length. create — This keyword is mandatory while creating a service endpoint. |

description

| | |
|-------------------|--|
| Syntax | description <i>description-string</i> 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 | <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. |

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 429](#).

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

Syntax restrict-protected-src alarm-only
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 **alarm-only** — Specifies that the SAP will be left up and only a notification, sapReceivedProtSrcMac, will be generated.

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. <table> <tr> <td>Values</td><td>1 — 511999 Chassis-mode A or B limit: 131071 Chassis-mode D limit: 511999</td></tr> </table> | Values | 1 — 511999 Chassis-mode A or B limit: 131071 Chassis-mode D limit: 511999 |
| Values | 1 — 511999 Chassis-mode A or B limit: 131071 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 | <p>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.</p> <p>When the arp-populate and lease-populate commands are enabled on an interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured arp-timeout value has no effect.</p> <p>The default value for arp-timeout is 14400 seconds (4 hours).</p> <p>The no form of this command restores arp-timeout to the default value.</p> |

| | |
|-------------------|---|
| 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 |

mac

| | |
|--------------------|--|
| Syntax | mac <i>ieee-address</i> no mac |
| Context | config>service>vpls>interface |
| Description | <p>This command assigns a specific MAC address to a VPLS IP interface.</p> <p>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.</p> <p>The no form of the command returns the MAC address of the IP interface to the default value.</p> |
| Default | The system chassis MAC address. |
| Parameters | <i>ieee-address</i> — 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 <i>ieee-mac-addr unnumbered</i> no static-arp <i>unnumbered</i> |
| Context | config>service>vpls>interface |
| Description | <p>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.</p> <p>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.</p> <p>The no form of the command removes a static ARP entry.</p> |
| Default | None |
| Parameters | <p><i>ip-address</i> — Specifies the IP address for the static ARP in dotted decimal notation.</p> <p><i>ieee-mac-address</i> — 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.</p> |

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 | <p>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.</p> <p>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.</p> <p>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.</p> <p>Static MACs configured in a bgp-evpn service are advertised as protected (EVPN will signal the mac as protected).</p> |

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 | <p>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.</p> <p>This command also assigns a conditional static MAC address entry to an EVPN VPLS SAP/spoke-SDP.</p> <p>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.</p> |
| Default | none |
| Parameters | <p>ieee-address — Specifies the static MAC address to an SPBM/sdp-binding interface.</p> <p>Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx). It cannot be all zeros.</p> <p>create — This keyword is mandatory while creating a static MAC.</p> <p>monitor fwd-status — Specifies that this static mac is based on the forwarding status of the SAP or spoke SDP for multi-homed operation.</p> |

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, or Frame Relay encapsulation. It is also supported on an Ethernet port. It is not supported on a TDM port or channel. |
| 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 | This command configures isid-policies for individual ISIDs or ISID ranges in a B-VPLS using SPBM. The ISIDs may belong to I-VPLS services or may be static-isids defined on this node. Multiple entry statements are allowed under a isid-policy . ISIDs that are declared as static do not require and isid-policy unless the ISIDs are not to be advertised. isid-policy allows finer control of ISID multicast but is not typically required for SPBM operation. Use of ISID policies can cause additional flooding of multicast traffic. |
| Default | no default |

entry

| | |
|--------------------|---|
| | entry <i>id</i> create no entry |
| Context | config>service>vpls>isid-policy |
| Description | This command creates or edits an isid-policy entry. Multiple entries can be created using unique entry-id numbers within the isid-policy. Default: No entry 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. The following rules govern the usage of multiple entry statements: |

- overlapping values are allowed:
 - isid from 301 to 310
 - isid from 305 to 315
 - isid 316
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “isid from 301 to 316” statement.
- there is no consistency check with the content of ISID statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation to execute the associated action for that entry.

no isid - removes all the previous statements under one entry.

no isid value | from value to higher-value - removes a specific ISID value or range. Must match a previously used positive statement: for example, if the command “isid 16 to 100” was used using “no isid 16 to 50”, it will not work but “no isid 16 to 100 will be successful.

Values 1-65535

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

to isid — Identifies upper value in a range of ISIDs to be used as matching criteria.

use-def-mcast

| | |
|--------------------|--|
| Syntax | [no] use-def-mcast |
| Context | config>service>vpls>isid-policy>entry |
| Description | The use-def-mcast option prevents local installation of the ISIDs in the range in the MFIB and uses the default multicast tree instead for a B-VPLS. In a node that does not have I-VPLS or static-isids, this command prevents the building of an MFIB entry for this ISID when received in a SPBM TLV and allows the broadcast of ISID based traffic on the default multicast tree. If an isid-policy exists, the core nodes can have this policy to prevent connectivity problems when some nodes are advertizing 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 | <p>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.</p> <p>The following algorithm describes the hash-key used for hashing when the new option is enabled:</p> <ul style="list-style-type: none"> • 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 <i>aging-timer</i> no local-age |
|---------------|--|

| | |
|--------------------|---|
| Context | config>service>vpls config>template>vpls-template |
| Description | <p>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.</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 FDB. The local-age timer specifies the aging time for local learned MAC addresses.</p> <p>The no form of this command returns the local aging timer to the default value.</p> |
| 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</p> |

relearn per second has been reached. This can already occur in the first second if the real relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

mac-protect

| | |
|--------------------|---|
| Syntax | mac-protect |
| Context | config>service>vpls |
| Description | This command indicates whether or not this MAC is protected on the MAC protect list. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP, spoke SDP or mesh-SDP that has restricted learning enabled. The MAC protect list is used in conjunction with restrict-protected-src , restrict-unprotected-dst and auto-learn-mac-protect . |
| Default | disabled |

mac

| | |
|--------------------|---|
| Syntax | [no] mac <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. |

mac-notification

| | |
|--------------------|--|
| Syntax | mac-notification |
| Context | config>service>vpls bvpls |
| Description | <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"> 1. When enabled in the BVPLS using no shutdown, a MAC notification will be sent for every active MC-LAG link. The following 3 cases assume no shutdown in the BVPLS. 2. Whenever a related MC-LAG link becomes active (related MC-LAG link = has at least 1 SAP associated with the BVPLS) if the MC-LAG peering is initialized and the PE peers are synchronized. 3. 1st SAP on an active MC-LAG is associated (via IVPLS/Epipe) with the BVPLS 4. The link between IVPLS/Epipe and BVPLS is configured and there are I-SAPs configured on an active MC-LAG link. <p>The MAC notification is not sent for the following events:</p> <ol style="list-style-type: none"> 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 | <p><i>value</i> — Specifies the frequency of subsequent MAC notification messages.</p> <p>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.</p> |

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 *value* — Specifies the time interval between re-notification in seconds.

Values 240—840 seconds

count

Syntax [no] count *value*

Context config>service>vpls>mac-notification

Description This command configures how often MAC notification messages are sent.

Parameters *value* — Specifies, in seconds, how often MAC notification messages are sent.

Values 1—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 | <i>number-retries</i> 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] [capture-sap] 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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 1319](#) for command syntax.

spoke-sdp

Syntax **[no] spoke-sdp** *spoke-id*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>secondary-ports

Description This command declares a given spoke SDP as a primary (or secondary) VPLS port.

Parameters *spoke-id* — Specifies the SDP ID to configure as the primary VPLS port.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

cumulative-factor

Syntax **[no] cumulative-factor** *factor*

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 *factor* — 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 | <p>This indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.</p> <p>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.</p> <p>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 effective retry timeout is doubled in order to avoid thrashing.</p> <p>The no form of the command reverts to the default value.</p> |
| 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 | <p>This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.</p> <p>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 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.</p> <p>The no form of this command removes the configured maximum MFIB table size.</p> |
| Default | none |
| Parameters | <i>size</i> — 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 <i>seconds</i> no remote-age |
| Context | config>service>vpls config>template>vpls-template |

| | |
|--------------------|--|
| Description | <p>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.</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 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.</p> <p>The no form of this command returns the remote aging timer to the default value.</p> |
| Default | remote age 900 — Remote MACs aged after 900 seconds |
| Parameters | <i>seconds</i> — 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 | 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. The no form of this command disables the generation of LDP MAC flush in I-VPLS on failure of SAP/endpoint/spoke-SDP in the related B-VPLS. |
| Default | no send-flush-on-bvpls-failure |

propagate-mac-flush-from-bvpls

| | |
|--------------------|--|
| Syntax | [no] propagate-mac-flush-from-bvpls |
| Context | config>service>vpls ivpls |
| Description | 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. 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 <i>octets</i> 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 <i>service-name</i> no service-name |
| Context | config>service>vpls |
| Description | <p>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 7750 SR, 7450 ESS and 7710 SRplatforms.</p> <p>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.</p> |
| Parameters | <i>service-name</i> — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9). |

site

| | |
|--------------------|---|
| Syntax | site <i>name</i> [create] no site <i>name</i> |
| Context | config>service>vpls |
| Description | <p>This command configures a VPLS site.</p> <p>The no form of the command removes the name from the configuration.</p> |
| Parameters | <p><i>name</i> — Specifies a site name up to 32 characters in length.</p> <p>create — This keyword is mandatory while creating a VPLS service.</p> |

boot-timer

| | |
|--------------------|--|
| Syntax | boot-timer <i>seconds</i> no boot-timer |
| Context | config>service>vpls>site |
| Description | <p>This command configures for how long the service manager 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.</p> <p>The no form of the command reverts the default.</p> |
| Default | 10 |
| Parameters | <p><i>seconds</i> — Specifies the site boot-timer in seconds.</p> <p>Values 0 — 100</p> |

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 1319](#) 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 **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

VPLS Service Commands

| | |
|--------------------|--|
| Syntax | site-id <i>value</i> no site-id |
| Context | config>service>vpls>site |
| Description | This command configures the identifier for the site in this service. |
| Parameters | <i>value</i> — Specifies the site identifier. |
| Values | 1 — 65535 |

split-horizon-group

| | |
|--------------------|---|
| Syntax | split-horizon-group <i>group-name</i> 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 | <i>group-name</i> — Specifies a split-horizon group name. |

spoke-sdp

| | |
|--------------------|--|
| Syntax | spoke-sdp <i>sdp-id:vc-id</i> 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. |

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 |

lsp-lifetime

| | |
|--------------------|--|
| Syntax | lsp-lifetime <i>seconds</i> no lsp-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 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.</p> <p>The LSP refresh timer is derived from this formula: $\text{lsp-lifetime}/2$</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> |

lsp-refresh-interval

| | |
|--------------------|--|
| Syntax | lsp-refresh-interval <i>seconds</i> no lsp-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 lsp-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 lsp-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> |

lsp-wait

| | |
|--------------------|--|
| Syntax | lsp-wait <i>lsp-wait</i> [<i>lsp-initial-wait</i> [<i>lsp-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 |

LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

| | |
|-------------------|--|
| Parameters | <i>lsp-max-wait</i> — Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated. |
| | Values 1 — 120 |
| | Default 5 |
| | <i>lsp-initial-wait</i> — Specifies the initial LSP generation delay in seconds. |
| | Values 0 — 100 |
| | Default 0 |
| | <i>lsp-second-wait</i> — 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 | <p>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:</p> <ol style="list-style-type: none"> 1. The timeout timer expires. 2. A manual override of the current overload state is entered with the config>router>isis>no overload command. <p>The no overload command does not affect the overload-on-boot function.</p> <p>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:</p> <pre>L1 LSDB Overload : Manual on boot (Indefinitely in overload) L2 LSDB Overload : Manual on boot (Indefinitely in overload)</pre> <p>This state can be cleared with the config>router>isis>no overload command.</p> <p>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:</p> <pre>L1 LSDB Overload : Manual on boot (Overload Time Left : 17) L2 LSDB Overload : Manual on boot (Overload Time Left : 17)</pre> <p>The overload state can be cleared before the timeout expires with the config>router>isis>no overload command.</p> <p>The no form of the command removes the overload-on-boot functionality from the configuration.</p> |

| | |
|-------------------|---|
| 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 | <p>This command administratively sets the router to operate in the overload state for a specific time period, in seconds, or indefinitely.</p> <p>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.</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 the router is overloaded or used prior to executing a shutdown command to divert traffic around the router.</p> <p>The no form of the command causes the router to exit the overload state.</p> |
| 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 | <p>This command configures the interval between LSP packets are sent from the interface.</p> <p>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.</p> <p>The no form of the command reverts to the default value.</p> |
| 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 | <p>This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.</p> <p>The no form of the command reverts to the default value.</p> |
| 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 | <p>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.</p> <p>A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group.</p> <p>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.</p> <p>Up to 30 split horizon groups can be defined per VPLS instance. Half are supported in i-VPLS.</p> |

The **no** form of the command removes the group name from the configuration.

| | |
|-------------------|---|
| Parameters | <p><i>group-name</i> — Specifies the name of the split horizon group to which the SDP belongs.</p> <p><i>residential-group</i> — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:</p> <ul style="list-style-type: none"> a) SAPs which are members of this Residential Split Horizon Group will have: <ul style="list-style-type: none"> – 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: <ul style="list-style-type: none"> – 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) <p>Default A split horizon group is by default not created as a residential-group.</p> |
|-------------------|---|

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 | <p>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.</p> <p>The no version of the command means a SAP over a LAG that is not operational will be removed from the forwarding process.</p> <p>Default no process-cpm-traffic-on-sap-down</p> |

pppoe-policy

| | |
|--------------------|---|
| Syntax | pppoe-policy <i>pppoe-policy-name</i> no pppoe-policy |
| Context | config>service>vpls>sap |
| Description | This command specifies an existing PPPoE policy. These policies are referenced from interfaces configured for PPPoE. Multiple PPPoE policies may be configured. |
| Default | none |
| Parameters | <i>pppoe-policy-name</i> — Specifies an existing PPPoE policy name up to 32 characters in length. |

auto-learn-mac-protect

| | |
|--------------------|---|
| Syntax | [no] auto-learn-mac-protect |
| Context | config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp config>service>vpls>split-horizon-group config>service>vpls>endpoint config>service>pw-template config>service>pw-template>split-horizon-group |
| Description | <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 This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the `mac-protect` command or automatically added using the `auto-learn-mac-protect` command. While enabled all packets entering the configured SAP, spoke-SDP, mesh-SDP, or any SAP that is part of the configured split horizon group (SHG) will be verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the `restrict-protected-src` command, namely:

- No parameter

The packet will be discarded, an alarm will be generated and the SAP, spoke-SDP or mesh-SDP will be set operationally down. The SAP, spoke-SDP or mesh-SDP must be shutdown and enabled (no shutdown) for this state to be cleared.

- alarm-only

The packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP/spoke-SDP/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. “**restrict-protected-src discard-frame**” can only be enabled on SAPs on FP2 or later hardware or on SDPs where all network interfaces are on FP2 or later hardware.

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 | <p>This command indicates how the system will forward packets destined to an unprotected MAC address, either manually added using the <code>mac-protect</code> command or automatically added using the <code>auto-learn-mac-protect</code> 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.</p> <p>If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with <code>restrict-unprotected-dst</code> enabled, it will be flooded.</p> |
| Default | no restrict-unprotected-dst |

vpls-group

| | |
|--------------------|---|
| Syntax | [no] vpls-group <i>id</i> |
| Context | config>service>vpls |
| Description | <p>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.</p> |
| Default | no vpls-group |
| Parameters | <i>id</i> — Specifies the ID associated with the VPLS group. |
| Values | 1 — 4094 |

service-range

| | |
|--------------------|---|
| Syntax | service-range <i>startid-endid</i> [start-vlan-id <i>startvid</i>] no service-range <i>startid-endid</i> |
| Context | config>service>vpls>vpls-group |
| Description | <p>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.</p> <p>Modifications of the service id and vlan ranges are allowed with the following restrictions.</p> <ul style="list-style-type: none"> • service-range increase can be achieved in two ways: <ul style="list-style-type: none"> – 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 | <i>startid-endid</i> — Specifies the range of service IDs. |
| Values | 1—2147483647 |
| | <i>startvid</i> — 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 <i>name/id</i> no vpls-template-binding |
| Context | config>service>vpls>vpls-group |
| Description | <p>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.</p> <p>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.</p> |
| Default | no vpls-template-binding |
| Parameters | <i>name/id</i> — 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>mvrp config>service>vpls>sap>mvrp |
| 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>mvrp>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>mvrp>mvrp |

| | |
|--------------------|--|
| Description | <p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of the command removes the specified group id.</p> |
| Default | no endstation-vid-group |
| Parameters | <p><i>id</i> — Specifies the range index.</p> <p>Values 1—4094</p> <p><i>starvid-endvid</i> — Specifies the range of VLANs to be controlled by MVRP.</p> <p>Values 1—4094</p> |

root-guard

| | |
|--------------------|--|
| Syntax | [no] root-guard |
| Context | config>service>vpls>sap>stp config>service>vpls>spoke-sdp>stp |
| Description | This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity. |
| Default | no root-guard |

static-host

| | |
|--------------------|--|
| Syntax | static-host ip <i>ip-address</i> [mac <i>ieee-address</i>] [create] static-host mac <i>ieee-address</i> [create] no static-host [ip <i>ip-address</i>>] mac <i>ieee-address</i>> no static-host all [force] no static-host ip <i>ip-address</i> |
| Context | config>service>vpls>sap |
| Description | This command configures a static host on this SAP. |
| Syntax | <p>ip <i>ip-address</i> — Specifies the IPv4 unicast address.</p> <p>mac <i>ieee-address</i> — Specify this optional parameter when defining a static host. Every static host definition must have at least one address defined, IP or MAC.</p> |

force — Specifies the forced removal of the static host addresses.

sla-profile sla-profile-name — This optional parameter is used to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

create — This keyword is mandatory while configuring a static host.

ancp-string

| | |
|--------------------|---|
| Syntax | ancp-string <i>ancp-string</i> no ancp-string |
| Context | config>service>vpls>sap>static-host |
| Description | This command specifies the ANCP string associated to this SAP host. |
| Parameters | <i>ancp-string</i> — Specifies the ANCP string up to 63 characters in length. |

app-profile

| | |
|--------------------|---|
| Syntax | app-profile <i>app-profile-name</i> no app-profile |
| Context | config>service>vpls>sap>static-host |
| Description | This command specifies an application profile name. |
| Parameters | <i>app-profile-name</i> — Specifies the application profile name up to 32 characters in length. |

inter-dest-id

| | |
|--------------------|--|
| Syntax | inter-dest-id <i>intermediate-destination-id</i> no inter-dest-id |
| Context | config>service>vpls>sap>static-host |
| Description | Specifies to which intermediate destination (for example a DSLAM) this host belongs. |
| Parameters | <i>intermediate-destination-id</i> — Specifies the intermediate destination ID. |

sla-profile

| | |
|----------------|---|
| Syntax | sla-profile <i>sla-profile-name</i> no sla-profile |
| Context | config>service>vpls>sap>static-host |

| | |
|--------------------|---|
| Description | This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the config>subscr-mgmt>sla-profile context. |
| Parameters | <i>sla-profile-name</i> — Specifies the SLA profile name. |

sub-profile

| | |
|--------------------|--|
| Syntax | sub-profile <i>sub-profile-name</i> no sub-profile |
| Context | config>service>vpls>sap>static-host |
| Description | This command specifies an existing subscriber profile name to be associated with the static subscriber host. |
| Parameters | <i>sub-profile-name</i> — Specifies the sub-profile name. |

subscriber

| | |
|--------------------|--|
| Syntax | subscriber <i>sub-ident</i> no subscriber |
| Context | config>service>vpls>sap>static-host |
| Description | This command specifies an existing subscriber identification profile to be associated with the static subscriber host. |
| Parameters | <i>sub-ident</i> — Specifies the subscriber identification/ |

subscriber-sap-id

| | |
|--------------------|--|
| Syntax | [no] subscriber-sap-id |
| Context | config>service>vpls>sap>static-host |
| Description | This command enables using the SAP ID as subscriber id. |
| Parameters | subscriber-sap-id — Specifies to use the sap-id as the subscriber-id. |

tod-suite

| | |
|--------------------|---|
| Syntax | tod-suite <i>tod-suite-name</i> 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. |

VPLS Service Commands

| | |
|-------------------|--|
| 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. |

trigger-packet

| | |
|--------------------|--|
| Syntax | trigger-packet [dhcp] [pppoe] [arp] [dhcp6] [ppp] no trigger-packet |
| Context | config>service>vpls>sap |
| Description | This command enables triggering packet to initiate RADIUS authentication that provides a service context. The authentication, together with the service context for this request, creates a managed SAP. The VLAN is the same as the triggering packet. This SAP behaves as a regular SAP but the configuration is not user-editable and not maintained in the configuration file. The managed SAP remains active as long as the session is active. |
| Default | none |
| Parameters | dhcp — Specifies whether the receipt of DHCP trigger packets on this VPLS SAP when the keyword capture-sap is specified in the sap command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of 'managed'. pppoe — Specifies whether the receipt of PPPoE trigger packets on this VPLS SAP when the keyword capture-sap is specified in the sap command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of 'managed'. arp — Indicates that ARP is the type of trigger packets for this entry. dhcp6 — Indicates that DHCP6 is the type of trigger packets for this entry. ppp — Indicates that PPP is the type of trigger packets for this entry. |

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

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 |


```
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 <i>multiplier</i> no hold-multiplier |
| Context | config>service>vpls>gsmp>group |
| Description | This command configures the hold-multiplier for the GSMP connections in this group. |
| Parameters | <i>multiplier</i> — Specifies the GSMP hold multiplier value. |
| Values | 1 — 100 |

keepalive

| | |
|--------------------|--|
| Syntax | keepalive <i>seconds</i> no keepalive |
| Context | config>service>vpls>gsmp>group |
| Description | This command configures keepalive values for the GSMP connections in this group. |

General Switch Management Protocol Commands

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 DHCP Commands

dhcp

| | |
|--------------------|--|
| Syntax | dhcp |
| Context | config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp |
| Description | This command enables the context to configure DHCP parameters. |

lease-populate

| | |
|--------------------|---|
| Syntax | lease-populate [<i>nmb-of-entries</i>] no lease-populate |
| Context | config>service>vpls>sap>dhcp |
| Description | <p>This command enables and disables dynamic host lease state management for VPLS SAPs. For VPLS, DHCP snooping must be explicitly enabled (using the snoop command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.</p> <p>The optional number-of-entries parameter is used to define the number of lease state table entries allowed for this SAP or IP interface. If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.</p> <p>The retained lease state information representing dynamic hosts may be used to:</p> <ul style="list-style-type: none">• populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding a new lease state entry or updating an existing lease state entry.• generate dynamic ARP replies if arp-reply-agent is enabled. |
| Default | no lease-populate |
| Parameters | <i>nbr-of-entries</i> — Specifies the number of DHCP leases allowed. Values 1 — 8000 |

option

| | |
|--------------------|---|
| Syntax | [no] option |
| Context | config>service>vpls>sap>dhcp config>service>vpls>sap>dhcp6 |
| Description | This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options. The no form of this command returns the system to the default. |
| Default | no option |

action

| | |
|--------------------|--|
| Syntax | action [<i>dhcp-action</i>] no action |
| Context | config>service>vpls>sap>dhcp>option |
| Description | This command configures the Relay Agent Information Option (Option 82) processing. The no form of this command returns the system to the default value. |
| Default | The default is to keep the existing information intact. |
| Parameters | <p><i>dhcp-action</i> — Specifies the DHCP option action.</p> <p>replace — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).</p> <p>drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.</p> <p>keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.</p> <p>The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.</p> <p>If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.</p> |

circuit-id

| | |
|----------------|--|
| Syntax | circuit-id [<i>ascii-tuple</i> <i>vlan-ascii-tuple</i>] |
| Context | config>service>vpls>sap>dhcp>option |

General Switch Management Protocol Commands

| | |
|--------------------|---|
| Description | <p>When enabled, the router sends an ASCII-encoded tuple in the circuit-id suboption of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by “ ”. If no keyword is configured, then the circuit-id suboption will not be part of the information option (Option 82).</p> <p>If disabled, the circuit-id suboption of the DHCP packet will be left empty.</p> |
| Default | no circuit-id |
| Parameters | <p>ascii-tuple — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.</p> <p>vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq encapsulated ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.</p> |

remote-id

| | |
|--------------------|---|
| Syntax | remote-id [mac string <i>string</i>] no remote-id |
| Context | config>service>vpls>sap>dhcp>option config>service>vpls>sap>dhcp6>option |
| Description | <p>This command specifies what information goes into the remote-id suboption in the DHCP Relay packet.</p> <p>If disabled, the remote-id suboption of the DHCP packet will be left empty.</p> <p>The no form of this command returns the system to the default.</p> |
| Default | no remote-id |
| Parameters | <p>mac — This keyword specifies the MAC address of the remote end is encoded in the suboption.</p> <p>string <i>string</i> — Specifies the remote-id.</p> |

vendor-specific-option

| | |
|--------------------|---|
| Syntax | [no] vendor-specific-option |
| Context | config>service>vpls>sap>dhcp>option config>service>ies>if>dhcp>option |
| Description | This command configures the vendor specific suboption of the DHCP relay packet. |

client-mac-address

| | |
|--------------------|---|
| Syntax | [no] client-mac-address |
| Context | config>service>vpls>sap>dhcp>option>vendor |
| Description | <p>This command enables the sending of the MAC address in the vendor specific suboption of the DHCP relay packet.</p> <p>The no form of the command disables the sending of the MAC address in the vendor specific suboption of the DHCP relay packet.</p> |

sap-id

| | |
|--------------------|---|
| Syntax | [no] sap-id |
| Context | config>service>vpls>sap>dhcp>option>vendor |
| Description | <p>This command enables the sending of the SAP ID in the vendor specific suboption of the DHCP relay packet.</p> <p>The no form of the command disables the sending of the SAP ID in the vendor specific suboption of the DHCP relay packet.</p> |

service-id

| | |
|--------------------|---|
| Syntax | [no] service-id |
| Context | config>service>vpls>sap>dhcp>option>vendor |
| Description | <p>This command enables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.</p> <p>The no form of the command disables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.</p> |

string

| | |
|--------------------|--|
| Syntax | [no] string <i>text</i> |
| Context | config>service>vpls>sap>dhcp>option>vendor |
| Description | <p>This command specifies the string in the vendor specific suboption of the DHCP relay packet.</p> <p>The no form of the command returns the default value.</p> |
| Parameters | <i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”). |

system-id

General Switch Management Protocol Commands

| | |
|--------------------|---|
| Syntax | [no] system-id |
| Context | config>service>vpls>sap>dhcp>option>vendor |
| Description | This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82. |

proxy-server

| | |
|--------------------|--|
| Syntax | proxy-server |
| Context | config>service>vpls>sap>dhcp |
| Description | This command configures the DHCP proxy server. |

emulated-server

| | |
|--------------------|---|
| Syntax | emulated-server <i>ip-address</i> no emulated-server |
| Context | config>service>vpls>sap>dhcp>proxy |
| Description | <p>This command configures the IP address which will be used as the DHCP server address in the context of this VPLS SAP. Typically, the configured address should be in the context of the subnet represented by the VPLS.</p> <p>The no form of of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.</p> |
| Parameters | <i>ip-address</i> — Specifies the emulated server address. |

lease-time

| | |
|--------------------|--|
| Syntax | lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] [radius-override] no lease-time |
| Context | config>service>vpls>sap>dhcp>proxy |
| Description | <p>This command defines the length of lease time that will be provided to DHCP clients. By default, the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.</p> <p>The no form of this command disables the use of the lease-time command. The local proxy server will use the lease time offered by either a RADIUS or DHCP server.</p> |
| Default | 7 days 0 hours 0 seconds |
| Parameters | <i>days</i> — Specifies the number of days that the given IP address is valid. Values 0 — 3650 <i>hours</i> — Specifies the number of hours that the given IP address is valid. |

Values 0 — 23

minutes — Specifies the number of minutes that the given IP address is valid.

Values 0 — 59

seconds — Specifies the number of seconds that the given IP address is valid.

Values 0 — 59

snoop

| | |
|--------------------|---|
| Syntax | [no] snoop |
| Context | config>service>vpls>sap>dhcp6 config>service>vpls>sap>dhcp config>service>vpls>spoke-sdp>dhcp config>service>vpls>mesh-sdp>dhcp |
| Description | <p>This command enables DHCP snooping of DHCP messages on the SAP or SDP. Enabling DHCP snooping on VPLS interfaces (SAPs and SDP bindings) is required where DHCP messages important to lease state table population are received, or where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.</p> <p>Use the no form of the command to disable DHCP snooping on the specified VPLS SAP or SDP binding.</p> |
| Default | no snoop |

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 726). 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. 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 This command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

General Switch Management Protocol Commands

The hello time parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time can also be used to calculate the forward delay. See [auto-edge on page 726](#).

The **no** form of this command returns the hello time to the default value.

| | |
|-------------------|---|
| Default | 2 seconds |
| Parameters | <i>hello-time</i> — The hello time for the STP instance in seconds. |
| Values | 1 — 10 |

hold-count

| | |
|--------------------|---|
| Syntax | hold-count <i>BDPU tx hold count</i> no hold-count |
| Context | config>service>vpls>stp config>template>vpls-template>stp |
| Description | This command configures the peak number of BPDUs that can be transmitted in a period of one second. The no form of this command returns the hold count to the default value |
| Default | 6 |
| Parameters | <i>BDPU tx hold count</i> — The hold count for the STP instance in seconds. |
| Values | 1 — 10 |

link-type

| | |
|--------------------|---|
| Syntax | link-type {pt-pt shared} no link-type |
| Context | config>service>vpls>sap>stp config>service>vpls>spoke-sdp>stp |
| Description | This command instructs STP on the maximum number of bridges behind this SAP or spoke SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP or spoke SDPs should all be configured as shared, and timer-based transitions are used. The no form of this command returns the link type to the default value. |
| Default | pt-pt |

mst-instance

| | |
|--------------------|---|
| Syntax | mst-instance <i>mst-inst-number</i> |
| 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 | <i>mst-inst-number</i> — Specifies an existing Multiple Spanning Tree Instance number. |
| Values | 1 — 4094 |

mst-path-cost

| | |
|--------------------|---|
| Syntax | mst-path-cost <i>inst-path-cost</i> no mst-path-cost |
| Context | config>service>vpls>sap>stp>mst-instance |
| Description | This commands 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 | <i>inst-path-cost</i> — 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 <i>stp-priority</i> no mst-priority |
| Context | config>service>vpls>sap>stp>mst-instance |
| Description | This commands 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 | <i>stp-priority</i> — 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 434 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

| | |
|----------------|--|
| Syntax | [no] mst-instance mst-inst-number |
| Context | config>service>vpls>stp |

| | |
|--------------------|---|
| Description | This command creates the context to configure MST instance (MSTI) related parameters. Up to 16 instances will be supported by MSTP. The instance 0 is mandatory by protocol and therefore, it cannot be created by the CLI. The software will maintain this instance automatically. |
| Default | none |
| Parameters | <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 | <p>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.</p> <p>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.</p> <p>The no form of this command sets the bridge-priority to its default value.</p> |
| 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 | <p>This command specifies a range of VLANs associated with a certain MST-instance. This range applies to all SAPs of the mVPLS.</p> <p>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.</p> <p>The no form of this command removes the vlan-range from given mst-instance.</p> |
| Parameters | <i>vlan-range</i> — 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 <i>hops-count</i> no mst-max-hops |
| Context | config>service>vpls>stp |
| Description | <p>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 <<i>max-hops</i>>. 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.</p> <p>The no form of this command sets the <i>hops-count</i> to its default value.</p> |
| Default | 20 |
| Parameters | <i>hops-count</i> — Specifies the maximum number of hops. |
| | Values 1 — 40 |

mst-name

| | |
|--------------------|--|
| Syntax | mst-name <i>region-name</i> no mst-name |
| Context | config>service>vpls>stp |
| Description | <p>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.</p> <p>The no form of this command removes <i>region-name</i> from the configuration.</p> |
| Default | no mst-name |
| Parameters | <i>region-name</i> — Specifies an MST-region name up to 32 characters in length. |

mst-revision

| | |
|----------------|--|
| 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 | <p><i>revision-number</i> — Specifies the MSTP region revision number to define the MSTP region.</p> <p>Values 0 — 65535</p> |

path-cost

| | |
|--------------------|---|
| Syntax | <p>path-cost <i>sap-path-cost</i></p> <p>no path-cost</p> |
| Context | <p>config>service>vpls>sap>stp</p> <p>config>service>vpls>spoke-sdp>stp</p> |
| 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 7450 ESS 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 | <p>config>service>vpls>sap>stp</p> <p>config>service>vpls>spoke-sdp>stp</p> |
| 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 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.</p> |

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 | <p>The bridge-priority command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.</p> <p>The no form of this command returns the bridge priority to the default value.</p> |
| 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 | <p>This command configures the Alcatel-Lucent Spanning Tree Protocol (STP) priority for the SAP or spoke SDP.</p> <p>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.</p> <p>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.</p> <p>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.</p> |

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][capture-sap] [create] [eth-ring ring-index]
[root-leaf-tag | leaf-ac]
no sap sap-id
```

Context config>service>vpls

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7450. 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 or SONET/SDH ports. 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 1319](#) for command syntax.

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

eth-ring — When used with Ethernet Rings control the split-horizon-group accepts the major ring instance "value". The split horizon group prevents loops in the cases where a Ethernet Virtual

Ring is miss configured on the main ring. Each path a and path b major ring are configured in the group and associated with the sub-ring control instance in the VPLS service.

ring-index — Specifies the ring index of the Ethernet ring.

root-leaf-tag — specifies a SAP as a root leaf tag SAP. Only SAPs of the form dot1q (for example, 1/1/1:X) or qinq (for example, 1/1/1:X.Y, 1/1/1:X.*) are supported. The default E-Tree SAP type is a root AC, if *root-leaf-tag* (or *leaf-ac*) is not specified at SAP creation. This option is only available when the VPLS is designated as an E-Tree VPLS.

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.

split-horizon-group *group-name* — Specifies the name of the split horizon group to which the SAP belongs.

capture-sap — Specifies a capturing SAP in which triggering packets will be sent to the CPM. Non-triggering packets captured by the capture SAP will be dropped.

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 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 749) has been reached. If max-nbr-mac-addr has not been set for</p> |

General Switch Management Protocol Commands

the SAP or spoke SDP, enabling discard-unknown-source has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.

Default **no discard-unknown**

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</i> tag <i>qtag</i> [<i>qtag</i>] no path <i>path-index</i> |
| Context | config>service>vpls>sap>eth-tunnel |

| | |
|--------------------|--|
| Description | This command configures Ethernet tunnel SAP path parameters. The no form of the command removes the values from the configuration. |
| Default | none |
| Parameters | <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 <i>mac-address</i>] primary-vlan-enable [vlan <i>vlan-id</i>] 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 | <p><i>mac-address</i> — Specifies the MAC address of the MEP.</p> <p>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.</p> <p>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.</p> <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 cannot be changed from or to primary vlan functions without first being deleted. 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> |
| Default | no mip |

mip

| | |
|--------------------|--|
| Syntax | mip primary-vlan-enable [vlan <i>vlan-id</i>] 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. |

| | |
|-------------------|--|
| 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>vlan-id — 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 [[/level /level ...]] no client-meg-level |
| Context | config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable config>service>vpls>spoke-sdp>eth-cfm>mep>ais-enable |

| | |
|--------------------|---|
| Description | This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level. |
| Parameters | <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 |

| | |
|--------------------|---|
| 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> [priority <i>priority</i>] [data-length <i>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. |

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 <i>bit-errors</i> |
| Context | config>service>vpls>mesh-sdp |
| Description | This command specifies the lowest priority defect that is allowed to generate a fault alarm. |
| Default | 1 |
| Parameters | <i>bit-errors</i> — Specifies the lowest priority defect. |
| Values | 0 — 11840 |

fault-propagation-enable

| | |
|--------------------|---|
| Syntax | fault-propagation-enable { <i>use-if-tlv</i> <i>suspend-ccm</i> } 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 | <i>use-if-tlv</i> — Specifies to use the interface TLV. <i>suspend-ccm</i> — Specifies to suspend the continuity check messages. |

low-priority-defect

| | | | | | | | |
|----------------------|---|---------------|--|----------------------|--|-------------------|--|
| Syntax | low-priority-defect { <i>allDef</i> <i>macRemErrXcon</i> <i>remErrXcon</i> <i>errXcon</i> <i>xcon</i> <i>noXcon</i> } | | | | | | |
| 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 | <table> <tr> <td><i>allDef</i></td><td>DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM</td></tr> <tr> <td><i>macRemErrXcon</i></td><td>Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM</td></tr> <tr> <td><i>remErrXcon</i></td><td>Only DefRemoteCCM, DefErrorCCM, and DefXconCCM</td></tr> </table> | <i>allDef</i> | DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM | <i>macRemErrXcon</i> | Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM | <i>remErrXcon</i> | Only DefRemoteCCM, DefErrorCCM, and DefXconCCM |
| <i>allDef</i> | DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM | | | | | | |
| <i>macRemErrXcon</i> | Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM | | | | | | |
| <i>remErrXcon</i> | 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 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 disabled 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). Note that MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit. |
| Default | When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise, MAC pinning is not enabled by default. |

max-nbr-mac-addr

| | |
|--------------------|---|
| Syntax | max-nbr-mac-addr <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. |

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 737](#)), 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 | <i>table-size</i> — Specifies the maximum number of learned and static entries allowed in the FDB of this service. |
| Values | 1 — 511999 Chassis-mode D limit: 511999 |

mc-endpoint

| | |
|--------------------|--|
| Syntax | mc-endpoint <i>mc-ep-id</i> 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 | <i>mc-ep-id</i> — Specifies a multi-chassis endpoint ID. |
| Values | 1 — 4294967295 |

mc-ep-peer

| | |
|--------------------|--|
| Syntax | mc-ep-peer <i>name</i> mc-ep-peer <i>ip-address</i> 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 | <i>name</i> — Specifies the name of the multi-chassis end-point peer. <i>ip-address</i> — Specifies the IP address of multi-chassis end-point peer. |

msap-defaults

| | |
|--------------------|--|
| Syntax | msap-defaults |
| Context | config>service>vpls>sap |
| Description | This command configures the msap-defaults. |

service

| | |
|--------------------|---|
| Syntax | [no] service <i>service-id</i> |
| Context | config>service>vpls>sap>msap-defaults |
| Description | This command sets default service for all subscribers created based on trigger packets received on the given capture SAP in case the corresponding VSA is not included in RADIUS authentication response. This command is applicable to capture SAP only. |
| Default | no service. |

policy

| | |
|--------------------|---|
| Syntax | policy <i>msap-policy-name</i> no policy |
| Context | config>service>vpls>sap>msap-defaults |
| Description | This command sets default msap-policy for all subscribers created based on trigger packets received on the given capture-sap in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only. |
| Default | no policy |

multi-service-site

| | |
|--------------------|--|
| Syntax | multi-service-site <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> |

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.

precedence

Syntax **precedence** [*precedence-value* | primary]
no precedence

Context config>service>vpls>spoke-sdp

Description This command configures the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint. When an SDP bind goes down, the next highest precedence SDP bind begins forwarding traffic.

Parameters *precedence-value* — Specifies the precedence of this SDP bind.

Values 1 — 4

primary — Assigns this as the primary spoke-sdp.

static-isid

Syntax [**no**] **static-isid range** *entry-id isid* [**to** *isid*] [**create**]

Context config>service>vpls><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 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** *vlan-range*
- Context** config>service>vpls>sap>managed-vlan-list
- Description** This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.
- This command is only valid when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or qinq, or on a Sonet/SDH port with encapsulation type of bcp-dot1q.
- To modify the range of VLANs, first the new range should be entered and afterwards the old range removed. See [Modifying VPLS Service Parameters on page 613](#).
- Default** None
- Parameters** *vlan-range* — Specify the VLAN start value and VLAN end value. The end-vlan must be greater than start-vlan. The format is <start-vlan>-<end-vlan>

Values start-vlan: 0 — 4094
 end-vlan: 0 — 4094

VPLS Filter and QoS Policy Commands

egress

| | |
|--------------------|---|
| Syntax | egress |
| Context | config>service>vpls>sap |
| Description | <p>This command enables the context to configure egress filter policies.</p> <p>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.</p> |

ingress

| | |
|--------------------|---|
| Syntax | ingress |
| Context | config>service>vpls>sap |
| Description | <p>This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.</p> <p>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.</p> |

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 enabled (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 per IOM 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> |

hsmdda-queue-override

| | |
|----------------|-----------------------------------|
| Syntax | [no] hsmdda-queue-override |
| Context | config>service>vpls>sap>egress |

Description This command enables the context to configure HSMDBA 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 HSMDBA 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 HSMDBA 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 | <p>add <i>add-bytes</i> — 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.</p> <p>Values 0 — 31</p> <p>subtract <i>sub-bytes</i> — 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.</p> <p>Values 1 — 64</p> |
|-------------------|---|

slope-policy

| | |
|--------------------|--|
| Syntax | slope-policy <i>hsmda-slope-policy-name</i> 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 <i>pir-rate</i> no rate |
| Context | config>service>vpls>sap>egress>hsmda-queue-over |
| Description | This command specifies the administrative PIR by the user. |
| Parameters | <p><i>pir-rate</i> — Configures the administrative PIR specified by the user.</p> <p>Values 1 — 40000000, max</p> |

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 | <i>group-name</i> — The <i>group-name</i> 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 | <p>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.</p> <p>The no form of this command disables the command.</p> |
| 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 | <p>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.</p> <p>The no form of the command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.</p> |
| 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 | <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 *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 | <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><i>level</i> — 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> |
| Values | 1 — 8 |

mbs-contribution

| | |
|--------------------|---|
| Syntax | mbs-contribution <i>size</i> [bytes kilobytes] |
| Context | <pre>config>service>vpls>sap>egress</pre> <pre>config>service>vpls>sap>ingress</pre> |
| 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 | <p>bytes — This keyword signifies that size is expressed in bytes.</p> <p>kilobytes — The optional kilobytes keyword signifies that size is expressed in kilobytes.</p> |
| Values | 0 – 16777216 or default |

policer-control-policy

| | |
|--------------------|--|
| Syntax | policer-control-policy <i>policy-name</i> [create] no policer-control-policy |
| Context | <pre>config>service>vpls>sap>egress</pre> <pre>config>service>vpls>sap>ingress</pre> |
| Description | <p>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</p> |

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) on an Ethernet MDA attached to an IOM3-XP or IMM module.

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> |

- 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 | <p>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.</p> <p>The no packet-byte-offset command is used to restore the policer? packet-byte-offset setting to the policy defined value.</p> |
| Default | no packet-byte-offset |
| Parameters | <p>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.</p> <p>Values 1 — 31</p> <p>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.</p> <p>Values 1 — 64</p> |

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 | <p>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.</p> <p>The no rate command is used to restore the policy defined metering and profiling rate to a policer.</p> |
| Parameters | <p>{<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</p> |

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 — 20000000000, max

stat-mode

| | |
|--------------------|---|
| Syntax | stat-mode <i>stat-mode</i> no stat-mode |
| Context | config>service>vpls>sap>egress>policer-override>policer config>service>vpls>sap>ingress>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> |

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 | <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. |
| | When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones. |
| Values | Multipoint or not present. |
| Default | Present (the queue is created as non-multipoint). |
| | fp-redirect-group — 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. |

queue-group-name — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid ingress queue group template name, configured under `config>qos>queue-group-templates`.

instance-id — Specifies the instance of the named queue group to be created on the IOM/IMMXMA ingress forwarding plane.

qos

| | |
|--------------------|---|
| Syntax | qos policy-id [port-redirect-group queue-group-name instance instance-id] no qos |
| Context | config>service>vpls>sap>egress |
| Description | <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 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 instance-id — 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 | <p>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.</p> <p>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.</p> <p><i>adaptation-rule</i> — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.</p> |
| Values | <p>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.</p> <p>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.</p> <p>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.</p> |

avg-frame-overhead

| | |
|--------------------|---|
| Syntax | avg-frame-overhead percent no avg-frame-overhead |
| Context | config>service>vpls>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. |

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

| | |
|--------------------|---|
| 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</i> [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 — 20000000000, **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 | <p>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.</p> <p>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.</p> |
| Default | no adaptation-rule |

| | |
|-------------------|---|
| Parameters | <p>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.</p> <p>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.</p> <p><i>adaptation-rule</i> — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.</p> |
| Values | <p>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.</p> <p>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.</p> <p>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.</p> |

avg-frame-overhead

| | |
|--------------------|---|
| Syntax | avg-frame-overhead percent no avg-frame-overhead |
| Context | config>service>vpls>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. |

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 | <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 — 100 |

cbs

| | |
|--------------------|--|
| Syntax | cbs <i>size-in-kbytes</i> no cbs |
| 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 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> <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 drop packets.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.</p> <p>The no form of this command returns the CBS size to the default value.</p> |
| Default | no cbs |
| Parameters | <i>size-in-kbytes</i> — 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 |

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 | <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 | <p>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.</p> |
| 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 — 100000000</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> <p>Values 0 — 100000000, max, sum</p> <p>Default 0</p> |

wred-queue-policy

| | |
|----------------|---|
| Syntax | <p>wred-queue-policy <i>slope-policy-name</i> no wred-queue-policy</p> |
| Context | <p>config>service>vpls>sap>egress>queue-override>queue</p> |

| | |
|--------------------|--|
| Description | <p>The <code>wred-queue-policy</code> 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 <i>queue-id</i> must be a WRE- enabled queue to be successful.</p> <p>The no form of the command removes the slope policy override for the WRED queue on the egress SAP.</p> |
| Parameters | <p><i>slope-policy-name</i> — 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.</p> |

scheduler-override

| | |
|--------------------|--|
| Syntax | [no] scheduler-override |
| Context | <p>config>service>vpls>sap>egress</p> <p>config>service>vpls>sap>ingress</p> |
| Description | <p>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.</p> |

scheduler

| | |
|--------------------|--|
| Syntax | <p>scheduler <i>scheduler-name</i></p> <p>no scheduler <i>scheduler-name</i></p> |
| Context | config>service>vpls>sap>egress>sched-override |
| Description | <p>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.</p> <p>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 <i>scheduler-name</i> already exists within the policy tier level (regardless of the inclusion of the keyword <code>create</code>), 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).</p> <p>If the <i>scheduler-name</i> exists within the policy on a different tier (regardless of the inclusion of the keyword <code>create</code>), an error occurs and the current CLI context will not change.</p> |

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 <i>weight</i>] [cir-weight <i>cir-weight</i>] no parent |
| Context | config>service>vpls>sap>ingress>sched-override>scheduler config>service>vpls>sap>egress>sched-override>scheduler |
| Description | <p>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.</p> <p>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.</p> <p>The no form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.</p> |
| 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 | <p>rate <i>pir-rate</i> [cir <i>cir-rate</i>]</p> <p>no rate</p> |
| Context | config>service>vpls>sap>egress>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 maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.</p> <p>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.</p> |

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 — 100000000 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 — 100000000, **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 19](#) defines the default behavior for Dot1P evaluation.

Table 19: 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 20: 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 1/1:10.***.
- For QinQ SAP type, both inner (bottom) and outer (top) tags are explicitly specified. For example, **sap 1/1:10.100**.

policer-control-policy

| | |
|--------------------|--|
| Syntax | policer-control-policy <i>policy-name</i> [create] no policer-control-policy |
| Context | config>service>vpls>sap>egress |
| Description | <p>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.</p> <p>Policer Control Policy Instances</p> <p>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.</p> |

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> |

authentication-policy

| | |
|--------------------|--|
| Syntax | authentication-policy <i>name</i> no authentication-policy |
| Context | config>service>vpls>sap |
| Description | This command defines which subscriber authentication policy must be applied when a DHCP message is received on the interface. The authentication policies must already be defined. The policy will only be applied when DHCP snooping is enabled on the SAP. |

accounting-policy

| | |
|--------------------|---|
| Syntax | accounting-policy <i>acct-policy-id</i> no accounting-policy |
| Context | config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp config>service>vpls>sap |
| Description | <p>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 <i>policy-id</i> does not exist, an error message is generated.</p> <p>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.</p> <p>The no form of this command removes the accounting policy association from the SAP or SDP, and the accounting policy reverts to the default.</p> |
| 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 |

app-profile

| | |
|--------------------|--|
| Syntax | app-profile <i>app-profile-name</i> no app-profile |
| Context | config>service>vpls>spoke-sdp |
| 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. |

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 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 |

| | |
|--------------------|---|
| Description | <p>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</p> <p>Following SAP attributes will be set in the instantiated saps (no configuration allowed):</p> <p>description: "Sap <sap-id> controlled by MVRP service <svc id>" – auto generated</p> <p>shutdown: no shutdown</p> |
| Parameters | <p><i>name/id</i> — Specifies the name in ASCII or the template ID.</p> <p>Values 1..2147483647</p> |

mac-move-level

| | |
|--------------------|--|
| Syntax | <p>mac-move-level {primary secondary}</p> <p>no mac-move-level</p> |
| Context | config>service>template>vpls-sap-template |
| Description | <p>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.</p> |
| Default | no mac-move-level; SAP is populated as a tertiary-port |

temp-flooding

| | |
|--------------------|---|
| Syntax | <p>temp-flooding flood-time</p> <p>no temp-flooding</p> |
| Context | <p>config>service>vpls</p> <p>config>service>template>vpls-template</p> |
| Description | <p>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.</p> |

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 <i>p2mp-lsp-template-name</i></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. |

| | |
|----------------|---------------|
| 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 | <p>This command configures the node to operate as both root and leaf of the I-PMSI in a given VPLS/B-VPLS instance.</p> <p>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.</p> <p>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..</p> <p>The no version of this command re-instates the default value.</p> |

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 | <p><i>p2mp-lsp-template-name</i> — The name of the P2MP LSP template. This is a string of 32 characters maximum.</p> <p>Default None</p> |

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 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.</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 <i>sdp-id</i> is bound to a service. |
| Special Cases | <p>VPLS — Several SDPs can be bound to a VPLS service. Each SDP must use unique <i>vc-ids</i>. An error message is generated if two SDP bindings with identical <i>vc-ids</i> terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.</p> |
| Parameters | <p><i>sdp-id</i> — The SDP identifier.</p> <p>Values 1 — 17407</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. 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. <p>Values ether, vlan</p> <p>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)</p> <p>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.</p> <p>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.</p> <p>split-horizon-group <i>group-name</i> — Specifies the name of the split horizon group to which the SDP belongs.</p> <p>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.</p> <p>no endpoint — removes the association of a spoke SDP with an explicit endpoint name.</p> <p>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</p> |

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/VP RN 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 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 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 | <i>network-policy-id</i> — Specifies the network policy identification. The value uniquely identifies the policy on the system. |
| Values | 1 — 65535 |
| fp- redirect-group | <i>queue-group-name</i> — Specifies the name of the queue group template up to 32 characters in length. |
| ingress-instance | <i>instance-id</i> — 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 | <p>This command enables the context to configure MFIB-allowed MDA destinations.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>By default, the MDA inclusion list is empty.</p> <p>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.</p> |

mda

| | |
|--------------------|---|
| Syntax | [no] mda mda-id |
| Context | config>service>vpls>mesh-sdp>egress>mfib-allowed-mda-destinations config>service>vpls>spoke-sdp>egress>mfib-allowed-mda-destinations |
| Description | This command specifies an MFIB-allowed MDA destination for an SDP binding configured in the system. |
| Parameters | <i>mda-id</i> — Specifies an MFIB-allowed MDA destination. |
| Values | slot/mda slot: 1 — 10 mda: 1 — 2 |

mda

| | |
|--------------------|---|
| Syntax | [no] mda mda-id |
| Context | config>service>vpls>mesh-sdp>egress>mfib-allowed-mda-destinations config>service>vpls>spoke-sdp>egress>mfib-allowed-mda-destinations |
| Description | This command specifies an MFIB-allowed MDA destination for a configured SDP binding. |
| 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 egress-vc-label no vc-label [egress-vc-label] |
| Context | config>service>vpls>mesh-sdp>egress config>service>vpls>spoke-sdp>egress |
| Description | This command configures the egress VC label. |
| Parameters | <i>vc-label</i> — A VC egress value that indicates a specific connection. |
| Values | 16 — 1048575 |

vc-label

| | |
|----------------|---|
| Syntax | vc-label ingress-vc-label no vc-label [ingress-vc-label] |
| Context | config>service>vpls>mesh-sdp>ingress config>service>vpls>spoke-sdp>ingress |

| | |
|--------------------|--|
| Description | This command configures the ingress VC label. |
| Parameters | <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 |
| 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 0..4094 no vlan-vc-tag [0..4094] |
| 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 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.</p> <p>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.</p> <p>The no form of this command disables the command.</p> |
| Default | no vlan-vc-tag |
| Parameters | 0..4094 — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID. |

SAP Subscriber Management Commands

cpu-protection

| | |
|--------------------|--|
| Syntax | cpu-protection <i>policy-id</i> [mac-monitoring] no cpu-protection |
| Context | config>service>vpls>sap config>template>vpls-sap-template |
| Description | <p>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.</p> <p>If no CPU protection policy is assigned to a service SAP, then a the default policy is used to limit the overall-rate.</p> |
| Default | <p>cpu-protection 254 (for access interfaces)</p> <p>cpu-protection 255 (for network interfaces)</p> <p>The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.</p> |
| Parameters | <p><i>policy-id</i> — Specifies an existing CPU protection policy.</p> <p>Values 1 — 255</p> <p>mac-monitoring — When specified, the per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy.</p> |

default-msap-policy

| | |
|--------------------|--|
| Syntax | default-msap-policy <i>policy-name</i> no default-msap-policy |
| Context | config>service>vpls>sap |
| Description | <p>This command specifies an existing managed SAP policy. Managed SAPs allow the use of policies and a SAP template for the creation of a SAP. Managed SAP policies are created in the config>subscr-mgmt context. This command is only applicable to SAPs created as a capture-sap.</p> |
| Default | none |
| Parameters | <i>msap-policy-name</i> — Specifies an existing managed SAP policy name up to 32 characters in length. |

sub-sla-mgmt

| | |
|--------------------|--|
| Syntax | [no] sub-sla-mgmt |
| Context | config>service>vpls>sap |
| Description | This command enables the context to configure subscriber management parameters for this SAP. |
| Default | no sub-sla-mgmt |

def-inter-dest-id

| | |
|--------------------|---|
| Syntax | def-inter-dest-id {string <i>string</i> use-top-q} no def-inter-dest-id |
| Context | config>service>vpls>sap>sub-sla-mgmt |
| Description | <p>This command specifies a default destination string for all subscribers associated with the SAP. The command also accepts the use-top-q flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.</p> <p>The no form of the command removes the default subscriber identification string from the configuration.</p> <p>no def-sub-id</p> |
| Default | no def-inter-dest-id |
| Parameters | <p>use-top-q — Derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.</p> <p>string <i>string</i> — Specifies the subscriber identification applicable for a subscriber host.</p> |

def-sla-profile

| | |
|--------------------|---|
| Syntax | def-sla-profile <i>default-sla-profile-name</i> no def-sla-profile |
| Context | config>service>vpls>sap>sub-sla-mgmt |
| Description | <p>This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the config>subscriber-mgmt>sla-profile context.</p> <p>An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.</p> <p>The no form of the command removes the default SLA profile from the SAP configuration.</p> |

| | |
|-------------------|---|
| Default | no def-sla-profile |
| Parameters | <i>default-sla-profile-name</i> — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the config>subscriber-mgmt>sla-profile context. |

def-sub-profile

| | |
|--------------------|--|
| Syntax | def-sub-profile <i>default-subscriber-profile-name</i> |
| Context | config>service>vpls>sap>sub-sla-mgmt |
| Description | <p>This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the config>subscriber-mgmt>sub-profile context.</p> <p>A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.</p> <p>The no form of the command removes the default SLA profile from the SAP configuration.</p> |
| Parameters | <i>default-sub-profile</i> — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the config>subscriber-mgmt>sub-profile context. |

mac-da-hashing

| | |
|--------------------|--|
| Syntax | [no] mac-da-hashing |
| Context | config>service>vpls>sap>sub-sla-mgmt |
| Description | <p>This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.</p> <p>This command is only meaningful if subscriber management is enabled and can be configured for this VPLS service.</p> |

multi-sub-sap

| | |
|--------------------|---|
| Syntax | multi-sub-sap [<i>subscriber-limit</i>] no multi-sub-sap |
| Context | config>service>vpls>sap>sub-sla-mgmt |
| Description | <p>This command configures the maximum number of subscribers for this SAP.</p> <p>The no form of this command returns the default value.</p> |
| Default | 1 |

Parameters *number-of-sub* — Specifies the maximum number of subscribers for this SAP.

Values 2 — 8000

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]
no non-sub-traffic

Context config>service>vpls>sap>sub-sla-mgmt>single-sub

Description This command configures non-subscriber traffic profiles. It is used in conjunction with the **profiled-traffic-only** command on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes the profiles and disables the feature.

Parameters **sub-profile** *sub-profile-name* — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile *sla-profile-name* — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

subscriber *sub-ident-string* — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's *sub-ident-string* is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the service destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

profiled-traffic-only

| | |
|--------------------|---|
| Syntax | [no] profiled-traffic-only |
| Context | config>service>vpls>sap>sub-sla-mgmt>single-sub |
| Description | <p>This command enables profiled traffic only for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).</p> <p>The no form of the command disables the command.</p> |

single-sub-parameters

| | |
|--------------------|--|
| Syntax | single-sub-parameters |
| Context | config>service>vpls>sap>sub-sla-mgmt |
| Description | This command enables the context to configure single subscriber parameters for this SAP. |

sub-ident-policy

| | |
|--------------------|--|
| Syntax | sub-ident-policy <i>sub-ident-policy-name</i> |
| Context | config>service>vpls>sap>sub-sla-mgmt |
| Description | <p>This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the config>subscriber-mgmt>sub-ident-policy context.</p> <p>Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.</p> <p>For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.</p> <p>When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.</p> <p>The no form of the command removes the default subscriber identification policy from the SAP configuration.</p> |
| Default | no sub-ident-policy |
| Parameters | <i>sub-ident-policy-name</i> — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the config>subscriber-mgmt>sub-ident-policy context. |

VPLS Multicast Commands

fast-leave

| | |
|--------------------|---|
| Syntax | [no] fast-leave |
| Context | config>service>vpls>sap>igmp-snooping config>service>vpls>spoke-sdp>igmp-snooping config>service>vpls>mesh-sdp>igmp-snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping |
| Description | <p>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').</p> <p>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.</p> |
| 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 either as a (*, g) or as one or more (s,g) records. When a static MLD or IGMP group is added, multicast data for that (*,g) or (s,g) is forwarded to the specific SAP or SDP without receiving any membership report from a host. |
| Default | none |
| Parameters | <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). |

- 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 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.

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/VP RN 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 ESS-Series 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 ESS-Series must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the pseudowire ID FEC element.

The **no** form of this command disables the use of the hash label.

| | |
|-------------------|--|
| Default | no hash-label |
| Parameters | signal-capability — Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The signal-capability option is not supported on a VPRN spoke-sdp. |

igmp-snooping

| | |
|--------------------|---|
| Syntax | igmp-snooping |
| Context | config>service>vpls config>service>vpls>sap config>service>vpls>spoke-sdp config>service>vpls>mesh-sdp |
| 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 |

| | |
|--------------------|--|
| 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 |

| | |
|--------------------|--|
| | <pre>config>service>vpls>sap>igmp-snooping cconfig>service>vpls>spoke-sdp>igmp-snooping</pre> |
| Description | <p>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.</p> <p>The no form of this command disables the check.</p> |
| Default | no max-num-grp-sources |
| Parameters | 1..32000 — Specifies the maximum number of multicast sources allowed to be tracked per group |

import

| | |
|--------------------|--|
| Syntax | <pre>import <i>policy-name</i> no import</pre> |
| Context | <pre>config>service>vpls>sap>igmp-snooping config>service>vpls>spoke-sdp>igmp-snooping config> service>vpls> mesh-sdp>igmp-snooping config>service>vpls>spoke-sdp>snooping config>service>vpls>mesh-sdp>snooping config>service>vpls>sap>mld-snooping config>service>vpls>spoke-sdp>mld-snooping config>service>vpls>mesh-sdp>mld-snooping</pre> |
| Description | <p>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.</p> <p>The no form of the command removes the policy association from the SAP or SDP.</p> |
| Default | no import — No import policy is specified. |
| Parameters | <i>policy-name</i> — 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 | <pre>last-member-query-interval <i>tenths-of-seconds</i> no last-member-query-interval</pre> |
| Context | <pre>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</pre> |

| | |
|--------------------|---|
| Description | This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP or SDP. |
| Default | 10 |
| Parameters | <i>seconds</i> — 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 | <p><i>bandwidth</i> — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).</p> <p>Values 0 — 2147483647</p> <p>mandatory-bw <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).</p> <p>If the <i>bandwidth</i> value is 0, no mandatory channels are allowed. If <i>bandwidth</i> is not configured, then all mandatory and optional channels are allowed.</p> <p>If the value of <i>mandatory-bw</i> is equal to the value of <i>bandwidth</i>, 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.</p> <p>The value of <i>mandatory-bw</i> should always be less than or equal to that of <i>bandwidth</i>. An attempt to set the value of <i>mandatory-bw</i> greater than that of <i>bandwidth</i>, will result in inconsistent value error.</p> <p>Values 0 — 2147483647</p> |

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. |

| | |
|-------------------|--|
| Parameters | <i>level-id</i> — Specifies has an entry for each multicast CAC policy constraint level configured on this system. |
| Values | 1 — 8 |
| | <i>bandwidth</i> — Specifies the bandwidth in kilobits per second (kbps) for the level. |
| Values | 1 — 2147483647 |

number-down

| | |
|--------------------|--|
| Syntax | number-down <i>number-lag-port-down</i> 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 | <i>number-lag-port-down</i> — 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 <i>count</i> 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 | <i>count</i> — 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 |
| Description | This command specifies whether a multicast router is attached behind this SAP or SDP. Configuring a SAP as an mrouter-port will have a double effect. Firstly, all multicast traffic received on another SAP or SDP will be copied to this SAP or SDP. Secondly, IGMP reports generated by the system as a result of someone joining or leaving a multicast group, will be sent to this SAP or SDP. If two multicast routers exist in the 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 or SDPs 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 or spoke SDP, 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 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 | <p>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.</p> <p>The configured query-response-interval must be smaller than the configured query-interval.</p> <p>If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.</p> |
| 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 | <p>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.</p> |
| 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 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 | <p>This command controls the number of attributes accepted on a per BVPLS basis. When the limit is reached, no new attributes will be registered.</p> <p>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.</p> | | | | |
| Default | maximum number of attributes | | | | |
| Parameters | <i>value</i> — Specifies the number of attributes accepted on a per BVPLS basis. | | | | |
| Values | <table> <tr> <td>ESS-6/7/12:</td><td>1 — 2047</td></tr> <tr> <td>ESS-1</td><td>1 — 1023</td></tr> </table> | ESS-6/7/12: | 1 — 2047 | ESS-1 | 1 — 1023 |
| ESS-6/7/12: | 1 — 2047 | | | | |
| ESS-1 | 1 — 1023 | | | | |

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

| | |
|----------------|---|
| 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> | |
| | config>service>vpls>sap>mrp config>service>vpls>spoke-sdp>mrp config>service>vpls>mesh-sdp>mrp | Context |
| 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 $\text{LeaveAllTime} < T < 1.5 * \text{leave-all-time}$ 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-sdp>mrp config>service>vpls>mesh-sdp>mrp | |
| Description | <p>This command instructs MMRP to use the mrp-policy specified in the command to control which Group BMAC attributes will be declared and registered on the egress SAP/Mesh-SDP/Spoke-SDP.</p> | |

The Group BMACs will be derived from the ISIDs using the procedure used in the PBB solution. The Group MAC = standard OUI with the last 24 bits being the ISID value. If the policy-name refers to a non-existing mrp-policy the command should return error. Changes to a mrp-policy are allowed and applied to the SAP/SDPs under which the policy is referenced.

| | |
|-------------------|--|
| Default | no mrp-policy is defined |
| Parameters | <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 The max number of MFIBs is 16000 for a 7450 router. |

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 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 <table> <tr> <td>Values</td><td>0 — 300</td></tr> <tr> <td>Default</td><td>90</td></tr> </table> | Values | 0 — 300 | Default | 90 |
| 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. <table> <tr> <td>Values</td><td>snooping, proxy</td></tr> <tr> <td>Default</td><td>proxy</td></tr> </table> | Values | snooping, proxy | Default | proxy |
| 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
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** *sap-id*
no to-sap

Context config>service>vpls>sap>igmp-snooping>mvr

Description In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behaviour) but to another SAP.

This command configures the SAP to which the multicast data needs to be copied.

Default no to-sap

Parameters *sap-id* — Specifies the SAP to which multicast channels should be copied. See [Common CLI Command Descriptions on page 1319](#) for command syntax.

VPLS DHCP and Anti-Spoofing Commands

anti-spoof

| | |
|--------------------|---|
| Syntax | anti-spoof { ip mac ip-mac } no anti-spoof |
| Context | config>service>vpls>sap |
| Description | <p>This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.</p> <p>The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (ip, mac, ip-mac) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The no form of the command disables anti-spoof filtering on the SAP.</p> |
| Default | no anti-spoof |
| Parameters | <p>ip — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof ip command will fail.</p> <p>mac — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the anti-spoof mac command will fail.</p> <p>ip-mac — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the anti-spoof ip-mac command will fail.</p> |

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 | <p>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.</p> <p>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.</p> <p>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.</p> |

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

| | |
|--------------------|--|
| 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

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.

Note: Efficient multicast replication will only be performed on IOMs that support chassis mode b. If an IOM does not support mode b operation, egress-multicast-group membership is ignored on that IOM's egress forwarding planes. The chassis need not be placed into mode b for efficient multicast replication to be performed on the capable IOMs.

| | |
|-------------------|---|
| 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 |

Egress Multicast Group Commands

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 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** *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*]

Context config>service>egress-multicast-group>sap-common-requirements

Description 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.

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.

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.

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.

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 **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

ipv6 *ipv6-filter-id* — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 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

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 | <p>dot1q — For an egress-multicast-group.</p> <p>null — If member SAPs are on a null encapsulated access port.</p> |
| 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 |
| | <p><i>ether-type</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.</p> <p>Values [0x0600 — 0xffff]: [1536 — 65535] in decimal or hex</p> |

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 | <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> |
| 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. |

- 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 | <i>value</i> — 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 | <p>This command administratively enables/disables the local BGP VPLS instance. On de-activation an MP-UNREACH-NLRI must be sent for the local NLRI.</p> <p>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.</p> |
| 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 <i>policy-id</i> [split-horizon-group <i>group-name</i>] [import-rt { <i>ext-community</i> , ...(up to 5 max)}] no pw-template-bind <i>policy-id</i> |
| Context | config>service>vpls>bgp-ad config>service>vpls>bgp |

Description This command binds the advertisements received with the route target (RT) that matches the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present the pw-template is used for all of them.

The pw-template-binding applies to both BGP-AD and BGP-VPLS if these features are enabled in the VPLS.

For BGP VPLS the following additional rules govern the use of pseudowire-template:

- On transmission the settings for the L2-Info extended community in the BGP Update are derived from the pseudowire template attributes. If multiple pseudowire templates (with or without import-rt) are specified for the same VPLS instance the first pw-template entry will be used.
- On reception the values of the parameters in the L2-Info extended community of the BGP Update are compared with the settings from the corresponding pw-template. The following steps are used to determine the local pw-template:
 - The RT values are matched to determine the pw-template.
 - If multiple pw-templates matches are found from the previous steps, the first configured pw-template entry will be considered.
 - If the values used for Layer 2 MTU or C Flag do not match the pseudowire setup fails.

The tools perform commands can be used to control the application of changes in pw-template for both BGP-AD and BGP-VPLS.

The **no** form of the command removes the values from the configuration.

Default none

Parameters *policy-id* — Specifies an existing policy ID.

Values 1 — 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 comand 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

| | |
|--------------------|--|
| Description | This command associates the context to which it is configured to the operational group specified in the <i>group-name</i> . The oper-group <i>group-name</i> 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 | <i>group-name</i> — Specifies a character string of maximum 32 ASCII characters identifying the group instance. |

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>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: <ul style="list-style-type: none"> • 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 <i>ext-community</i> — Specify communities allowed to be sent to remote PE neighbors. import <i>ext-community</i> — Specify communities allowed to be accepted from remote PE neighbors. |

vpls-id

| | | | | | | | | | | |
|--------------------|--|---|-----------|---|--|---------|---------|--|----------|-----------|
| Syntax | vpls-id <i>vpls-id</i> | | | | | | | | | |
| Context | config>service>vpls>bgp-ad | | | | | | | | | |
| Description | <p>This command configures the VPLS ID component that will be signaled in one of the extended community attributes (<i>ext-comm</i>).</p> <p>Values and format (6 bytes, other 2 bytes of type-subtype will be automatically generated)</p> | | | | | | | | | |
| Parameters | <p><i>vpls-id</i> — Specifies a globally unique VPLS ID for BGP auto-discovery in this VPLS service.</p> <table><tr><td>Values</td><td>vpls-id :</td><td><ip-addr:comm-val> <as-number:ext-comm-val></td></tr><tr><td></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> | Values | vpls-id : | <ip-addr:comm-val> <as-number:ext-comm-val> | | ip-addr | a.b.c.d | | comm-val | 0 — 65535 |
| 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 <i>policy-name</i> [<i>policy-name...</i> (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 <i>low-order-vsi-id</i> 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 | <i>low-order-vsi-id</i> — Specifies a unique VSI ID. |
| Values | 0— 4294967295 |

route-distinguisher

| | | | | | | | | | | | | | |
|-------------|--|----------------|---------|---------|--|----------|-----------|--------|-----------|-----------|--|--------------|----------------|
| Syntax | route-distinguisher rd route-distinguisher auto-rd no route-distinguisher | | | | | | | | | | | | |
| Context | config>service>vpls>bgp | | | | | | | | | | | | |
| Description | <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 sytem genrates 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. | | | | |
| | <table> <tr> <td>Values</td><td>3..10 minutes</td></tr> <tr> <td>Default</td><td>3 minutes</td></tr> </table> | Values | 3..10 minutes | Default | 3 minutes |
| Values | 3..10 minutes | | | | |
| Default | 3 minutes | | | | |

retry

| | |
|--------------------|---|
| 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 | [no] shutdown |
| Context | config>service>vpls>bgp-evpn.vxlan |
| Description | This command enables/disables the automatic creation of VXLAN auto-bindings by BGP-EVPN. |
| Default | shutdown |

VPLS Show Commands

active-subscribers

| | |
|--------------------|---|
| Syntax | active-subscribers summary active-subscribers [subscriber <i>sub-ident-string</i> [sap <i>sap-id</i> sla-profile <i>sla-profile-name</i>]] [detail] active-subscribers hierarchy [subscriber <i>sub-ident-string</i>] |
| Context | show>service |
| Description | This command displays active subscriber information. |
| Parameters | sap <i>sap-id</i> — Displays SAP information for the specified SAP ID. See Common CLI Command Descriptions on page 1319 for <i>sap-id</i> command syntax. sla-profile <i>sla-profile-name</i> — Displays information for the specified SLA profile name. summary — Displays active subscriber information in a brief format. subscriber <i>sub-ident-string</i> — Displays information for the specified subscriber identification string. hierarchy — Displays the subscriber hierarchy. detail — Displays detailed output. |

Sample Output

```
A:Dut-A# show service active-subscribers summary
=====
Active Subscriber table summary
=====
Total Count      : 6
=====
A:Dut-A#

A:Dut-A# show service active-subscribers hierarchy
=====
Active Subscriber hierarchy
=====
-- alcatel_100 (sub_default)
|
|-- sap:1/2/1:100 - sla:sla_default
|
|-- 10.100.1.3 - 00:10:00:00:00:01 (-/D/-)
|
|-- sap:1/2/1:101 - sla:sla_default
|
|-- 10.100.1.4 - 00:10:00:00:00:02 (-/D/-)
|
|
```

Show, Clear, Debug Commands

```
|-- sap:1/2/1:102 - sla:sla_default
|
|
|-- 10.100.1.5 - 00:10:00:00:00:03 (-/D/-)
|
|

-- alcatel_110 (sub_default)
|
|-- sap:1/2/1:110 - sla:sla_default
|
|
|-- 10.110.1.3 - 00:10:10:00:00:01 (-/D/-)
|
|
|-- sap:1/2/1:111 - sla:sla_default
|
|
|-- 10.110.1.4 - 00:10:10:00:00:02 (-/D/-)
|
|
|-- sap:1/2/1:112 - sla:sla_default
|
|
|-- 10.110.1.5 - 00:10:10:00:00:03 (-/D/-)
|
|

-- alcatel_120 (sub_default)
|
|-- sap:1/2/1:120 - sla:sla_default
|
|
|-- 10.120.1.3 - 00:10:20:00:00:01 (-/D/-)
|
|
|-- sap:1/2/1:121 - sla:sla_default
|
|
|-- 10.120.1.4 - 00:10:20:00:00:02 (-/D/-)
|
|
|-- sap:1/2/1:122 - sla:sla_prof120_VOIP
|
|
|-- 10.120.1.5 - 00:10:20:00:00:03 (-/D/-)
|
|

-- alcatel_130 (sub_default)
|
|-- sap:1/2/1:130 - sla:sla_default
|
|
|-- 10.130.1.3 - 00:10:30:00:00:01 (-/D/-)
|
|

-- alcatel_140 (sub_default)
|
|-- sap:1/2/1:140 - sla:sla_default
|
|
|-- 10.140.1.3 - 00:10:40:00:00:01 (-/D/-)
|
|

-- alcatel_80 (sub_default)
|
|-- sap:1/2/1:80 - sla:sla_default
|
|
|-- 10.80.1.3 - 00:80:00:00:00:01 (-/D/-)
|
|
|-- 10.80.1.4 - 00:80:00:00:00:02 (-/D/-)
|
|
```

```

| |-- 10.80.1.5 - 00:80:00:00:00:03 (-/D/-)
| |
-- alcatel_81 (sub_prof81)
|
|-- sap:1/2/1:80 - sla:sla_prof81_VOIP
| |
| |-- 10.80.1.6 - 00:80:00:00:00:04 (-/D/-)
| |
-- alcatel_90 (sub_default)
|
|-- sap:1/2/1:90 - sla:sla_default
| |
| |-- 10.90.1.3 - 00:90:00:00:00:01 (-/D/-)
| |
-- client_PC1 (sub_profPC1)
|
|-- sap:1/2/2:4000 - sla:sla_profPC1
| |
| |-- 0.0.0.0 - 00:00:00:00:00:00 (-/-/N)
| |
| |-- 10.24.1.253 - 00:13:21:67:a4:cd (-/D/-)
| |
|-- sap:lag-1 - sla:sla_profPC1
| |
| |-- 1.2.3.4 - 00:05:04:03:02:01 (S/-/-)
| |
-- static (sub_default)
|
|-- sap:1/2/1:80 - sla:sla_default
| |
| |-- 10.80.123.123 - 00:00:12:34:56:78 (S/-/-)
| |
=====
A:Dut-A#

A:Dut-A# show service active-subscribers subscriber alcatel_100 hierarchy
=====
Active Subscriber hierarchy
=====
-- alcatel_100 (sub_prof100)
|
|-- sap:1/2/1:101 - sla:sla_prof100_VOIP
| |
| |-- 10.100.1.4 - 00:10:00:00:00:02 (-/D/-)
| |
|
|-- sap:1/2/1:102 - sla:sla_default
| |
| |-- 10.100.1.5 - 00:10:00:00:00:03 (-/D/-)
| |
=====
A:Dut-A#

```

Show, Clear, Debug Commands

```
A:Dut-A# show service active-subscribers subscriber alcatel_100
=====
Active Subscribers
-----
Subscriber alcatel_100 (sub_default)
-----
(1) SLA Profile Instance sap:1/2/1:100 - sla:sla_default
-----
IP Address      MAC Address      Origin(*)
-----
10.100.1.3      00:10:00:00:00:01 -/D/-
-----
(2) SLA Profile Instance sap:1/2/1:101 - sla:sla_default
-----
IP Address      MAC Address      Origin(*)
-----
10.100.1.4      00:10:00:00:00:02 -/D/-
-----
(3) SLA Profile Instance sap:1/2/1:102 - sla:sla_default
-----
IP Address      MAC Address      Origin(*)
-----
10.100.1.5      00:10:00:00:00:03 -/D/-
=====
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
=====
A:Dut-A#

A:Dut-A# show service active-subscribers subscriber alcatel_100 sap 1/2/1:100 sla-
profile sla_default
=====
Active Subscribers
=====
Subscriber alcatel_100 (sub_default)
-----
(1) SLA Profile Instance sap:1/2/1:100 - sla:sla_default
-----
IP Address      MAC Address      Origin(*)
-----
10.100.1.3      00:10:00:00:00:01 -/D/-
=====
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
=====
A:Dut-A#

A:Dut-A# show service active-subscribers subscriber alcatel_100 sap 1/2/1:100 sla-
profile sla_default detail
=====
Active Subscribers
=====
Subscriber alcatel_100 (sub_default)
-----
I. Sched. Policy : service_all
E. Sched. Policy : service_all
Acct. Policy      : N/A                                Collect Stats : Disabled
-----
(1) SLA Profile Instance
    - sap:1/2/1:100 (VPLS 100)
    - sla:sla_default
```



```

-----
Host Limit           : No Limit
Ingress Qos-Policy   : 1000
Ingress Queuing Type : Service-queuing
Ingress Filter-Id    : N/A
Egress Qos-Policy    : 1000
Egress Filter-Id     : N/A
-----

```

```

-----
IP Address      MAC Address      Origin(*)
-----
10.100.1.3      00:10:00:00:00:01  -/D/-
-----

```

SLA Profile Instance statistics

```

-----
Packets      Octets
Off. HiPrio   : 0                0
Off. LowPrio  : 42361            8639977
Off. Uncolor  : 0                0

```

Queueing Stats (Ingress QoS Policy 1000)

```

Dro. HiPrio   : 0                0
Dro. LowPrio  : 6783             1392451
For. InProf   : 0                0
For. OutProf  : 35392            7211148

```

Queueing Stats (Egress QoS Policy 1000)

```

Dro. InProf   : 6599            1347340
Dro. OutProf  : 0                0
For. InProf   : 34364           7011246
For. OutProf  : 0                0

```

SLA Profile Instance per Queue statistics

```

-----
Packets      Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio   : 0                0
Off. LowPrio  : 0                0
Off. Uncolor  : 0                0
Dro. HiPrio   : 0                0
Dro. LowPrio  : 0                0
For. InProf   : 0                0
For. OutProf  : 0                0

```

Ingress Queue 2 (Unicast) (Priority)

```

Off. HiPrio   : 0                0
Off. LowPrio  : 0                0
Off. Uncolor  : 0                0
Dro. HiPrio   : 0                0
Dro. LowPrio  : 0                0
For. InProf   : 0                0
For. OutProf  : 0                0

```

Ingress Queue 3 (Unicast) (Priority)

```

Off. HiPrio   : 0                0
Off. LowPrio  : 42361            8639977
Off. Uncolor  : 0                0
Dro. HiPrio   : 0                0
Dro. LowPrio  : 6783             1392451
For. InProf   : 0                0
For. OutProf  : 35392            7211148

```

Ingress Queue 11 (Multipoint) (Priority)

```

Off. HiPrio   : 0                0

```

```
Off. LowPrio           : 0                0
Off. Uncolor           : 0                0
Dro. HiPrio            : 0                0
Dro. LowPrio           : 0                0
For. InProf            : 0                0
For. OutProf           : 0                0

Egress Queue 1
Dro. InProf            : 0                0
Dro. OutProf           : 0                0
For. InProf            : 0                0
For. OutProf           : 0                0

Egress Queue 2
Dro. InProf            : 0                0
Dro. OutProf           : 0                0
For. InProf            : 0                0
For. OutProf           : 0                0

Egress Queue 3
Dro. InProf            : 6599             1347340
Dro. OutProf           : 0                0
For. InProf            : 34364            7011246
For. OutProf           : 0                0
=====
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:Dut-A#
```

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> |

Output **Show Service Egress Command Output** — The following table describes show service egress label output fields.

| Label | Description |
|--------------------------|--|
| Svc Id | The ID that identifies a service. |
| Sdp Id | The ID that identifies an SDP. |
| Type | Indicates whether the SDP binding is a spoke or a mesh. |
| I. Lbl | The VC label used by the far-end device to send packets to this device in this service by the SDP. |
| E. Lbl | The VC label used by this device to send packets to the far-end device in this service by the SDP. |
| Number of bindings found | The total number of SDP bindings that exist within the specified egress label range. |

Sample Output

```
*A:ALA-12# show service egress-label 0 10000
=====
Martini Service Labels
=====
Svc Id      Sdp Id      Type I.Lbl      E.Lbl
-----
1           10:1        Mesh 0         0
1           20:1        Mesh 0         0
1           30:1        Mesh 0         0
1           100:1       Mesh 0         0
...
1           107:1       Mesh 0         0
1           108:1       Mesh 0         0
1           300:1       Mesh 0         0
1           301:1       Mesh 0         0
1           302:1       Mesh 0         0
1           400:1       Mesh 0         0
1           500:2       Spok 131070    2001
1           501:1       Mesh 131069    2000
100         300:100     Spok 0         0
200         301:200     Spok 0         0
300         302:300     Spok 0         0
400         400:400     Spok 0         0
-----
Number of Bindings Found : 23
=====
*A:ALA-12#
```

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. |
| 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. |
| Remote Age | 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 | 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. |
| MAC Pinning | Specifies whether MAC pinning is enabled in this service. |

| Label | Description (Continued) |
|---------------------------|--|
| 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

```
*A:ALA-12# show service fdb-info
=====
Forwarding Database (FDB) Information
=====
Service Id      : 700          Mac Move      : Disabled
Mac Move Rate   : 10          Mac Move Timeout : 10
Table Size      : 250         Total Count    : 0
Learned Count   : 0          Static Count   : 0
Remote Age      : 900         Local Age      : 300
High WaterMark  : 95%        Low Watermark  : 90%
Mac Aging       : Enabl       Relearn Only   : False
Service Id      : 725          Mac Move      : Disabled
Mac Move Rate   : 10          Mac Move Timeout : 10
Table Size      : 250         Total Count    : 0
Learned Count   : 0          Static Count   : 0
Remote Age      : 900         Local Age      : 300
High WaterMark  : 95%        Low Watermark  : 90%
Mac Learning    : Enabl       Discard Unknown : Dsabl
Mac Aging       : Enabl       Relearn Only   : False
Service Id      : 740          Mac Move      : Disabled
Mac Move Rate   : 10          Mac Move Timeout : 10
Table Size      : 250         Total Count    : 0
Learned Count   : 0          Static Count   : 0
Remote Age      : 900         Local Age      : 300
High WaterMark  : 95%        Low Watermark  : 90%
Mac Learning    : Enabl       Discard Unknown : Dsabl
Mac Aging       : Enabl       Relearn Only   : False
...
-----
Total Service FDBs : 7
Total FDB Configured Size : 1750
Total FDB Entries In Use : 0
=====
A:*A:ALA-48#
```

fdb-mac

| | |
|--------------------|---|
| Syntax | fdb-mac <i>ieee-address</i> [expiry] |
| Context | show>service |
| Description | This command displays the FDB entry for a given MAC address. |
| Parameters | <p><i>ieee-address</i> — The 48-bit MAC address for which to display the FDB entry 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.</p> <p>expiry — Shows the time until the MAC is aged out.</p> |
| 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 | <p>Static — FDB entries created by management.</p> <p>Learned — Dynamic entries created by the learning process.</p> <p>OAM — Entries created by the OAM process.</p> <p>H — Host, the entry added by the system for a static configured subscriber host.</p> <p>D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease.</p> <p>P — Indicates the MAC is protected by the MAC protection feature.</p> |

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 Age | Last Change |
|--------|-------------------|-------------------|-------------|-------------------|
| 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 <i>start-label</i> [<i>end-label</i>] |
| Context | show>service |
| Description | <p>Display services using the range of ingress labels.</p> <p>If only the mandatory <i>start-label</i> parameter is specified, only services using the specified label are displayed.</p> <p>If both <i>start-label</i> and <i>end-label</i> parameters are specified, the services using the range of labels X where <i>start-label</i> <= X <= <i>end-label</i> are displayed.</p> <p>Use the show router ldp bindings command to display dynamic labels.</p> |

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 or mesh. |
| I.Lbl | The ingress label used by the far-end device to send packets to this device in this service by the SDP. |
| 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. |

Sample Output

```
*A:ALA-12# show service ingress-label 0
=====
Martini Service Labels
=====
```

| Svc Id | Sdp Id | Type | I.Lbl | E.Lbl |
|--------|---------|------|-------|-------|
| 1 | 10:1 | Mesh | 0 | 0 |
| 1 | 20:1 | Mesh | 0 | 0 |
| 1 | 30:1 | Mesh | 0 | 0 |
| 1 | 50:1 | Mesh | 0 | 0 |
| 1 | 100:1 | Mesh | 0 | 0 |
| 1 | 101:1 | Mesh | 0 | 0 |
| 1 | 102:1 | Mesh | 0 | 0 |
| 1 | 103:1 | Mesh | 0 | 0 |
| 1 | 104:1 | Mesh | 0 | 0 |
| 1 | 105:1 | Mesh | 0 | 0 |
| 1 | 106:1 | Mesh | 0 | 0 |
| 1 | 107:1 | Mesh | 0 | 0 |
| 1 | 108:1 | Mesh | 0 | 0 |
| 1 | 300:1 | Mesh | 0 | 0 |
| 1 | 301:1 | Mesh | 0 | 0 |
| 1 | 302:1 | Mesh | 0 | 0 |
| 1 | 400:1 | Mesh | 0 | 0 |
| 100 | 300:100 | Spok | 0 | 0 |
| 200 | 301:200 | Spok | 0 | 0 |
| 300 | 302:300 | Spok | 0 | 0 |


```

400          400:400          Spok 0          0
-----
Number of Bindings Found : 21
-----
*A:ALA-12#

```

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.
qos-policy *qos-policy-id* — The ingress or egress 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
dyn-script — Displays dynamic service SAPs information.
authentication *auth-ply-name* — The session authentication policy for which to display matching SAPs.
sap-id — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 1319](#) for command syntax.
interface — Specifies matching SAPs with the specified IP interface.

ip-address — The IP address of the interface for which to display matching SAPs.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display matching SAPs.

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. |
| SapMTU | The SAP MTU value. |
| I.QoS | The SAP ingress QoS policy number specified on the ingress SAP. |
| I.MAC/IP | The MAC or IP filter policy ID applied to the ingress SAP. |
| Egr. Fltr | The filter policy ID applied to the egress SAP. |
| A.Pol | The accounting policy ID assigned to the SAP. |
| Adm | The administrative state of the SAP. |
| Opr | The actual state of the SAP. |

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

```
A:ALA-701# show service sap-using
=====
Service Access Points
=====
PortId          SvcId          Ing.   Ing.   Egr.   Egr.   Anti   Adm   Opr
                  QoS    Fltr   QoS    Fltr   Spoof
-----
1/1/3           10203041       1      ip4    1      none   none   Up    Up
1/1/4           10203042       1      none   1      ip4    none   Up    Up
-----
Number of SAPs : 2
-----
```

A:ALA-701#

show service sap-using process-cpm-traffic-on-sap-down

=====

SAP Ignore Sap Lag Down Information

=====

| SAP | Svc Id | Ignore SapLag Down |
|--------------|--------|--------------------|
| lag-1:1100.* | 1100 | enabled |

=====

Number of lag saps: 1

=====

Sample Output

The following is sample output for VPLS E-Tree configured SAPs.

```
*A:DutA# show service sap-using etree
=====
Etree SAP Information
=====
Svc Id      SAP                               Leaf-Tag  Root-   Leaf-Ac
                               leaf-tag
-----
2005        1/1/1:2005                          0         Disabled Enabled
2005        1/1/7:2006.200                      2007      Enabled  N/A
2005        1/1/7:0.*                          0         Disabled Disabled
2005        1/1/7:2005.*                       0         Disabled Disabled
2005        1/1/8:1                            0         Disabled Disabled
-----
Number of etree saps: 5
=====
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]
sdp [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]
sdp [*sdp-id* | **far-end** *ip-addr*] [**detail** | **keep-alive-history**]

Context show>service>id

Description This command displays information for the SDPs associated with the service.
If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Displays only information for the specified SDP ID. An SDP is a logical mechanism that ties a far-end ESS-Series 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 ESS-Series 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 or a mesh. |

| Label | Description (Continued) |
|---------------------|--|
| Split Horizon Group | Name of the split horizon group where the SDP belongs. |
| VC Type | Displays the VC type, ether or vlan. |
| VC Tag | Displays the explicit dot1Q value used when encapsulating to the SDP far end. |
| I. Lbl | The VC label used by the far-end device to send packets to this device in this service by the SDP. |
| Admin Path MTU | The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.) |
| Oper Path MTU | The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |
| Far End | Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP. |
| Delivery | Specifies the type of delivery used by the SDP: GRE or MPLS. |
| Admin State | The administrative state of this SDP. |
| Oper State | The operational state of this SDP. |
| Ingress Label | The label used by the far-end device to send packets to this device in this service by this SDP. |
| Egress Label | The label used by this device to send packets to the far-end device in this service by the SDP. |
| Last Changed | The date and time of the most recent change to the SDP. |
| Signaling | Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP. |
| Admin State | The administrative state of the Keepalive process. |
| Oper State | The operational state of the Keepalive process. |
| Hello Time | Specifies how often the SDP echo request messages are transmitted on this SDP. |
| Max Drop Count | Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault. |
| 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. |

| Label | Description (Continued) |
|---------------------|---|
| I. Fwd. Pkts. | Specifies the number of forwarded ingress packets. |
| I. Dro. Pkts | Specifies the number of dropped ingress packets. |
| E. Fwd. Pkts. | Specifies the number of forwarded egress packets. |
| E. Fwd. Octets | Specifies the number of forwarded egress octets. |
| Associated LSP List | When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS. |

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 [<i>sdp-id</i> [: <i>vc-id</i>] far-end <i>ip-address</i>] sdp-using etree |
| Context | show>service |
| Description | This command displays services using SDP or far-end address options. |
| Parameters | <i>sdp-id</i> — Displays only services bound to the specified SDP ID. Values 1 — 17407 <i>vc-id</i> — 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 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. |

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 |
|-------|---------|------|-----------|-----------|---------|---------|
| 1 | 300:1 | Mesh | 10.0.0.13 | Up | 131071 | 131071 |
| 2 | 300:2 | Spok | 10.0.0.13 | Up | 131070 | 131070 |
| 100 | 300:100 | Mesh | 10.0.0.13 | Up | 131069 | 131069 |
| 101 | 300:101 | Mesh | 10.0.0.13 | Up | 131068 | 131068 |
| 102 | 300:102 | Mesh | 10.0.0.13 | Up | 131067 | 131067 |

```
-----
Number of SDPs : 5
-----
*A:ALA-1#
```

Sample Output

The following is sample output for VPLS E-Tree configured SDP bindings.

```
*A:DutA# show service sdp-using etree
=====
Etree SDP-BIND Information
=====
```

| Svc Id | SDP-BIND Information | Type | Root-leaf-tag | Leaf-Ac |
|--------|----------------------|-------|---------------|---------|
| 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] [ies] [vpls] [mirror] [b-vpls] [i-vpls] [m-vpls] [ipipe] sdp <i>sdp-id</i> [<i>customer customer-id</i>] service-using etree |
| Context | show>service |
| Description | This command displays the services matching certain usage properties. If no optional parameters are specified, all services defined on the system are displayed. |
| Parameters | epipe — Displays matching Epipe services. ies — Displays matching IES instances. vpls — Displays matching VPLS instances. mirror — Displays matching mirror services. b-vpls — Displays matching B-VPLS services. |

i-vpls — Displays matching I-VPLS services.

ipipe — Displays matching Ipipe services.

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 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
```

Show, Clear, Debug Commands

```

103          Epipe    Up    Up          6          09/22/2006 23:05:58
-----
Matching Services : 4
=====
*A:ALA-12#

*A:ALA-14# show service service-using
=====
Services
=====
ServiceId    Type      Adm    Opr      CustomerId    Last Mgmt Change
-----
10           mVPLS     Down   Down     1             10/26/2006 15:44:57
11           mVPLS     Down   Down     1             10/26/2006 15:44:57
100          mVPLS     Up     Up       1             10/26/2006 15:44:57
101          mVPLS     Up     Up       1             10/26/2006 15:44:57
102          mVPLS     Up     Up       1             10/26/2006 15:44:57
-----
Matching Services : 5
-----
*A:ALA-14#

*A:SetupCLI# show service service-using
- service-using [epipe] [ies] [vpls] [mirror] [ipipe] [b-vpls] [i-vpls] [m-vpls]
[sdp <sdp-id>] [customer <customer-id>]

<epipe>      : keyword - displays epipe services
<ies>        : keyword - displays ies services
<vpls>       : keyword - displays vpls services
<mirror>     : keyword - displays mirror services
<ipipe>      : keyword - displays ipipe services
<sdp-id>     : [1..17407] - display services using this sdp
<customer-id> : [1..2147483647] - display services using this customer
<b-vpls>     : keyword - displays b-vpls services
<i-vpls>     : keyword - displays i-vpls services
<m-vpls>     : keyword - displays m-vpls services

*A:SetupCLI# show service service-using
=====
Services
=====
ServiceId    Type      Adm    Opr      CustomerId    Last Mgmt Change
-----
23           mVPLS     Up     Down     2             09/25/2007 21:45:58
100          Epipe     Up     Down     2             09/25/2007 21:45:58
101          Epipe     Up     Down     2             09/25/2007 21:45:58
102          Epipe     Up     Down     2             09/25/2007 21:45:58
105          Epipe     Up     Down     2             09/25/2007 21:45:58
110          Epipe     Up     Down     1             09/25/2007 21:45:58
990          IES       Up     Down     1             09/25/2007 21:45:58
1000         Mirror    Up     Down     1             09/25/2007 21:45:59
1001         Epipe     Up     Down     1             09/25/2007 21:45:58
1002         Epipe     Up     Down     1             09/25/2007 21:45:58
1003         Epipe     Up     Down     1             09/25/2007 21:45:58
1004         Epipe     Up     Down     1             09/25/2007 21:45:58
2000         Mirror    Up     Down     1             09/25/2007 21:45:59
2001         i-VPLS    Up     Down     1             09/25/2007 21:45:59
2002         b-VPLS    Up     Down     1             09/25/2007 21:45:59
2003         i-VPLS    Down   Down     1             09/25/2007 21:45:59

```

```

2004      b-mVPLS   Down   Down      1          09/25/2007 21:45:59
2005      i-mVPLS   Down   Down      1          09/25/2007 21:45:59
8787      IES       Up     Down      2          09/25/2007 21:45:58
8888      IES       Up     Down      1          09/25/2007 21:45:58
10000     IES       Down   Down      1          09/25/2007 21:45:59
10001     VPLS      Up     Down      1          09/25/2007 21:45:58
483000    Ipipe      Down   Down      2          09/25/2007 21:45:59
483001    Ipipe      Up     Down      2          09/25/2007 21:45:59
483004    Ipipe      Down   Down      2          09/25/2007 21:45:59
483007    VPLS      Down   Down      2          09/25/2007 21:45:59
483010    Ipipe      Down   Down      1          09/25/2007 21:45:59
...

```

```

-----
Matching Services : 27
-----

```

```

*A:SetupCLI#

```

```

*A:SetupCLI# show service service-using

```

```

=====
Services
=====
ServiceId      Type      Adm      Opr      CustomerId      Last Mgmt Change
-----
23             mVPLS     Up       Down     2               09/25/2007 21:45:58
100            Epipe     Up       Down     2               09/25/2007 21:45:58
101            Epipe     Up       Down     2               09/25/2007 21:45:58
102            Epipe     Up       Down     2               09/25/2007 21:45:58
105            Epipe     Up       Down     2               09/25/2007 21:45:58
110            Epipe     Up       Down     1               09/25/2007 21:45:58
990            IES       Up       Down     1               09/25/2007 21:45:58
1000           Mirror    Up       Down     1               09/25/2007 21:45:59
1001            Epipe     Up       Down     1               09/25/2007 21:45:58
1002            Epipe     Up       Down     1               09/25/2007 21:45:58
1003            Epipe     Up       Down     1               09/25/2007 21:45:58
1004            Epipe     Up       Down     1               09/25/2007 21:45:58
2000           Mirror    Up       Down     1               09/25/2007 21:45:59
2001            i-VPLS    Up       Down     1               09/25/2007 21:45:59
2002            b-VPLS    Up       Down     1               09/25/2007 21:45:59
2003            i-VPLS    Down    Down     1               09/25/2007 21:45:59
2004            b-mVPLS   Down    Down     1               09/25/2007 21:45:59
2005            i-mVPLS   Down    Down     1               09/25/2007 21:45:59
8787            IES       Up       Down     2               09/25/2007 21:45:58
8888            IES       Up       Down     1               09/25/2007 21:45:58
10000          IES       Down    Down     1               09/25/2007 21:45:59
10001          VPLS      Up       Down     1               09/25/2007 21:45:58
483000         Ipipe     Down    Down     2               09/25/2007 21:45:59
483001         Ipipe     Up       Down     2               09/25/2007 21:45:59
483004         Ipipe     Down    Down     2               09/25/2007 21:45:59
483007         VPLS      Down    Down     2               09/25/2007 21:45:59
483010         Ipipe     Down    Down     1               09/25/2007 21:45:59
...

```

```

-----
Matching Services : 27
-----

```

```

*A:SetupCLI#

```

```

*A:term17>config>service>epipe# show service id 2000 epipe

```

```

=====
Related Epipe services for bVpls service 2000
=====

```

Show, Clear, Debug Commands

```
EpipE SvcId      Oper ISID      Admin      Oper
-----
1              1              Down       Down
-----
Number of Entries : 1
-----
*A:term17>config>service>epipe#
```

The following sample outputs show VPLS Services configured as E-Tree.

```
*A:DutA# show service service-using
=====
Services
=====
ServiceId      Type      Adm  Opr  CustomerId Service Name
-----
1              VPLS      Up   Up   1          evpn-vxlan-1
2              VPRN      Up   Up   1
2005           VPLS-Etr* Up   Up   1
2006           VPRN      Up   Up   1
2147483648     IES       Up   Down 1          _tmnx_InternalIesService
2147483649     intVpls   Up   Down 1          _tmnx_InternalVplsService
-----
Matching Services : 6
-----
* indicates that the corresponding row element may have been truncated.

*A:DutA# show service service-using etree
=====
Services [etree]
=====
ServiceId      Type      Adm  Opr  CustomerId Service Name
-----
2005           VPLS-Etr* Up   Up   1
-----
Matching Services : 1
-----
* indicates that the corresponding row element may have been truncated.
```

subscriber-using

| | |
|--------------------|---|
| Syntax | subscriber-using [service-id <i>service-id</i>] [sap-id <i>sap-id</i>] [interface <i>ip-int-name</i>] [ip <i>ip-address[/mask]</i>] [mac <i>ieee-address</i>] [sub-profile <i>sub-profile-name</i>] [sla-profile <i>sla-profile-name</i>] |
| Context | show>service>subscriber-using |
| Description | This command displays subscribers using specified options. |
| Parameters | service-id <i>service-id</i> — Display subscriber information about the specified service ID. Values service-id: 1 — 214748364 svc-name: A string up to 64 characters in length. sap-id <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1319 for command syntax. interface <i>ip-int-name</i> — Display subscriber information about the specified interface. ip <i>ip-address[/mask]</i> — Display subscriber information about the specified IP address. mac <i>ieee-address</i> — Display subscriber information about the specified MAC address. sub-profile <i>sub-profile-name</i> — Display subscriber information about the specified subscriber profile name. sla-profile <i>sla-profile-name</i> — Display subscriber information about the specified SLA profile name. |

id

| | |
|--------------------|---|
| Syntax | id <i>service-id</i> |
| Context | show>service |
| Description | This command displays information for a particular service-id. |
| Parameters | <p><i>service-id</i> — The unique service identification number that identifies the service in the service domain.</p> <p>Values</p> <p>service-id: 1 — 214748364</p> <p>svc-name: A string up to 64 characters in length.</p> <p>all — Display detailed information about the service.</p> <p>arp — Display ARP entries for the service.</p> <p>authentication — Display subscriber authentication information</p> <p>base — Display basic service information.</p> <p>dhcp — Display DHCP information.</p> <p>endpoint — Display service endpoint information.</p> <p>epipe — Display e-pipe services associated with the b-vpls service.</p> <p>fdb — Display FDB entries.</p> <p>gsmp — Display GSMP information.</p> <p>host — Display static hosts configured on the service.</p> <p>i-vpls — Display i-vpls services associated with the b-vpls.</p> <p>igmp-snooping — Display IGMP snooping information.</p> <p>interface — Display service interfaces.</p> <p>l2-route-table — Display Layer-2 route information associated with the service.</p> <p>l2pt — Display L2PT information of SAPs and Spokes.</p> <p>labels — Display labels being used by this service.</p> <p>mac-move — Display Mac Move related information about the service.</p> <p>mac-protect — Display MAC protect information.</p> <p>mfib — Display MFIB related information.</p> <p>mld-snooping — Display MLD snooping information.</p> <p>mmrp — Display MMRP information.</p> <p>mrp — Display MRP information</p> <p>msap — Display MSAPs associated to the service.</p> <p>pim-snooping — Display PIM snooping information.</p> <p>pppoe — Display PPPoE information.</p> <p>retailers — Display service retailer information.</p> |

sap — Display SAPs associated to the service.

sdp — Display SDPs associated with the service.

source-address — Display source-address configured for applications.

split-horizon-group — Display split horizon group information.

stp — Display STP information.

subscriber-host — Display subscriber host information.

wholesalers — Display service wholesaler 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. |
| SDP Id | The SDP identifier. |
| Description | Generic information about the service. |
| Customer Id | The customer identifier. |
| Last Mgmt Change | The date and time of the most recent management-initiated change to this customer. |
| SAP Count | The number of SAPs specified for this service. |
| SDP Bind Count | The number of SDPs bound to this service. |
| Split Horizon Group | Name of the split horizon group for this service. |
| Description | Description of the split horizon group. |
| Last Changed | The date and time of the most recent management-initiated change to this split horizon group. |
| SDP Id | The SDP identifier. |
| Type | Indicates whether this service SDP binding is a spoke or a mesh. |
| Admin Path MTU | The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |
| Oper Path MTU | The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |
| Delivery | Specifies the type of delivery used by the SDP: GRE or MPLS. |
| Admin State | The administrative state of this SDP. |
| Oper State | The operational state of this SDP. |

| Label | Description (Continued) |
|------------------------|--|
| Ingress Label | The label used by the far-end device to send packets to this device in this service by this SDP. |
| Egress Label | The label used by this device to send packets to the far-end device in this service by this SDP. |
| Ingress Filter | The ID of the ingress filter policy. |
| Egress Filter | The ID of the egress filter policy. |
| Far End | Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP. |
| Last Changed | The date and time of the most recent change to this customer. |
| Hello Time | Specifies how often the SDP echo request messages are transmitted on this SDP. |
| Hello Msg Len | Specifies the length of the SDP echo request messages transmitted on this SDP. |
| Max Drop Count | Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault. |
| Hold Down Time | Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state. |
| SDP Delivery Mechanism | When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS |
| Number of SDPs | The total number SDPs applied to this service ID. |
| Service Id | The service identifier. |
| Port Id | The ID of the access port where this SAP is defined. |
| Description | Generic information about the SAP. |
| Encap Value | The value of the label used to identify this SAP on the access port. |
| Admin State | The administrative state of the SAP. |
| Oper State | The operating state of the SAP. |
| Last Changed | The date and time of the last change. |
| Admin MTU | The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |
| Oper MTU | The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |

| Label | Description (Continued) |
|-----------------------|---|
| Ingress qos-policy | The SAP ingress QoS policy ID. |
| Egress qos-policy | The SAP egress QoS policy ID. |
| Ingress Filter-Id | The SAP ingress filter policy ID. |
| Egress Filter-Id | The SAP egress filter policy ID. |
| Multi Svc Site | Indicates the multi-service site that the SAP is a member. |
| Ingress sched-policy | Indicates the ingress QoS scheduler for the SAP. |
| Egress sched-policy | Indicates the egress QoS scheduler for the SAP. |
| Acct. Pol | Indicates the accounting policy applied to the SAP. |
| Collect Stats | Specifies whether accounting statistics are collected on the SAP. |
| Dropped | The number of packets or octets dropped. |
| Offered Hi Priority | The number of high priority packets, as determined by the SAP ingress QoS policy. |
| Offered Low Priority | The number of low priority packets, as determined by the SAP ingress QoS policy. |
| Offered Low Priority | The number of low priority packets, as determined by the SAP ingress QoS policy. |
| Forwarded In Profile | The number of in-profile packets or octets (rate below CIR) forwarded. |
| Forwarded Out Profile | The number of out-of-profile packets or octets (rate above CIR) forwarded. |
| Dropped In Profile | The number of in-profile packets or octets discarded. |
| Dropped Out Profile | The number of out-of-profile packets or octets discarded. |
| Forwarded In Profile | The number of in-profile packets or octets (rate below CIR) forwarded. |
| Forwarded Out Profile | The number of out-of-profile packets or octets (rate above CIR) forwarded. |
| Ingress Queue 1 | The index of the ingress QoS queue of this SAP. |
| High priority offered | The packets or octets count of the high priority traffic for the SAP. |
| High priority dropped | The number of high priority traffic packets/octets dropped. |
| Low priority offered | The packets or octets count of the low priority traffic. |

| Label | Description (Continued) |
|-----------------------|---|
| Low priority dropped | The number of low priority traffic packets/octets dropped. |
| In profile forwarded | The number of in-profile packets or octets (rate below CIR) forwarded. |
| Out profile forwarded | The number of out-of-profile octets (rate above CIR) forwarded. |
| Egress Queue 1 | The index of the egress QoS queue of the SAP. |
| In profile forwarded | The number of in-profile packets or octets (rate below CIR) forwarded. |
| In profile dropped | The number of in-profile packets or octets dropped for the SAP. |
| Out profile forwarded | The number of out-of-profile packets or octets (rate above CIR) forwarded. |
| Out profile dropped | The number of out-of-profile packets or octets discarded. |
| State | Specifies whether DHCP Relay is enabled on this SAP. |
| Info Option | Specifies whether Option 82 processing is enabled on this SAP. |
| Action | Specifies the Option 82 processing on this SAP or interface: keep, replace or drop. |
| Circuit ID | Specifies whether the If Index is inserted in Circuit ID sub-option of Option 82. |
| Remote ID | Specifies whether the far-end MAC address is inserted in Remote ID sub-option of Option 82. |
| Managed by Service | Specifies the service-id of the management VPLS managing this SAP. |
| Managed by MSTI | Specifies the MST instance inside the management VPLS managing this SAP. |
| Last BPDU from | The bridge ID of the sender of the last BPDU received on this SAP. |
| Managed by SAP | Specifies the sap-id inside the management VPLS managing this SAP. |
| Prune state | Specifies the STP state inherited from the management VPLS. |
| Managed by Service | Specifies the service-id of the management VPLS managing this spoke SDP. |
| Last BPDU from | The bridge ID of the sender of the last BPDU received on this SAP. |
| Managed by Spoke | Specifies the sap-id inside the management VPLS managing this spoke SDP. |
| Prune state | Specifies the STP state inherited from the management VPLS. |

| Label | Description (Continued) |
|--------------|--|
| Peer Pw Bits | <p>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.</p> <p>pwNotForwarding — Pseudowire not forwarding lacIngressFault Local — Attachment circuit RX fault lacEgresssFault Local — Attachment circuit TX fault psnIngressFault Local — PSN-facing PW RX fault psnEgressFault Local — PSN-facing PW TX fault pwFwdingStandby — Pseudowire in standby mode</p> |

Sample Output

```
*A:ALA-48# show service id 700 all
=====
Service Detailed Information
=====
Service Id       : 700                Vpn Id           : 0
Service Type     : VPLS
Description      : IMA VPLS
Customer Id      : 7
Last Status Change: 02/02/2009 09:27:55
Last Mgmt Change : 02/02/2009 09:27:57
Admin State      : Up                  Oper State        : Down
MTU              : 1514                Def. Mesh VC Id   : 700
SAP Count        : 1                  SDP Bind Count    : 2
Snd Flush on Fail : Disabled           Host Conn Verify   : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
-----
BGP Auto-discovery Information
-----
Admin State       : Down                Vpls Id           : None
Route Dist        : None                Prefix            : 10.10.10.103
Rte-Target Import : None                Rte-Target Export : None
Vsi-Import        : None
Vsi-Export        : None
PW-Template Id    : None
-----
Split Horizon Group specifics
-----
Split Horizon Group : DSL-group1
-----
Description       : (Not Specified)
Instance Id       : 1                  Last Change        : 02/02/2009 09:27:57
-----
Split Horizon Group : SHG_test
-----
Description       : test
Instance Id       : 2                  Last Change        : 02/02/2009 09:27:57
```

Service Destination Points(SDPs)

Sdp Id 2:222 -(10.10.10.104)

| | | | |
|-----------------------|---|--------------------|-------------|
| Description | : GRE-10.10.10.104 | | |
| SDP Id | : 2:222 | Type | : Spoke |
| Split Horiz Grp | : (Not Specified) | | |
| VC Type | : Ether | VC Tag | : n/a |
| Admin Path MTU | : 0 | Oper Path MTU | : 0 |
| Far End | : 10.10.10.104 | Delivery | : GRE |
| Admin State | : Up | Oper State | : Down |
| Acct. Pol | : None | Collect Stats | : Disabled |
| Ingress Label | : 0 | Egress Label | : 0 |
| Ing mac Fltr | : n/a | Egr mac Fltr | : n/a |
| Ing ip Fltr | : n/a | Egr ip Fltr | : n/a |
| Ing ipv6 Fltr | : n/a | Egr ipv6 Fltr | : n/a |
| Admin ControlWord | : Not Preferred | Oper ControlWord | : False |
| Last Status Change | : 02/02/2009 09:27:55 | Signaling | : TLDP |
| Last Mgmt Change | : 02/02/2009 09:27:57 | Force Vlan-Vc | : Disabled |
| Endpoint | : N/A | Precedence | : 4 |
| Class Fwding State | : Down | | |
| Flags | : SdpOperDown NoIngVCLabel NoEgrVCLabel PathMTUTooSmall | | |
| Time to RetryReset | : never | Retries Left | : 3 |
| Mac Move | : Blockable | Blockable Level | : Tertiary |
| Peer Pw Bits | : None | | |
| Peer Fault Ip | : None | | |
| Max Nbr of MAC Addr | : No Limit | Total MAC Addr | : 0 |
| Learned MAC Addr | : 0 | Static MAC Addr | : 0 |
| MAC Learning | : Enabled | Discard Unkwn Srce | : Disabled |
| BPDU Translation | : Disabled | | |
| L2PT Termination | : Disabled | | |
| MAC Pinning | : Disabled | | |
| Ignore Standby Sig | : False | Block On Mesh Fail | : False |
| KeepAlive Information | : | | |
| Admin State | : Disabled | Oper State | : Disabled |
| Hello Time | : 10 | Hello Msg Len | : 0 |
| Max Drop Count | : 3 | Hold Down Time | : 10 |
| Statistics | : | | |
| I. Fwd. Pkts. | : 0 | I. Dro. Pkts. | : 0 |
| I. Fwd. Octs. | : 0 | I. Dro. 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 :
SDP Delivery Mechanism is not MPLS

Stp Service Destination Point specifics

| | | | |
|-------------------|------------|----------------|--------------|
| Stp Admin State | : Up | Stp Oper State | : Down |
| Core Connectivity | : Down | | |
| Port Role | : Disabled | Port State | : Discarding |

Show, Clear, Debug Commands

```

Port Number      : 2049
Port Path Cost   : 10
Admin Edge       : Disabled
Link Type        : Pt-pt
Root Guard       : Disabled
Last BPDUs from  : N/A
Designated Bridge : N/A

Port Priority     : 128
Auto Edge        : Enabled
Oper Edge        : False
BPDU Encap       : Dot1d
Active Protocol  : Rstp

Designated Port Id: 0

Fwd Transitions  : 0
Cfg BPDUs rcvd   : 0
TCN BPDUs rcvd   : 0
RST BPDUs rcvd   : 0

Bad BPDUs rcvd   : 0
Cfg BPDUs tx     : 0
TCN BPDUs tx     : 0
RST BPDUs tx     : 0
-----
Sdp Id 2:700  -(10.10.10.104)
-----
Description      : GRE-10.10.10.104
SDP Id           : 2:700
Split Horiz Grp  : (Not Specified)
VC Type          : Ether
Admin Path MTU   : 0
Far End          : 10.10.10.104

Type             : Mesh
VC Tag           : n/a
Oper Path MTU    : 0
Delivery         : GRE

Admin State       : Up
Acct. Pol        : None
Ingress Label    : 0
Ing mac Fltr     : n/a
Ing ip Fltr      : n/a
Ing ipv6 Fltr    : n/a
Admin ControlWord : Not Preferred
Last Status Change : 02/02/2009 09:27:55
Last Mgmt Change  : 02/02/2009 09:27:57
Endpoint         : N/A
Class Fwding State : Down
Flags            : SdpOperDown
                  NoIngVCLabel NoEgrVCLabel
                  PathMTUTooSmall

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  : False
Signaling         : TLDP
Force Vlan-Vc     : Disabled
Precedence        : 4

Peer Pw Bits      : None
Peer Fault Ip     : None
MAC Pinning       : Disabled

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
E. Fwd. Octets     : 0
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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-----
Number of SDPs : 2
-----
Service Access Points
-----

```

SAP 1/1/9:0

```
-----
Service Id      : 700
SAP             : 1/1/9:0                      Encap           : q-tag
Description     : (Not Specified)
Admin State     : Up                          Oper State      : Down
Flags           : PortOperDown
Multi Svc Site  : None
Last Status Change : 02/02/2009 09:27:55
Last Mgmt Change  : 02/02/2009 09:27:57
Sub Type        : regular
Dot1Q Ethertype : 0x8100                      QinQ Ethertype  : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit                  Total MAC Addr   : 0
Learned MAC Addr   : 0                        Static MAC Addr  : 0
Admin MTU          : 1518                     Oper MTU         : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id   : 10
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id  : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id : n/a
tod-suite          : None                      qinq-pbit-marking : both
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
Q Frame-Based Acct : Disabled
ARP Reply Agent    : Enabled                   Host Conn Verify  : Enabled
Mac Learning       : Enabled                   Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                   Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol         : None                       Collect Stats     : Disabled

Anti Spoofing      : Ip                       Avl Static Hosts  : 1
                                                         Tot Static Hosts  : 1

Calling-Station-Id : n/a
Application Profile: None

MCAC Policy Name   :                         MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited
Restr MacProt Src  : Enabled                   Restr MacUnpr Dst : Disabled
Time to RetryReset : never                     Retries Left      : 3
Mac Move           : Blockable                 Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy         : none
-----
```

Stp Service Access Point specifics

```
-----
Stp Admin State   : Up                          Stp Oper State    : Down
Core Connectivity : Down
Port Role         : Disabled
Port Number       : 2048
Port Path Cost    : 10
Admin Edge        : Disabled
Link Type         : Pt-pt
Root Guard        : Disabled
Last BPDU from    : N/A
CIST Desig Bridge : N/A

Forward transitions: 0

Port State        : Discarding
Port Priority      : 128
Auto Edge         : Enabled
Oper Edge         : False
BPDU Encap        : Dot1d
Active Protocol    : Rstp
Designated Port   : N/A

Bad BPDUs rcvd    : 0
-----
```

Show, Clear, Debug Commands

```
Cfg BPDUs rcvd      : 0                      Cfg BPDUs tx       : 0
TCN BPDUs rcvd      : 0                      TCN BPDUs tx       : 0
RST BPDUs rcvd      : 0                      RST BPDUs tx       : 0
MST BPDUs rcvd      : 0                      MST BPDUs tx       : 0
-----
ARP host
-----
Admin State          : outOfService
Host Limit           : 1                      Min Auth Interval  : 15 minutes
-----
QOS
-----
Ingress qos-policy   : 100                    Egress qos-policy   : 1
Shared Q plcy        : default                 Multipoint shared    : Enabled
I. Sched Pol         : SLA1
E. Sched Pol         : SLA1
-----
Ingress Queue Override
-----
Queue Id             : 1 (no overrides)
-----
Egress Queue Override
-----
Queue Id             : 1 (no overrides)
-----
DHCP
-----
Description          : (Not Specified)
Admin State          : Down                    Lease Populate       : 0
DHCP Snooping        : Down                    Action               : Keep
Proxy Admin State    : Down
Proxy Lease Time     : N/A
Emul. Server Addr    : Not Configured
-----
Subscriber Management
-----
Admin State          : Down                    MAC DA Hashing      : False
Def Sub-Id           : None
Def Sub-Profile      : None
Def SLA-Profile      : None
Def App-Profile      : None
Sub-Ident-Policy     : None

Subscriber Limit     : 1
Single-Sub-Parameters
  Prof Traffic Only  : False
  Non-Sub-Traffic    : N/A
-----
Sap Statistics
-----
Last Cleared Time    : N/A

                                Packets          Octets
Forwarding Engine Stats
Dropped              : 0                      0
Off. HiPrio          : 0                      0
Off. LowPrio         : 0                      0
Off. Uncolor         : 0                      0

Queueing Stats(Ingress QoS Policy 100)
```


| | | |
|--------------|-----|---|
| Dro. HiPrio | : 0 | 0 |
| Dro. LowPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |

Queueing Stats(Egress QoS Policy 1)

| | | |
|--------------|-----|---|
| Dro. InProf | : 0 | 0 |
| Dro. OutProf | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |

Sap per Queue stats

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

Ingress Queue 10 (Unicast) (Priority)

| | | |
|--------------|-----|---|
| Off. HiPrio | : 0 | 0 |
| Off. LoPrio | : 0 | 0 |
| Dro. HiPrio | : 0 | 0 |
| Dro. LoPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |

Ingress Queue 12 (Unicast) (Priority)

| | | |
|--------------|-----|---|
| Off. HiPrio | : 0 | 0 |
| Off. LoPrio | : 0 | 0 |
| Dro. HiPrio | : 0 | 0 |
| Dro. LoPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |

Ingress Queue 13 (Unicast) (Priority)

| | | |
|--------------|-----|---|
| Off. HiPrio | : 0 | 0 |
| Off. LoPrio | : 0 | 0 |
| Dro. HiPrio | : 0 | 0 |
| Dro. LoPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |

...

VPLS Spanning Tree Information

| | | | |
|-----------------|--------|-------------------|--------|
| VPLS oper state | : Down | Core Connectivity | : Down |
| Stp Admin State | : Up | Stp Oper State | : Down |
| Mode | : Rstp | Vcp Active Prot. | : N/A |

| | | | |
|---------------------|---------------------------|--------------------|------|
| Bridge Id | : 10:02.90:30:ff:00:00:00 | Bridge Instance Id | : 2 |
| Bridge Priority | : 4096 | Tx Hold Count | : 5 |
| Topology Change | : Inactive | Bridge Hello Time | : 5 |
| Last Top. Change | : 0d 00:00:00 | Bridge Max Age | : 25 |
| Top. Change Count | : 0 | Bridge Fwd Delay | : 20 |
| MST region revision | : 0 | Bridge max hops | : 20 |
| MST region name | : | | |

Show, Clear, Debug Commands

Root Bridge : N/A
Primary Bridge : N/A

Root Path Cost : 0
Rcvd Hello Time : 5
Root Priority : 4098

Root Forward Delay: 20
Root Max Age : 25
Root Port : N/A

Forwarding Database specifics

Service Id : 700
Primary Factor : 3
Mac Move Rate : 2
Mac Move Retries : 3
Table Size : 250
Learned Count : 0
OAM-learned Count : 0
Host-learned Count: 1
Remote Age : 900
High Watermark : 95%
Mac Learning : Enabled
Mac Aging : Enabled
Mac Subnet Len : 48

Mac Move : Disabled
Secondary Factor : 2
Mac Move Timeout : 10
Total Count : 1
Static Count : 0
DHCP-learned Count: 0
Local Age : 300
Low Watermark : 90%
Discard Unknown : Disabled
Relearn Only : False

IGMP Snooping Base info

Admin State : Up
Querier : No querier found

| Sap/Sdp Id | Oper State | MRtr Port | Pim Port | Send Queries | Max Num Grps | Max Num Srcs | MVR From-VPLS | Num Grps |
|-------------|------------|-----------|----------|--------------|--------------|--------------|---------------|----------|
| sap:1/1/9:0 | Down | No | No | No | None | None | Local | 0 |
| sdp:2:222 | Down | No | No | No | None | None | N/A | 0 |
| sdp:2:700 | Down | No | No | No | None | None | N/A | 0 |

MLD Snooping Base info

Admin State : Down
Querier : No querier found

| Sap/Sdp Id | Oper State | MRtr Port | Send Queries | Max Num Groups | MVR From-VPLS | Num Groups |
|-------------|------------|-----------|--------------|----------------|---------------|------------|
| sap:1/1/9:0 | Down | No | Disabled | No Limit | Local | 0 |
| sdp:2:222 | Down | No | Disabled | No Limit | N/A | 0 |
| sdp:2:700 | Down | No | Disabled | No Limit | N/A | 0 |

DHCP Summary, service 700

| Sap/Sdp | Snoop | Used/ Provided | Arp Reply Agent | Info Option | Admin State |
|-------------|-------|----------------|-----------------|-------------|-------------|
| sap:1/1/9:0 | No | 0/0 | Yes | Keep | Down |
| sdp:2:222 | No | N/A | N/A | N/A | N/A |
| sdp:2:700 | No | N/A | N/A | N/A | N/A |

Number of Entries : 3

ARP host Summary, service 700

```

Sap                Used        Provided    Admin State
-----
sap:1/1/9:0        0          1          outOfService
-----

Number of SAPs : 1
-----

Service Endpoints
-----

No Endpoints found.
=====
*A:ALA-48#

*A:SetupCLI# show service id 2001 all
=====
Service Detailed Information
=====
Service Id         : 2001                Vpn Id             : 0
Service Type       : i-VPLS
Customer Id        : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change   : 09/25/2007 21:45:59
Admin State        : Up                  Oper State          : Down
MTU                : 1514                Def. Mesh VC Id    : 2001
SAP Count          : 1                  SDP Bind Count     : 0
Snd Flush on Fail  : Disabled            Host Conn Verify    : Disabled
b-vpls Id          : 2002                Oper ISID           : 122
Snd Flush in bVpls: Disabled
Snd Flush in bVpls: All-from-me          b-vpls-status       : Up
                  : All-but-mine

-----
Split Horizon Group specifics
-----
Service Destination Points(SDPs)
-----
No Matching Entries
-----
Service Access Points
-----
SAP 1/1/12:2001.2001
-----
Service Id         : 2001
SAP                : 1/1/12:2001.2001    Encap               : qinq
Sub Type           : regular
QinQ Dot1p         : Default
Dot1Q Ethertype    : 0x8100              QinQ Ethertype      : 0x8100

Admin State        : Up                  Oper State          : Down
Flags              : PortOperDown
Last Status Change : 09/25/2007 21:12:01
Last Mgmt Change   : 09/25/2007 21:45:59
Max Nbr of MAC Addr: No Limit            Total MAC Addr      : 0
Learned MAC Addr   : 0                  Static MAC Addr     : 0
Admin MTU          : 1522                Oper MTU            : 1522
Ingress qos-policy : 1                  Egress qos-policy   : 1
Shared Q plcy      : n/a                Multipoint shared    : Disabled
Ingr IP Fltr-Id    : n/a                Egr IP Fltr-Id      : n/a
Ingr Mac Fltr-Id   : n/a                Egr Mac Fltr-Id     : n/a
tod-suite          : None                qinq-pbit-marking   : both
Egr Agg Rate Limit : max

```

Show, Clear, Debug Commands

```

Q Frame-Based Acct : Disabled
Mac Learning       : Enabled
Mac Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

```

```

Discard Unkwn Srce: Disabled
Mac Pinning        : Disabled

```

```

Multi Svc Site     : None
Acct. Pol          : None
Restr MacProt Src  : Disabled
Mac Move           : Non Blockable
Egr MCast Grp      :

```

```

Collect Stats      : Disabled
Restr MacUnpr Dst  : Disabled
Mac Move Block Lvl: Tertiary

```

----- Stp Service Access Point specifics -----

```

Stp Admin State    : Up
Core Connectivity   : Down
Port Role          : N/A
Port Number        : 2049
Port Path Cost     : 10
Admin Edge         : Disabled
Link Type          : Pt-pt
Root Guard         : Disabled
Last BPDU from     : N/A
CIST Desig Bridge  : N/A

```

```

Stp Oper State     : Down
Port State         : Unknown
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
RST BPDUs rcvd     : 0
MST BPDUs rcvd     : 0

```

```

Bad BPDUs rcvd     : 0
Cfg BPDUs tx       : 0
TCN BPDUs tx       : 0
RST BPDUs tx       : 0
MST BPDUs tx       : 0

```

----- SAP MRP Information -----

```

Rx Pdus           : 0
Dropped Pdus      : 0
Rx New Event      : 0
Rx In Event       : 0
Rx Empty Event    : 0
Tx New Event      : 0
Tx In Event       : 0
Tx Empty Event    : 0

```

```

Tx Pdus           : 0
Tx Pdus           : 0
Rx Join-In Event  : 0
Rx Join Empty Evt : 0
Rx Leave Event    : 0
Tx Join-In Event  : 0
Tx Join Empty Evt : 0
Tx Leave Event    : 0

```

----- SAP MMRP Information -----

```

MAC Address      Registered      Declared

```

```

Number of MACs=0 Registered=0 Declared=0

```

----- Sap Statistics -----

```

Last Cleared Time : N/A
                  Packets
Forwarding Engine Stats
Dropped           : 0
Off. HiPrio       : 0
Off. LowPrio      : 0
Off. Uncolor      : 0

```

```

Octets

```

```

0
0
0
0

```

```

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio       : 0

```

```

0

```

| | | |
|--------------|-----|---|
| Dro. LowPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |

Queueing Stats(Egress QoS Policy 1)

| | | |
|--------------|-----|---|
| Dro. InProf | : 0 | 0 |
| Dro. OutProf | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |

Sap per Queue stats

| | Packets | Octets |
|--|---------|--------|
| Ingress Queue 1 (Unicast) (Priority) | | |
| Off. HiPrio | : 0 | 0 |
| Off. LoPrio | : 0 | 0 |
| Dro. HiPrio | : 0 | 0 |
| Dro. LoPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |
| Ingress Queue 11 (Multipoint) (Priority) | | |
| Off. HiPrio | : 0 | 0 |
| Off. LoPrio | : 0 | 0 |
| Dro. HiPrio | : 0 | 0 |
| Dro. LoPrio | : 0 | 0 |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |
| Egress Queue 1 | | |
| For. InProf | : 0 | 0 |
| For. OutProf | : 0 | 0 |
| Dro. InProf | : 0 | 0 |
| Dro. OutProf | : 0 | 0 |

VPLS Spanning Tree Information

| | | | |
|---------------------|---------------------------|--------------------|--------|
| VPLS oper state | : Down | Core Connectivity | : Down |
| Stp Admin State | : Down | Stp Oper State | : Down |
| Mode | : Rstp | Vcp Active Prot. | : N/A |
| Bridge Id | : 80:00:70:ec:ff:00:00:00 | Bridge Instance Id | : 0 |
| Bridge Priority | : 32768 | Tx Hold Count | : 6 |
| Topology Change | : Inactive | Bridge Hello Time | : 2 |
| Last Top. Change | : 0d 00:00:00 | Bridge Max Age | : 20 |
| Top. Change Count | : 0 | Bridge Fwd Delay | : 15 |
| MST region revision | : 0 | Bridge max hops | : 20 |
| MST region name | : | | |
| Root Bridge | : N/A | | |
| Primary Bridge | : N/A | | |
| Root Path Cost | : 0 | Root Forward Delay | : 15 |
| Rcvd Hello Time | : 2 | Root Max Age | : 20 |
| Root Priority | : 32768 | Root Port | : N/A |

Forwarding Database specifics

| | | | |
|--------------------|--------|--------------------|------------|
| Service Id | : 2001 | Mac Move | : Disabled |
| Primary Factor | : 3 | Secondary Factor | : 2 |
| Mac Move Rate | : 2 | Mac Move Timeout | : 10 |
| Table Size | : 250 | Total Count | : 0 |
| Learned Count | : 0 | Static Count | : 0 |
| OAM-learned Count | : 0 | DHCP-learned Count | : 0 |
| Host-learned Count | : 0 | | |
| Remote Age | : 900 | Local Age | : 300 |

Show, Clear, Debug Commands

```

High WaterMark      : 95%                Low Watermark      : 90%
Mac Learning        : Enabl               Discard Unknown   : Dsabl
Mac Aging           : Dsabl               Relearn Only      : False
-----
IGMP Snooping Base info
-----
Admin State : Down
Querier      : No querier found
-----
Sap/Sdp      Oper  MRtr Pim  Send      Max Num  MVR      Num
Id           State Port Port  Queries  Groups   From-VPLS Groups
-----
sap:1/1/12:2001.2001  Down   No   No   Disabled No Limit Local    0
-----
DHCP Summary, service 2001
-----
Sap/Sdp      Snoop  Used/  Arp Reply  Info  Admin
              State Provided Agent      Option State
-----
sap:1/1/12:2001.2001  No    0/0    No        Keep   Down
-----
Number of Entries : 1
-----
MRP Information
-----
Admin State      : Down                Failed Register Cnt: 0
Max Attributes   : 2048                Attribute Count    : 0
-----
*A:SetupCLI#

*A:SetupCLI# show service id 2002 all
=====
Service Detailed Information
=====
Service Id      : 2002                Vpn Id          : 0
Service Type    : b-VPLS
Customer Id     : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change : 09/25/2007 21:45:59
Admin State     : Up                  Oper State      : Down
MTU             : 1530                Def. Mesh VC Id : 2002
SAP Count       : 2                  SDP Bind Count  : 2
Snd Flush on Fail : Disabled          Host Conn Verify : Disabled
Oper Backbone Src : 00:f7:f7:f7:f7:f7
-----
Related iVpls services for bVpls service 2002
-----
iVpls SvcId     Oper ISID                Admin              Oper
-----
2001            122                      Up                 Down
-----
Number of Entries : 1
-----
Split Horizon Group specifics
-----
Service Destination Points(SDPs)
-----
Sdp Id 2000:2001  -(101.101.101.101)
-----
SDP Id          : 2000:2001                Type              : Spoke

```

| | | | |
|----------------|-------------------|---------------|--------|
| VC Type | : Ether | VC Tag | : n/a |
| Admin Path MTU | : 1500 | Oper Path MTU | : 1500 |
| Far End | : 101.101.101.101 | Delivery | : MPLS |

| | | | |
|---------------------|--|------------------|------------|
| Admin State | : Down | Oper State | : Down |
| Acct. Pol | : None | Collect Stats | : Disabled |
| Ingress Label | : 0 | Egress Label | : 0 |
| Ing mac Fltr | : n/a | Egr mac Fltr | : n/a |
| Admin ControlWord | : Not Preferred | Oper ControlWord | : False |
| Last Status Change | : 09/25/2007 21:12:01 | Signaling | : TLDP |
| Last Mgmt Change | : 09/25/2007 21:45:59 | Force Vlan-Vc | : Disabled |
| Endpoint | : N/A | Precedence | : 4 |
| Class Fwding State | : Down | | |
| Flags | SdpOperDown SdpBindAdminDown NoIngVCLabel NoEgrVCLabel PathMTUTooSmall | | |
| Peer Pw Bits | : None | | |
| Peer Fault Ip | : None | | |
| Max Nbr of MAC Addr | : No Limit | Total MAC Addr | : 0 |
| Learned MAC Addr | : 0 | Static MAC Addr | : 0 |

| | | | |
|--------------------|------------|--------------------|------------|
| MAC Learning | : Enabled | Discard Unkwn Srce | : Disabled |
| BPDU Translation | : Disabled | | |
| L2PT Termination | : Disabled | | |
| Ignore Standby Sig | : False | | |

| | | | |
|-------------------------|-----------|----------------|---------------|
| KeepAlive Information : | | | |
| Admin State | : Enabled | Oper State | : No response |
| Hello Time | : 600 | Hello Msg Len | : 1500 |
| Max Drop Count | : 3 | Hold Down Time | : 10 |

| | | | |
|---------------|-----|----------------|-----|
| Statistics | : | | |
| I. Fwd. Pkts. | : 0 | I. Dro. Pkts. | : 0 |
| E. Fwd. Pkts. | : 0 | E. Fwd. Octets | : 0 |

Associated LSP LIST :

No LSPs Associate

Class-based forwarding :

| | | | |
|------------------|------------|---------------|--------|
| Class forwarding | : disabled | | |
| Default LSP | : Uknwn | Multicast LSP | : None |

=====

FC Mapping Table

=====

| | |
|---------|----------|
| FC Name | LSP Name |
|---------|----------|

No FC Mappings

Stp Service Destination Point specifics

| | | | |
|-------------------|-------------|--------------------|--------------|
| Mac Move | : Blockable | Blockable Level | : Tertiary |
| Stp Admin State | : Up | Stp Oper State | : Down |
| Core Connectivity | : Down | | |
| Port Role | : N/A | Port State | : Discarding |
| Port Number | : 2050 | Port Priority | : 128 |
| Port Path Cost | : 10 | Auto Edge | : Enabled |
| Admin Edge | : Disabled | Oper Edge | : N/A |
| Link Type | : Pt-pt | BPDU Encap | : Dot1d |
| Root Guard | : Disabled | Active Protocol | : N/A |
| Last BPDU from | : N/A | | |
| Designated Bridge | : N/A | Designated Port Id | : 0 |

Show, Clear, Debug Commands

```

Fwd Transitions      : 0
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0
Bad BPDUs rcvd       : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 0
RST BPDUs tx         : 0
-----
Sdp Id 2000:2001 MRP Information
-----
Rx Pdus              : 0
Dropped Pdus         : 0
Rx New Event         : 0
Rx In Event          : 0
Rx Empty Event       : 0
Tx New Event         : 0
Tx In Event          : 0
Tx Empty Event       : 0
Tx Pdus              : 0
Rx Join-In Event     : 0
Rx Join Empty Evt    : 0
Rx Leave Event       : 0
Tx Join-In Event     : 0
Tx Join Empty Evt    : 0
Tx Leave Event       : 0
-----
SDP MMRP Information
-----
MAC Address          Registered      Declared
-----
Number of MACs=0 Registered=0 Declared=0
-----
Sdp Id 2000:2002 -(101.101.101.101)
-----
SDP Id              : 2000:2002
VC Type             : Ether
Admin Path MTU      : 1500
Far End             : 101.101.101.101
Type                : Mesh
VC Tag              : n/a
Oper Path MTU       : 1500
Delivery            : MPLS

Admin State         : Down
Acct. Pol           : None
Ingress Label       : 2050
Ing mac Fltr        : n/a
Admin ControlWord   : Not Preferred
Last Status Change  : 09/25/2007 21:12:01
Last Mgmt Change    : 09/25/2007 21:45:58
Endpoint            : N/A
Class Fwding State  : Down
Flags               : SdpOperDown SdpBindAdminDown
                    PathMTUTooSmall
Peer Pw Bits        : None
Peer Fault Ip       : None
Ignore Standby Sig  : False

Oper State          : Down
Collect Stats       : Disabled
Egress Label        : 2050
Egr mac Fltr        : n/a
Oper ControlWord    : False
Signaling           : TLDP
Force Vlan-Vc       : Disabled
Precedence          : 4

KeepAlive Information :
Admin State          : Enabled
Hello Time           : 600
Max Drop Count       : 3
Oper State           : No response
Hello Msg Len        : 1500
Hold Down Time       : 10

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

Associated LSP LIST :
No LSPs Associated
Class-based forwarding :
-----
Class forwarding     : disabled
Default LSP          : Uknwn
Multicast LSP        : None
=====

```


FC Mapping Table

| FC Name | | LSP Name | |
|--|------------------------------|--------------------|------------|
| ----- | | | |
| No FC Mappings | | | |
| ----- | | | |
| Sdp Id 2000:2002 MRP Information | | | |
| ----- | | | |
| Rx Pdus | : 0 | Tx Pdus | : 0 |
| Dropped Pdus | : 0 | | |
| Rx New Event | : 0 | Rx Join-In Event | : 0 |
| Rx In Event | : 0 | Rx Join Empty Evt | : 0 |
| Rx Empty Event | : 0 | Rx Leave Event | : 0 |
| Tx New Event | : 0 | Tx Join-In Event | : 0 |
| Tx In Event | : 0 | Tx Join Empty Evt | : 0 |
| Tx Empty Event | : 0 | Tx Leave Event | : 0 |
| ----- | | | |
| SDP MMRP Information | | | |
| ----- | | | |
| MAC Address | Registered | Declared | |
| ----- | | | |
| Number of MACs=0 Registered=0 Declared=0 | | | |
| ----- | | | |
| Number of SDPs : 2 | | | |
| ----- | | | |
| Service Access Points | | | |
| ----- | | | |
| SAP 1/1/12:2002.2002 | | | |
| ----- | | | |
| Service Id | : 2002 | | |
| SAP | : 1/1/12:2002.2002 | Encap | : qinq |
| Sub Type | : regular | | |
| QinQ Dot1p | : Default | | |
| Dot1Q Ethertype | : 0x8100 | QinQ Ethertype | : 0x8100 |
| PBB Ethertype | : 0x88e7 | | |
| | | | |
| Admin State | : Down | Oper State | : Down |
| Flags | : SapAdminDown | | |
| | PortOperDown PortMTUTooSmall | | |
| Last Status Change | : 09/25/2007 21:12:01 | | |
| Last Mgmt Change | : 09/25/2007 21:45:58 | | |
| Max Nbr of MAC Addr | : No Limit | Total MAC Addr | : 0 |
| Learned MAC Addr | : 0 | Static MAC Addr | : 0 |
| Admin MTU | : 1522 | Oper MTU | : 1522 |
| Ingress qos-policy | : 1 | Egress qos-policy | : 1 |
| Shared Q plcy | : n/a | Multipoint shared | : Disabled |
| Ingr Mac Fltr-Id | : n/a | Egr Mac Fltr-Id | : n/a |
| tod-suite | : None | qinq-pbit-marking | : both |
| Egr Agg Rate Limit | : max | | |
| Q Frame-Based Acct | : Disabled | | |
| Mac Learning | : Enabled | Discard Unkwn Srce | : Disabled |
| Mac Aging | : Enabled | Mac Pinning | : Disabled |
| BPDU Translation | : Disabled | | |
| L2PT Termination | : Disabled | | |
| Vlan-translation | : None | | |
| | | | |
| Multi Svc Site | : None | | |
| Acct. Pol | : None | Collect Stats | : Disabled |
| Restr MacProt Src | : Disabled | Restr MacUnpr Dst | : Disabled |
| Mac Move | : Blockable | Mac Move Block Lvl | : Tertiary |
| Egr MCast Grp | : | | |

Show, Clear, Debug Commands

```

-----
Stp Service Access Point specifics
-----
Stp Admin State      : Up                      Stp Oper State      : Down
Core Connectivity    : Down
Port Role            : N/A                      Port State          : Unknown
Port Number          : 2049                     Port Priority        : 128
Port Path Cost       : 10                       Auto Edge           : Enabled
Admin Edge           : Disabled                  Oper Edge           : N/A
Link Type            : Pt-pt                     BPDU Encap          : Dot1d
Root Guard           : Disabled                  Active Protocol      : N/A
Last BPDU from       : N/A                      Designated Port      : N/A
CIST Desig Bridge    : N/A

Forward transitions: 0                      Bad BPDUs rcvd      : 0
Cfg BPDUs rcvd       : 0                   Cfg BPDUs tx        : 0
TCN BPDUs rcvd       : 0                   TCN BPDUs tx        : 0
RST BPDUs rcvd       : 0                   RST BPDUs tx        : 0
MST BPDUs rcvd       : 0                   MST BPDUs tx        : 0
-----

SAP MRP Information
-----
Rx Pdus              : 0                      Tx Pdus              : 0
Dropped Pdus         : 0                      Tx Pdus              : 0
Rx New Event         : 0                      Rx Join-In Event     : 0
Rx In Event          : 0                      Rx Join Empty Evt    : 0
Rx Empty Event       : 0                      Rx Leave Event       : 0
Tx New Event         : 0                      Tx Join-In Event     : 0
Tx In Event          : 0                      Tx Join Empty Evt    : 0
Tx Empty Event       : 0                      Tx Leave Event       : 0
-----

SAP MMRP Information
-----
MAC Address          Registered      Declared
-----
Number of MACs=0 Registered=0 Declared=0
-----

Sap Statistics
-----
Last Cleared Time    : N/A

                                Packets          Octets
Forwarding Engine Stats
Dropped              : 0                      0
Off. HiPrio          : 0                      0
Off. LowPrio         : 0                      0
Off. Uncolor         : 0                      0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio          : 0                      0
Dro. LowPrio         : 0                      0
For. InProf          : 0                      0
For. OutProf         : 0                      0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf          : 0                      0
Dro. OutProf         : 0                      0
For. InProf          : 0                      0
For. OutProf         : 0                      0
-----

Sap per Queue stats

```

```

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

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio                    : 0                                0
Off. LoPrio                    : 0                                0
Dro. HiPrio                    : 0                                0
Dro. LoPrio                    : 0                                0
For. InProf                    : 0                                0
For. OutProf                   : 0                                0

Egress Queue 1
For. InProf                    : 0                                0
For. OutProf                   : 0                                0
Dro. InProf                    : 0                                0
Dro. OutProf                   : 0                                0
-----
SAP 1/1/30:2002
-----
Service Id      : 2002
SAP             : 1/1/30:2002          Encap             : q-tag
Sub Type       : regular
Dot1Q Ethertype : 0x8100              QinQ Ethertype    : 0x8100
PBB Ethertype  : 0x88e7

Admin State     : Down                Oper State        : Down
Flags           : SapAdminDown
                  PortOperDown PortMTUTooSmall
Last Status Change : 09/25/2007 21:12:01
Last Mgmt Change  : 09/25/2007 21:45:58
Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
Admin MTU        : 1518
Ingress qos-policy : 1
Shared Q plcy    : n/a
Ingr Mac Fltr-Id : n/a
tod-suite        : None
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled
Mac Learning     : Enabled
Mac Aging        : Enabled
BPDU Translation : Disabled
L2PT Termination : Disabled
Vlan-translation : None

Total MAC Addr   : 0
Static MAC Addr  : 0
Oper MTU         : 1518
Egress qos-policy : 1
Multipoint shared : Disabled
Egr Mac Fltr-Id  : n/a
qing-pbit-marking : both

Discard Unkwn Srce: Disabled
Mac Pinning       : Disabled

Multi Svc Site   : None
Acct. Pol        : None
Restr MacProt Src : Disabled
Mac Move         : Blockable
Egr MCast Grp    :

Collect Stats     : Disabled
Restr MacUnpr Dst : Disabled
Mac Move Block Lvl: Tertiary
-----
Stp Service Access Point specifics

```

Show, Clear, Debug Commands

```

-----
Stp Admin State      : Up
Core Connectivity    : Down
Port Role            : N/A
Port Number          : 2048
Port Path Cost       : 10
Admin Edge           : Disabled
Link Type            : Pt-pt
Root Guard           : Disabled
Last BPDU from       : N/A
CIST Desig Bridge    : N/A

Stp Oper State       : Down
Port State           : Unknown
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
RST BPDUs rcvd       : 0
MST BPDUs rcvd       : 0

Bad BPDUs rcvd       : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 0
RST BPDUs tx         : 0
MST BPDUs tx         : 0
-----

SAP MRP Information
-----
Rx Pdus              : 0
Dropped Pdus         : 0
Rx New Event         : 0
Rx In Event          : 0
Rx Empty Event       : 0
Tx New Event         : 0
Tx In Event          : 0
Tx Empty Event       : 0

Tx Pdus              : 0
Tx Pdus              : 0
Rx Join-In Event     : 0
Rx Join Empty Evt    : 0
Rx Leave Event       : 0
Tx Join-In Event     : 0
Tx Join Empty Evt    : 0
Tx Leave Event       : 0
-----

SAP MMRP Information
-----
MAC Address          Registered      Declared
-----
Number of MACs=0 Registered=0 Declared=0
-----

Sap Statistics
-----
Last Cleared Time    : N/A

Packets              Octets
Forwarding Engine Stats
Dropped              : 0              0
Off. HiPrio          : 0              0
Off. LowPrio         : 0              0
Off. Uncolor         : 0              0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio          : 0              0
Dro. LowPrio         : 0              0
For. InProf          : 0              0
For. OutProf         : 0              0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf          : 0              0
Dro. OutProf         : 0              0
For. InProf          : 0              0
For. OutProf         : 0              0
-----

Sap per Queue stats
-----
Packets              Octets
Ingress Queue 1 (Unicast) (Priority)

```

```

Off. HiPrio          : 0          0
Off. LoPrio          : 0          0
Dro. HiPrio          : 0          0
Dro. LoPrio          : 0          0
For. InProf          : 0          0
For. OutProf         : 0          0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio          : 0          0
Off. LoPrio          : 0          0
Dro. HiPrio          : 0          0
Dro. LoPrio          : 0          0
For. InProf          : 0          0
For. OutProf         : 0          0

Egress Queue 1
For. InProf          : 0          0
For. OutProf         : 0          0
Dro. InProf          : 0          0
Dro. OutProf         : 0          0
-----
VPLS Spanning Tree Information
-----
VPLS oper state      : Down          Core Connectivity : Down
Stp Admin State      : Down          Stp Oper State   : Down
Mode                 : Rstp          Vcp Active Prot. : N/A

Bridge Id            : 80:00.70:ec:ff:00:00:00 Bridge Instance Id: 0
Bridge Priority       : 32768          Tx Hold Count    : 6
Topology Change      : Inactive        Bridge Hello Time : 2
Last Top. Change     : 0d 00:00:00     Bridge Max Age    : 20
Top. Change Count    : 0               Bridge Fwd Delay   : 15
MST region revision  : 0               Bridge max hops    : 20
MST region name      :

Root Bridge          : N/A
Primary Bridge       : N/A

Root Path Cost       : 0               Root Forward Delay: 15
Rcvd Hello Time      : 2               Root Max Age       : 20
Root Priority         : 32768           Root Port          : N/A
-----
Forwarding Database specifics
-----
Service Id           : 2002            Mac Move           : Disabled
Primary Factor        : 3               Secondary Factor    : 2
Mac Move Rate         : 2               Mac Move Timeout    : 10
Table Size            : 250             Total Count         : 0
Learned Count         : 0               Static Count        : 0
OAM-learned Count    : 0               DHCP-learned Count : 0
Host-learned Count    : 0
Remote Age            : 900             Local Age           : 300
High WaterMark        : 95%             Low Watermark       : 90%
Mac Learning          : Enabl            Discard Unknown     : Dsabl
Mac Aging             : Dsabl            Relearn Only        : False
-----
IGMP Snooping Base info
-----
Admin State          : Down
Querier              : No querier found
-----

```

Show, Clear, Debug Commands

| Sap/Sdp Id | Oper State | MRtr Port | Pim Port | Send Queries | Max Num Groups | MVR From-VPLS | Num Groups |
|----------------------|------------|-----------|----------|--------------|----------------|---------------|------------|
| sap:1/1/12:2002.2002 | Down | No | No | Disabled | No Limit | Local | 0 |
| sap:1/1/30:2002 | Down | No | No | Disabled | No Limit | Local | 0 |
| sdp:2000:2001 | Down | No | No | Disabled | No Limit | N/A | 0 |
| sdp:2000:2002 | Down | No | No | Disabled | No Limit | N/A | 0 |

DHCP Summary, service 2002

| Sap/Sdp | Snoop | Used/ Provided | Arp Reply Agent | Info Option | Admin State |
|----------------------|-------|----------------|-----------------|-------------|-------------|
| sap:1/1/12:2002.2002 | No | 0/0 | No | Keep | Down |
| sap:1/1/30:2002 | No | 0/0 | No | Keep | Down |
| sdp:2000:2001 | No | N/A | N/A | N/A | N/A |
| sdp:2000:2002 | No | N/A | N/A | N/A | N/A |

Number of Entries : 4

MRP Information

Admin State : Up
Max Attributes : 2048
Failed Register Cnt: 0
Attribute Count : 2

*A:SetupCLI#

*A:alcag1-R6# show service id 5000 all | match post-lines 15 eth-cfm
eth-cfm Configuration Information

```

Md-index          : 1                      Direction        : Up
Ma-index          : 1                      Admin            : Enabled
MepId             : 51                    CCM-Enable      : Enabled
LowestDefectPri   : allDef                HighestDefect    : none
Defect Flags      : None
Mac Address       : 00:ae:ae:ae:ae:ae
CcmTx             : 11548                 CcmSequenceErr  : 2
LbRxReply         : 1                     LbRxBadOrder    : 0
LbRxBadMsdu       : 0                     LbTxReply       : 2
LbNextSequence    : 3                     LtNextSequence  : 3
LtRxUnexplained   : 0

```

*A:alcmtul-R6#

*A:alcmtul-R6# show service id 5000 all | match post-lines 15 eth-cfm
eth-cfm Configuration Information

```

Md-index          : 1                      Direction        : Up
Ma-index          : 1                      Admin            : Enabled
MepId             : 56                    CCM-Enable      : Enabled
LowestDefectPri   : allDef                HighestDefect    : defMACstatus
Defect Flags      : bDefMACstatus
Mac Address       : 00:af:af:af:af:af
CcmTx             : 815                   CcmSequenceErr  : 0
LbRxReply         : 0                     LbRxBadOrder    : 0
LbRxBadMsdu       : 0                     LbTxReply       : 0
LbNextSequence    : 3                     LtNextSequence  : 1
LtRxUnexplained   : 0

```

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. |

authentication

| | |
|--------------------|--|
| Syntax | authentication |
| Context | show>service>id |
| Description | This command enables the context to show session authentication information. |

statistics

| | |
|--------------------|---|
| Syntax | statistics [policy name] [sap sap-id] |
| Context | show>service>id>auth |
| Description | This command displays subscriber authentication statistics. |

arp-host

| | |
|--------------------|--|
| Syntax | arp-host [wholesaler service-id] [sap sap-id interface interface-name ip-address ip-address[/mask] mac ieee-address {[port port-id] [no-inter-dest-id inter-dest-id inter-dest-id]]} [detail] arp-host statistics [sap sap-id interface interface-name] arp-host summary [interface interface-name] |
| Context | show>service>id |
| Description | This command displays ARP host related information. |

Sample Output

```
*A:Dut-C# show service id 2 arp-host
=====
ARP host table, service 2
=====
IP Address      Mac Address      Sap Id           Remaining      MC
                  Time                               Stdby
-----
128.128.1.2      00:80:00:00:00:01 2/1/5:2          00h04m41s
128.128.1.3      00:80:00:00:00:02 2/1/5:2          00h04m42s
128.128.1.4      00:80:00:00:00:03 2/1/5:2          00h04m43s
128.128.1.5      00:80:00:00:00:04 2/1/5:2          00h04m44s
128.128.1.6      00:80:00:00:00:05 2/1/5:2          00h04m45s
128.128.1.7      00:80:00:00:00:06 2/1/5:2          00h04m46s
128.128.1.8      00:80:00:00:00:07 2/1/5:2          00h04m47s
128.128.1.9      00:80:00:00:00:08 2/1/5:2          00h04m48s
128.128.1.10     00:80:00:00:00:09 2/1/5:2          00h04m49s
128.128.1.11     00:80:00:00:00:0a 2/1/5:2          00h04m50s
-----
Number of ARP hosts : 10
=====
*A:Dut-C#
```



```

*A:Dut-C# show service id 2 arp-host ip-address 128.128.1.2 detail
=====
ARP hosts for service 2
=====
Service ID           : 2
IP Address           : 128.128.1.2
MAC Address          : 00:80:00:00:00:01
SAP                  : 2/1/5:2
Remaining Time       : 00h04m58s

Sub-Ident            : "alu_1_2"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
RADIUS-User-Name     : "128.128.1.2"

Session Timeout (s)  : 301
Start Time           : 02/09/2009 16:35:07
Last Auth            : 02/09/2009 16:36:34
Last Refresh         : 02/09/2009 16:36:38
Persistence Key      : N/A
-----
Number of ARP hosts : 1
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host statistics
=====
ARP host statistics
=====
Num Active Hosts      : 20
Received Triggers     : 70
Ignored Triggers      : 10
Ignored Triggers (overload) : 0
SHCV Checks Forced    : 0
Hosts Created         : 20
Hosts Updated         : 40
Hosts Deleted         : 0
Authentication Requests Sent : 40
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host summary
=====
ARP host Summary, service 2
=====
Sap                  Used      Provided   Admin State
-----
sap:2/1/5:2         20       8000      inService
-----
Number of SAPs : 1
=====
*A:Dut-C#

```

base

| | |
|--------------------|--|
| Syntax | base [msap] |
| Context | show>service>id show>service>id>igmp-snooping |
| Description | This command displays basic information about the service ID including service type, description, SAPs and SDPs. |
| Parameters | msap — Displays management SAPs. |
| 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. |

| Label | Description (Continued) |
|-------|--------------------------------|
| Opr | The operating state of the SAP |

Sample Output

```

A:~# show service id 300 sap 1/1/1:300.* detail
=====
Service Access Points(SAP)
=====
Service Id      : 300
SAP             : 1/1/1:300.*      Encap           : qinq
QinQ Dot1p     : Default
Admin State     : Up               Oper State      : Up
Flags           : None
Multi Svc Site  : None
Last Status Change : 11/19/2007 20:42:34
Last Mgmt Change  : 11/19/2007 20:42:25
Sub Type        : regular
Dot1Q Ethertype : 0x8100 QinQ      Ethertype       : 0x8100

Admin MTU       : 1522             Oper MTU        : 1522
Ingr IP Fltr-Id : n/a             Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a            Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a           Egr IPv6 Fltr-Id : n/a
tod-suite       : None             qinq-pbit-marking : both
Egr Agg Rate Limit : max           Endpoint        : N/A
Q Frame-Based Acct : Disabled
Vlan-translation : None

Acct. Pol       : None             Collect Stats    : Disabled
Ingress qos-policy : 1             Egress qos-policy : 1
Shared Q plcy    : n/a             Multipoint shared : Disabled
-----
Sap Statistics
-----
Last Cleared Time      : 11/19/2007 21:23:45

Packets      Octets
Forwarding Engine Stats
Dropped      : 0      0
Off. HiPrio   : 0      0
Off. LowPrio  : 0      0
Off. Uncolor  : 0      0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio   : 0      0
Dro. LowPrio  : 0      0
For. InProf   : 0      0
For. OutProf  : 0      0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf   : 0      0
Dro. OutProf  : 0      0
For. InProf   : 0      0
For. OutProf  : 0      0
-----
Sap per Queue stats
-----

```

Show, Clear, Debug Commands

```

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

Egress Queue 1
For. InProf                    : 0              0
For. OutProf                   : 0              0
Dro. InProf                    : 0              0
Dro. OutProf                   : 0              0
=====
*A:Sr-4#
*A:SetupCLI# show service id 2001 base
=====
Service Basic Information
=====
Service Id      : 2001          Vpn Id      : 0
Service Type    : i-VPLS
Customer Id     : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change  : 09/25/2007 21:45:59
Admin State     : Up           Oper State   : Down
MTU             : 1514         Def. Mesh VC Id : 2001
SAP Count       : 1           SDP Bind Count : 0
Snd Flush on Fail : Disabled   Host Conn Verify : Disabled
b-vpls Id       : 2002         Oper ISID      : 122
Snd Flush in bVpls: Disabled
-----
Service Access & Destination Points
-----
Identifier              Type      AdmMTU  OprMTU  Adm    Opr
-----
sap:1/1/12:2001.2001    qinq      1522    1522    Up      Down
-----
[<sap-id>] indicates a Managed SAP
=====
*A:SetupCLI#

*A:SetupCLI# show service id 2002 base
=====
Service Basic Information
=====
Service Id      : 2002          Vpn Id      : 0
Service Type    : b-VPLS
Customer Id     : 1
Last Status Change: 09/25/2007 21:12:01
Last Mgmt Change  : 09/25/2007 21:45:59
Admin State     : Up           Oper State   : Down
MTU             : 1530         Def. Mesh VC Id : 2002
SAP Count       : 2           SDP Bind Count : 2
Snd Flush on Fail : Disabled   Host Conn Verify : Disabled
Oper Backbone Src : 00:f7:f7:f7:f7:f7
-----
Related iVpls services for bVpls service 2002
-----
iVpls SvcId      Oper ISID      Admin          Oper

```

```

-----
2001                122                Up                Down
-----
Number of Entries : 1
-----
Service Access & Destination Points
-----
Identifier                Type                AdmMTU    OprMTU    Adm        Opr
-----
sap:1/1/12:2002.2002      qinq                1522      1522      Down       Down
sap:1/1/30:2002           q-tag               1518      1518      Down       Down
sdp:2000:2001 S(101.101.101.101) n/a                1500      1500      Down       Down
sdp:2000:2002 M(101.101.101.101) n/a                1500      1500      Down       Down
-----
[<sap-id>] indicates a Managed SAP
=====

```

```

A:ALA-48>config>service>vpls# show service id 700 base
=====
Service Basic Information
=====
Service Id      : 700                Vpn Id          : 0
Service Type    : VPLS
Description     : IMA VPLS
Customer Id     : 7
Last Status Change: 11/21/2008 17:33:20
Last Mgmt Change  : 11/21/2008 17:33:34
Admin State     : Up                Oper State       : Down
MTU             : 1514              Def. Mesh VC Id  : 700
SAP Count       : 1                SDP Bind Count   : 2
Snd Flush on Fail : Disabled        Host Conn Verify : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP  : None
Def. Gateway MAC : None
-----
BGP Auto-discovery Information
-----
Admin State     : Down              Vpls Id         : None
Route Dist      : None              Prefix           : 10.10.10.103
Rte-Target Import : None            Rte-Target Export : None
Vsi-Import      : None
Vsi-Export      : None
PW-Template Id  : None
-----
Service Access & Destination Points
-----
Identifier                Type                AdmMTU    OprMTU    Adm        Opr
-----
sap:1/1/9:0              q-tag               1518      1518      Up          Down
sdp:2:222 S(10.10.10.104) n/a                0          0          Up          Down
sdp:2:700 M(10.10.10.104) n/a                0          0          Up          Down
=====
A:ALA-48>config>service>vpls#

```

epipe

| | |
|-------------|---|
| Syntax | epipe |
| Context | show>service>id |
| Description | This command displays Epipe services associated with the B-VPLS service. The command only applies when the service is a B-VPLS. |
| Output | <pre>*A:term17>show>service>id# epipe ===== Related Epipe services for bVpls service 2000 ===== Epipe SvcId Oper ISID Admin Oper ----- 100 100 Down Down ----- Number of Entries : 1 ----- *A:term17>show>service>id#</pre> |

fdb

| | |
|--------------------|--|
| Syntax | fdb [sap <i>sap-id</i> [expiry]] [sdp <i>sdp-id</i> [expiry]] [mac <i>ieee-address</i> [expiry]] endpoint <i>endpoint</i> [detail] [expiry] [pbb] |
| Context | show>service>id show>service>fdb-mac |
| Description | This command displays FDB entries for a given MAC address. |
| Parameters | sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP. See Common CLI Command Descriptions on page 1319 for command syntax. detail — Displays detailed information. expiry — Displays time until MAC is aged out. pbb — Displays PBB related information. This keyword is only applicable to b-vpls or i-vpls services. <i>endpoint-name</i> — Specifies an endpoint name up to 32 characters in length. 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. |

| Label | Description (Continued) |
|--------------------|--|
| Table Size | Specifies the maximum number of learned and static entries allowed in the FDB of this service. The maximum value is 511999 when the value of the chassis mode is d . |
| Total Count | Displays the total number of learned entries in the FDB of this service. |
| 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. |
| DHCP-learned Count | Displays the current number of DHCP-learned entries in the FDB of this service. |
| Host-learned Count | Displays the current number of host-learned 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. |

| Label | Description (Continued) |
|-------------|--|
| | Age — Specifies the number of seconds used to age out TLS FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs. |
| | L — Learned - Dynamic entries created by the learning process. |
| | OAM — Entries created by the OAM process. |
| | H — Host, the entry added by the system for a static configured subscriber host. |
| | D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease. |
| | P — Indicates the MAC is protected by the MAC protection feature. |
| | Static — Statically configured. |
| Last Change | Indicates the time of the most recent state changes. |

Sample Output

```
A:ALA-48>show>service>id# fdb mac detail
=====
Service Forwarding Database
=====
ServId      MAC                Source-Identifier    Type/Age  Last Change
-----
6           00:aa:00:00:00:00  sap:lag-2           L/0       06/27/2006 15:04:31
6           00:aa:00:00:00:01  sap:lag-2           L/0       06/27/2006 15:04:31
6           00:aa:00:00:00:02  sap:lag-2           L/0       06/27/2006 15:04:31
6           00:aa:00:00:00:03  sap:lag-2           L/0       06/27/2006 15:04:31
6           00:aa:00:00:00:04  sap:lag-2           L/0       06/27/2006 15:04:31
10          12:12:12:12:12:12  sap:1/1/1:100       S         06/26/2006 10:03:29
=====
A:ALA-48>show>service>id#

A:PE-1# show service id 1 fdb detail
=====
Forwarding Database, Service 1
=====
ServId      MAC                Source-Identifier    Type      Last Change
                        Age
-----
1           00:00:00:00:00:01  sap:1/1/1           LP/0      02/24/12 11:40:07
-----
No. of MAC Entries: 1
-----
Legend:  L=Learned O=Oam P=Protected-MAC
=====
A:PE-1#
```

Show, Clear, Debug Commands

```
A:ALA-48# show service id 700 fdb
=====
Forwarding Database, Service <service-id>
=====
ServId      MAC                      Source-Identifier      Type/Age  Last Change
-----
1           aa:aa:aa:aa:aa:aa sdp:100:1            P          11/02/2006 06:04:03
-----
No. of MAC Entries: 1
=====
A:ALA-48#

*A:cses-B0102>show>service>id# fdb detail

=====
Forwarding Database, Service 510
=====
ServId      MAC                      Source-Identifier      Type      Last Change
                        Age
-----
510         00:00:00:aa:aa:aa sdp:1/1/22:510        CStatic   06/14/13 20:16:19
510         00:00:00:bb:bb:bb sdp:1/1/22:510        CStatic   06/14/13 20:14:49
510         00:00:00:dd:dd:dd sdp:7:2              Spb        06/14/13 20:03:23
510         d8:da:ff:00:00:00 sdp:1/1/22:510        CStatic   06/14/13 21:06:38
510         d8:e0:ff:00:00:00 sdp:7:2              Spb        06/14/13 21:09:29
-----
No. of MAC Entries: 5
-----
Legend:  L=Learned O=Oam P=Protected-MAC
=====

A:term17>config>service# show service id 2000 fdb pbb
(BVPLS = 2000, IVPLS = 2100)
=====
Forwarding Database, bVpls Service 2000
=====
MAC          Source-Identifier      iVplsMACs  Type/Age  Last Change
-----
00:f4:f4:f4:f4:f4 sdp:100:2000          10          L/O        09/25/2007 15:34:19
=====
A:term17>config>service#

*A:SetupCLI# show service id 2100 fdb pbb
=====
Forwarding Database, iVpls Service 2100
=====
MAC          Source-Identifier      B-Svc      bVpls MAC  Type/Age
-----
76:55:ff:00:01:a4 b-sdp:100:2000          2000        00:f4:f4:f4:f4:ff L/O
76:55:ff:00:01:bb sap:1/1/1:2100          2000        N/A         Static
=====
*A:SetupCLI#

A:term17>config>service# show service id 2100 fdb pbb
=====
Forwarding Database, iVpls Service 2100
=====
MAC          Source-Identifier      B-Svc      bVpls MAC  Type/Age
-----
```

```

00:f4:f4:f4:00:00 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:01 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:02 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:03 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:04 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:05 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:06 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:07 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:08 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f4:f4:f4:00:09 b-sdp:100:2000      2000      00:f4:f4:f4:f4:f4 L/0
00:f7:f7:f7:00:00 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:01 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:02 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:03 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:04 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:06 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:07 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:08 sap:lag-1:2100      2000      N/A          L/0
00:f7:f7:f7:00:09 sap:lag-1:2100      2000      N/A          L/0
=====
A:term17>config>service#

*A:SetupCLI# show service id 2100 fdb pbb
=====
Forwarding Database, iVpls Service 2100
=====
MAC                Source-Identifier      B-Svc      bVpls MAC      Type/Age
-----
76:55:ff:00:01:a4 b-sdp:100:2000      2000      00:f4:f4:f4:f4:ff L/0
76:55:ff:00:01:bb sap:1/1/1:2100      2000      N/A          Static
=====
*A:SetupCLI#

*A:term17>config>service>epipe# show service id 2000 fdb detail pbb
=====
Forwarding Database, bVpls Service 2000
=====
MAC                Source-Identifier      iVplsMACs   Epipes      Type/Age
-----
No Matching Entries
=====
*A:term17>config>service>epipe#

*A:term17>config>service>epipe# show service id 2100 fdb detail
=====
Forwarding Database, Service 2100
=====
ServId   MAC                Source-Identifier      Type/Age   Last Change
-----
No Matching Entries
=====
*A:term17>config>service>epipe# show service id 2100 fdb detail pbb
=====
Forwarding Database, iVpls Service 2100
=====
MAC                Source-Identifier      B-Svc      bVpls MAC      Type/Age
-----
No Matching Entries

```

```
=====
*A:term17>config>service>epipe#
```

egress-multicast-group

- Syntax** egress-multicast-group [group-name]
- Context** show>service
- Description** This command displays egress multicast group information.
- Parameters** group-name — Specifies the name of the egress multicast group.

Sample Output

```
A:Dut-C# show service egress-multicast-group emg1
=====
Egress Multicast Group Entry
=====
Group                : emg1
-----
Chain Limit          : 16                Encap Type          : dot1q
Dot1q ether type     : 0x8100            Filter-Id           : n/a
-----
Service Access Points
1/1/1:100
=====
A:Dut-C#
```

gsmp

- Syntax** gsmp
- Context** show>service>id
- Description** This command displays GSMP information.

neighbors

- Syntax** neighbors group [name] [ip-address]
- Context** show>service>id>gsmp
- Description** This command displays GSMP neighbor information.
- Parameters** group — A GSMP group defines a set of GSMP neighbors which have the same properties.
name — Specifies a GSMP group name is unique only within the scope of the service in which it is defined.
ip-address — Specifies the ip-address of the neighbor.

Sample Output

These commands show the configured neighbors per service, regardless of the fact there exists an open TCP connection with this neighbor. The admin state is shown because for a neighbor to be admin enabled, the service, gsmp node, group node and the neighbor node in this service must all be in 'no shutdown' state. Session gives the number of session (open TCP connections) for each configured neighbor.

```
A:active>show>service>id>gsmp# neighbors
=====
GSMP neighbors
=====
Group                               Neighbor                               AdminState  Sessions
-----
dslaml                             192.168.1.2                           Enabled     0
dslaml                             192.168.1.3                           Enabled     0
-----
Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslaml
=====
GSMP neighbors
=====
Group                               Neighbor                               AdminState  Sessions
-----
dslaml                             192.168.1.2                           Enabled     0
dslaml                             192.168.1.3                           Enabled     0
-----
Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslaml 192.168.1.2
=====
GSMP neighbors
=====
Group                               Neighbor                               AdminState  Sessions
-----
dslaml                             192.168.1.2                           Enabled     0
=====
A:active>show>service>id>gsmp#
```

sessions

| | |
|--------------------|--|
| Syntax | sessions [group name] neighbor <i>ip-address</i>] [port <i>port-number</i>] [association] [statistics] |
| Context | show>service>id>gsmp |
| Description | This command displays GSMP sessions information. |
| Parameters | group — A GSMP group defines a set of GSMP neighbors which have the same properties. <i>name</i> — Specifies a GSMP group name within the scope of the service in which it is defined. |

ip-address — Specifies the ip-address of the neighbor.

port — Specifies the neighbor TCP port number use for this ANCP session.

Values 0 — 65535

association — Displays to what object the ANCP-string is associated.

statistics — Displays statistics information about an ANCP session known to the system.

Description **Show Sessions Neighbor Output** — The following table describes show sessions neighbor output fields.

| Label | Description |
|------------------|--|
| State | The current state of the ANCP session. |
| Peer Instance | The instance number of the ANCP session at the neighbor's side. |
| Sender Instance | The instance number of the ANCP session at our side. |
| Peer Port | The port number of the ANCP session at the neighbor's side. |
| Sender Port | The port number of the ANCP session at the local side. |
| Peer Name | The MAC address of the ANCP session at the neighbor's side. |
| Sender name | The MAC address of the ANCP session at the local side. |
| timeouts | The number of adjacency protocol message timeouts. |
| Max. Timeouts | The maximum allowed of the above timeouts before closing. |
| Peer Timer | The timer value for the neighbor periodic adjacency protocol messages. |
| Sender Timer | The timer value for the local periodic adjacency protocol messages. |
| Capabilities | The negotiated capabilities for the Established ANCP session (DTD: dynamic topology discovery - OAM: operation and maintenance). |
| Conf Cap | The configured local capabilities. |
| Priority Marking | The DSCP bits for the IP messages used in the ANCP session. |
| Local Addr. | The destination IP address for this ANCP session. |
| Conf Local Addr. | The destination IP address accepted for ANCP connections. |

Sample Output

This show command gives information about the open TCP connections with DSLAMs.

```
A:active>show>service>id>gsmp# sessions
=====
GSMP sessions for service 999 (VPRN)
=====
```

```

Port    Ngbr-IpAddr    GsmP-Group
-----
40590   192.168.1.2       dslam1
-----
Number of GSMP sessions : 1
=====
A:active>show>service>id>gsmp#

```

```

A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590
=====
GSMP sessions for service 999 ), neighbor 192.168.1.2, Port 40590
=====
State           : Established
Peer Instance   : 1                Sender Instance : a3cf58
Peer Port       : 0                Sender Port      : 0
Peer Name       : 12:12:12:12:12:12 Sender Name      : 00:00:00:00:00:00
timeouts       : 0                Max. Timeouts   : 3
Peer Timer      : 100             Sender Timer     : 100
Capabilities    : DTD OAM
Conf Capabilities : DTD OAM
Priority Marking : dscp nc2
Local Addr.     : 192.168.1.4
Conf Local Addr. : N/A
=====
A:active>show>service>id>gsmp#

```

```

A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
=====
ANCP-Strings
=====
ANCP-String                                           Assoc. State
-----
No ANCP-Strings found
=====
A:active>show>service>id>gsmp#

```

```

A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 statistics
=====
GSMP session stats, service 999 neighbor 192.168.1.2, Port 40590
=====
Event                                     Received    Transmitted
-----
Dropped                                  0           0
Syn                                       1           1
Syn Ack                                  1           1
Ack                                       14          14
Rst Ack                                  0           0
Port Up                                  0           0
Port Down                                0           0
OAM Loopback                             0           0
=====
A:active>show>service>id>gsmp#

```

Note: The association command gives an overview of each ANCP string received from this session.

```

A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
=====
ANCP-Strings

```

```
=====
ANCP-String                                     Assoc.
State
-----
7330-ISAM-E47 atm 1/1/01/01:19425.64048          ANCP    Up
-----
Number of ANCP-Strings : 1
=====
A:active>show>service>id>gsm#
```

host

- Syntax** **host** [**sap** *sap-id*] [**detail**]
host summary
- Context** show>service>id
- Description** This command displays static host information configured on this service.
- Parameters** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 1319](#) for command syntax.
summary — Displays summary host information.

host-connectivity-verify

- Syntax** **host-connectivity-verify statistics** [**sap** *sap-id*]
- Context** show>service>id
- Description** This command displays host connectivity check statistics.
- Parameters** **statistics** — Displays host connectivity verification data.
sap *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 1319](#) for command syntax.
- Output** **Show Service Id Host Connectivity Verify** — The following table describes show service-id host connectivity verification output fields:

| Label | Description |
|----------------|--|
| Svc Id | The service identifier. |
| SapId/SdpId | The SAP and SDP identifiers. |
| DestIp Address | The destination IP address. |
| Last Response | The time when the last response was received. |
| Time Expired | Displays whether the interval value has expired. |

| Label | Description (Continued) |
|------------|--|
| Oper State | Displays the current operational state of the service. |

Sample Output

```
A:ALA-48>show>service>id# host-connectivity-verify statistics sap 1/1/9:0
=====
Host connectivity check statistics
=====
Svc    SapId/      DestIp      Last      Time      Oper
Id     SdpId      Address     Response  Expired   State
-----
1000 1/2/3:0143.144.145.1                               Up
=====
A:ALA-48>show>service>id#
```

i-vpls

| | |
|--------------------|--|
| Syntax | i-vpls |
| Context | show>service>id |
| Description | Displays i-vpls services associated with the b-vpls service. This command only applies when the service is a b-vpls. |
| Output | <pre>*A:SetupCLI# show service id 2002 i-vpls ===== Related iVpls services for bVpls service 2002 ===== iVpls SvcId Oper ISID Admin Oper ----- 2001 122 Up Down ----- Number of Entries : 1 ----- *A:term17>show>service>id# i-vpls ===== Related iVpls services for bVpls service 2000 ===== iVpls SvcId Oper ISID Admin Oper ----- 2100 2100 Up Up 2110 123 Up Up ----- Number of Entries : 2 ----- =====</pre> |

isis-using

- Syntax** isid-using [ISID]
- Context** show>service
- Description** This command displays services using an ISID.
- Parameters** ISID — 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. |

| Label | Description |
|--------|--|
| Type | Indicates whether the SDP is spoke or mesh. |
| I. Lbl | The VC label used by the far-end device to send packets to this device in this service by the SDP. |
| 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. |

Sample Output

```
A:ALA-48>show>service>id# l2pt
=====
L2pt summary, Service id 700
=====
```

| | L2pt-term enabled | L2pt-term disabled | Bpdu-trans auto | Bpdu-trans disabled | Bpdu-trans pvst | Bpdu-trans 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- termination | Admin Bpdu- translation | Oper Bpdu- 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 enabled | L2pt-term disabled | Bpdu-trans auto | Bpdu-trans disabled | Bpdu-trans pvst | Bpdu-trans 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- termination | Admin Bpdu- translation | Oper Bpdu- translation |
|---------|----------------------|----------------------------|---------------------------|
| 1/1/9:0 | disabled | disabled | disabled |

```
-----
```

Show, Clear, Debug Commands

```
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
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 | <pre> *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 </pre> |

mld-snooping

| | |
|--------------------|---|
| Syntax | mld-snooping |
| Context | show>service>id |
| Description | This command displays MLD snooping information. |

all

| | |
|--------------------|--|
| Syntax | all |
| Context | show>service>id>mld-snooping |
| Description | This command displays detailed information about MLD snooping. |

base

| | |
|--------------------|---|
| Syntax | base |
| Context | show>service>id>mld-snooping |
| Description | This command displays basic MLD snooping information. |

mrouters

| | |
|--------------------|--|
| Syntax | mrouters [detail] |
| Context | show>service>id>mld-snooping |
| Description | This command displays all multicast routers. |

mvr

| | |
|--------------------|--|
| Syntax | mvr |
| Context | show>service>id>mld-snooping |
| Description | This command displays multicast VPLS registration information. |

port-db

| | |
|--------------------|---|
| Syntax | port-db sap <i>sap-id</i> port-db sap <i>sap-id detail</i> port-db sap <i>sap-id group</i> <i>grp-ipv6-address</i> port-db sdp <i>sdp-id:vc-id</i> [detail] port-db sdp <i>sdp-id:vc-id group</i> <i>grp-ipv6-address</i> |
| Context | show>service>id>mld-snooping |
| Description | This command displays MLD snooping information related to a specific SAP. |

proxy-db

| | |
|--------------------|--|
| Syntax | proxy-db [detail] proxy-db group <i>grp-ip-address</i> |
| Context | show>service>id>mld-snooping |
| Description | This command displays proxy-reporting database entries. |
| Parameters | <i>grp-ip-address</i> — Displays the IGMP snooping proxy reporting database for a specific multicast group address. detail — Displays detailed information about the proxy-reporting database, |

querier

| | |
|--------------------|--|
| Syntax | querier |
| Context | show>service>id>mld-snooping |
| Description | This command displays information about the current querier. |

static

| | |
|----------------|---|
| Syntax | static [sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i>] |
| Context | show>service>id>mld-snooping |

Description This command displays MLD snooping static group membership data.

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
Context show>service>id>mld-snooping
Description This command displays MLD snooping statistics.

mrp-policy

Syntax **mrp-policy** [*mrp-policy*]
mrp-policy *mrp-policy* [**association**]
mrp-policy *mrp-policy* [**entry** *entry-id*]
Context show>service>id
Description This command displays information on an MRP policy.
Parameters *mrp-policy* — Specifies the MRP policy name.

Values 32 chars max
entry-id — Specifies the entry ID number.

Values 1..65535

Output *A:PE-B# show service mrp-policy

```
=====
Mrp Policies
=====
Mrp-Policy                               Scope      Applied Description
-----
1                                         template  Yes
2                                         template  Yes
-----
Total: 2
=====
```

*A:PE-B# show service mrp-policy "1"

```
=====
Mrp Policy
=====
Policy Name : 1                               Applied      : Yes
Scope       : template                       Def. Action   : block
Entries     : 1
Description  : (Not Specified)
-----
```

Mrp Policy Entries

```
-----
Entry       : 1                               Match action   : end-station
```

```
Description : (Not Specified)
isid         : 10..11
```

```
=====
*A:PE-B#
```

mmrp

- Syntax** `mmrp mac [ieee-address]`
- Context** `show>service>id`
- Description** This command displays information on MACs. If a MAC address is specified, information will be displayed relevant to the specific group. No parameter will display information on all group MACs on a server.
- Parameters** *ieee-address* — 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#
```

mstp-configuration

| | |
|-------------|--|
| Syntax | mstp-configuration |
| Context | show>service>id |
| Description | This command displays the MSTP specific configuration data. This command is only valid on a management VPLS. |

provider-tunnel

| | |
|-------------|--|
| Syntax | provider-tunnel |
| Context | show>service>id |
| 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
=====
A:PE71(1)#

```

proxy-nd

| | |
|--------------------|--|
| Syntax | proxy-nd [ip-address <i>ip-address</i>] [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
-----

```

Show, Clear, Debug Commands

```
Admin State      : enabled
Dyn Populate     : enabled
Age Time        : 200 secs          Send Refresh      : 120 secs

Dup Detect
-----
Detect Window    : 3 mins           Num Moves         : 3
Hold down       : max
Anti Spoof MAC   : 00:ca:ca:ca:ca:ca

EVPN
-----
Garp Flood       : disabled         Req Flood         : disabled
-----
A:PE71(1)# show service id 600 proxy-nd detail
-----
Proxy nd
-----
Admin State      : enabled
Dyn Populate     : enabled
Age Time        : 200 secs          Send Refresh      : 120 secs

Dup Detect
-----
Detect Window    : 3 mins           Num Moves         : 3
Hold down       : max
Anti Spoof MAC   : 00:ca:ca:ca:ca:ca

EVPN
-----
Garp Flood       : disabled         Req Flood         : disabled
-----

=====
VPLS Proxy ND Entries
=====
IP Address      Mac Address      Type      Status      Last Update
-----
172.16.0.1      00:ca:fe:ca:fe:02  evpn      active      12/01/2014 12:02:27
172.16.0.61     00:ca:de:ba:ca:00  dyn       active      12/01/2014 15:40:10
172.16.0.100    00:00:00:00:00:01  stat      inActv      12/01/2014 12:01:57
172.16.0.102    00:00:00:00:00:02  stat      inActv      12/01/2014 12:01:57
-----
Number of entries : 4
=====
A:PE71(1)#
```

retailers

| | |
|--------------------|---|
| Syntax | retailers |
| Context | show>service>id |
| Description | This command displays the service ID of the retailer subscriber service to which this DHCP lease belongs. |

wholesalers

| | |
|--------------------|---|
| Syntax | wholesalers |
| Context | show>service>id |
| Description | This command displays service wholesaler information. |

vxlan

| | |
|--------------------|---|
| Syntax | vxlan |
| Context | show>service>id show>service |
| Description | This command displays the VXLAN bindings auto-created in a 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
-----
60                  60           Up
-----
```

=====

sap

| | |
|--------------------|---|
| Syntax | sap <i>sap-id</i> [<i>filter</i>] |
| 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 <i>sap-id</i> — The ID that displays SAPs for the service in the <i>slot/mdal/port[.channel]</i> form. See Common CLI Command Descriptions on page 1319 for command syntax. detail — Displays detailed information for the SAP. <i>filter</i> — Specifies a search term to narrow down the results. Values base, detail, dhcp, mc-ring, mcac, mrp, qos, sap-stats, stats, stp, sub-mgmt |
| 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 Ethertype value. |
| Admin State | The administrative state of the SAP. |
| Oper State | The operational state of the SAP. |
| Flags | Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode. |
| Last Status Change | Specifies the time of the most recent operating status change to this SAP |
| Last Mgmt Change | Specifies the time of the most recent management-initiated change to this SAP. |
| Admin MTU | The largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented. |

| Label | Description (Continued) |
|-------------------------------------|---|
| Oper MTU | The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented. |
| Ingress qos-policy | The ingress QoS policy ID assigned to the SAP. |
| Egress qos-policy | The egress QoS policy ID assigned to the SAP. |
| Ingress Filter-Id | The ingress filter policy ID assigned to the SAP. |
| Egress Filter-Id | The egress filter policy ID assigned to the SAP. |
| Acct. Pol | The accounting policy ID assigned to the SAP. |
| Collect Stats | Specifies whether collect stats is enabled. |
| Forwarding Engine Stats | |
| Dropped | The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, etc. |
| Received Valid | The number of valid packets and octets received on the SAP. |
| Off. HiPrio | The number of high priority packets and octets, as determined by the SAP ingress QoS policy. |
| Off. LowPrio | The number of low priority packets and octets, as determined by the SAP ingress QoS policy. |
| Off. Uncolor | The number of uncolored packets and octets, as determined by the SAP ingress QoS policy. |
| 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. |

| Label | Description (Continued) |
|---------------------|---|
| For. InProf | The number of in-profile packets and octets (rate below CIR) forwarded. |
| For. OutProf | The number of out-of-profile packets and octets (rate above CIR) forwarded. |
| Ingress TD Profile | The profile ID applied to the ingress SAP. |
| Egress TD Profile | The profile ID applied to the egress SAP. |
| Alarm Cell Handling | The indication that OAM cells are being processed. |
| AAL-5 Encap | The AAL-5 encapsulation type. |

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
Learned MAC Addr   : 0
OAM MAC Addr       : 0
Host MAC Addr      : 0
SPB MAC Addr       : 0
BGP EVPN Addr      : 0
Admin MTU          : 1518
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
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      : n/a
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 BPDU 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
Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0
TCN BPDUs rcvd    : 0
TC bit BPDUs rcvd : 0
RST BPDUs rcvd    : 0
MST 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 : 1
Ingress FP QGrp    : (none)
Ing FP QGrp Inst   : (none)
Shared Q plcy      : n/a
Egress qos-policy  : 30
Egress Port QGrp   : (none)
Egr Port QGrp Inst : (none)
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)

Show, Clear, Debug Commands

```

Off. HiPrio           : 0           0
Off. LowPrio          : 0           0
Dro. HiPrio           : 0           0
Dro. LowPrio          : 0           0
For. InProf           : 0           0
For. OutProf          : 0           0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0           0
Off. LowPrio          : 0           0
Off. Managed          : 0           0
Dro. HiPrio           : 0           0
Dro. LowPrio          : 0           0
For. InProf           : 0           0
For. OutProf          : 0           0

Egress Queue 1
For. InProf           : 0           0
For. OutProf          : 0           0
Dro. InProf           : 0           0
Dro. OutProf          : 0           0
=====
*A:PE#

*A:PE-A# show service id 10 sap 2/2/5:10 mrp
=====
Service Access Points(SAP)
=====
Service Id           : 10
SAP                  : 2/2/5:10          Encap                  : q-tag
Description          : Default sap description for service id 10
Admin State          : Up                Oper State              : Up
Flags                : None
Multi Svc Site       : None
Last Status Change   : 01/16/2008 09:37:57
Last Mgmt Change     : 01/16/2008 09:37:41
-----
SAP MRP Information
-----
Join Time            : 0.2 secs          Leave Time              : 1.0 secs
Leave All Time        : 10.0 secs         Periodic Time           : 1.0 secs
Periodic Enabled     : false
Rx Pdus              : 11                Tx Pdus                : 12
Dropped Pdus         : 0                 Tx Pdus                : 12
Rx New Event         : 0                 Rx Join-In Event       : 150
Rx In Event          : 10                Rx Join Empty Evt      : 10
Rx Empty Event       : 10                Rx Leave Event         : 0
Tx New Event         : 0                 Tx Join-In Event       : 140
Tx In Event          : 0                 Tx Join Empty Evt      : 20
Tx Empty Event       : 10                Tx Leave Event         : 0
-----
SAP MMRP Information
-----
MAC Address          Registered          Declared
-----
01:1e:83:00:00:65 Yes                    Yes
01:1e:83:00:00:66 Yes                    Yes
01:1e:83:00:00:67 Yes                    Yes
01:1e:83:00:00:68 Yes                    Yes
01:1e:83:00:00:69 Yes                    Yes
01:1e:83:00:00:6a Yes                    Yes

```



```

01:1e:83:00:00:6b Yes Yes
01:1e:83:00:00:6c Yes Yes
01:1e:83:00:00:6d Yes Yes
01:1e:83:00:00:6e Yes Yes
-----
Number of MACs=10 Registered=10 Declared=10
-----
*A:PE-A#

```

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 or mesh. |
| Split Horizon Group | Indicates the name of the split horizon group that the SDP belongs to. |
| VC Type | Displays the VC type: ether, vlan, or vpls. |
| VC Tag | Displays the explicit dot1Q value used when encapsulating to the SDP far end. |
| I. Lbl | The VC label used by the far-end device to send packets to this device in this service by the SDP. |
| Admin Path MTU | The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.) |
- 7450 ESS Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN

Page 973

| Label | Description (Continued) |
|---------------------|--|
| 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. |
| 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 |

| Label | Description (Continued) |
|--------------|---|
| Peer Pw Bits | <p>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.</p> <p>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</p> |

Sample Output

```
*A:Dut-C# show service id 1001 sdp 17407:4294967295 detail
=====
Service Destination Point (Sdp Id : 17407:4294967295) Details
=====
-----

Sdp Id 17407:4294967295  -(0.0.0.0)
-----

Description      : (Not Specified)

SDP Id           : 17407:4294967295      Type           : VplsPmsi

Split Horiz Grp  : (Not Specified)

VC Type          : Ether                  VC Tag           : n/a

Admin Path MTU   : 9194                  Oper Path MTU    : 9194

Far End          : not applicable         Delivery         : MPLS

Tunnel Far End   : n/a                   LSP Types        : None

Hash Label       : Disabled              Hash Lbl Sig Cap : Disabled

Oper Hash Label  : Disabled

Admin State      : Up                    Oper State        : Up

Acct. Pol        : None                  Collect Stats     : Disabled

Ingress Label    : 0                     Egress Label     : 3

Ingr Mac Fltr-Id : n/a                   Egr Mac Fltr-Id  : n/a
```

Show, Clear, Debug Commands

| | | | |
|---------------------|-----------------------|--------------------|------------|
| Ingr IP Fltr-Id | : n/a | Egr IP Fltr-Id | : n/a |
| Ingr IPv6 Fltr-Id | : n/a | Egr IPv6 Fltr-Id | : n/a |
| Admin ControlWord | : Not Preferred | Oper ControlWord | : False |
| Last Status Change | : 01/31/2012 00:51:46 | Signaling | : None |
| Last Mgmt Change | : 01/31/2012 00:49:58 | Force Vlan-Vc | : Disabled |
| Endpoint | : N/A | Precedence | : 4 |
| PW Status Sig | : Enabled | | |
| Class Fwding State | : Down | | |
| Flags | : None | | |
| Time to RetryReset | : never | Retries Left | : 3 |
| Mac Move | : Blockable | Blockable Level | : Tertiary |
| Local Pw Bits | : None | | |
| Peer Pw Bits | : None | | |
| Peer Fault Ip | : None | | |
| Application Profile | : None | | |
| Max Nbr of MAC Addr | : No Limit | Total MAC Addr | : 0 |
| Learned MAC Addr | : 0 | Static MAC Addr | : 0 |
| | | | |
| MAC Learning | : Enabled | Discard Unkwn Srce | : Disabled |
| MAC Aging | : Enabled | | |
| BPDU Translation | : Disabled | | |
| L2PT Termination | : Disabled | | |
| MAC Pinning | : Disabled | | |
| Ignore Standby Sig | : False | Block On Mesh Fail | : False |
| Oper Group | : (none) | Monitor Oper Grp | : (none) |
| Rest Prot Src Mac | : Disabled | | |
| Auto Learn Mac Prot | : Disabled | RestProtSrcMacAct | : Disable |
| | | | |
| Ingress Qos Policy | : (none) | Egress Qos Policy | : (none) |
| Ingress FP QGrp | : (none) | Egress Port QGrp | : (none) |

```

Ing FP QGrp Inst      : (none)                      Egr Port QGrp Inst: (none)

-----

ETH-CFM SDP-Bind specifics

-----

V-MEP Filtering      : Disabled

KeepAlive Information :

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

Statistics           :

I. Fwd. Pkts.        : 0                            I. Dro. Pkts.       : 0
I. Fwd. Octs.         : 0                            I. Dro. Octs.       : 0
E. Fwd. Pkts.        : 5937639                      E. Fwd. Octets      : 356258340

MCAC Policy Name     :

MCAC Max Unconst BW: no limit                        MCAC Max Mand BW    : no limit
MCAC In use Mand BW: 0                              MCAC Avail Mand BW  : unlimited
MCAC In use Opnl BW: 0                              MCAC Avail Opnl BW  : unlimited

-----

RSVP/Static LSPs

-----

Associated LSP List :

No LSPs Associated

-----

Class-based forwarding :

-----

Class forwarding      : Disabled                      EnforceDSTELspFc    : Disabled
Default LSP           : Uknwn                         Multicast LSP        : None

=====
FC Mapping Table

```

Show, Clear, Debug Commands

```
=====
FC Name          LSP Name
-----
No FC Mappings

-----
Stp Service Destination Point specifics
-----

Stp Admin State   : Down                Stp Oper State    : Down
Core Connectivity : Down
Port Role         : N/A                Port State        : Forwarding
Port Number       : 0                  Port Priority      : 128
Port Path Cost    : 10                 Auto Edge         : Enabled
Admin Edge        : Disabled           Oper Edge         : N/A
Link Type         : Pt-pt              BPDU Encap        : Dot1d
Root Guard        : Disabled           Active Protocol    : N/A
Last BPDU from    : N/A
Designated Bridge : N/A                Designated Port Id: N/A

Fwd Transitions   : 0                  Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0                  Cfg BPDUs tx      : 0
TCN BPDUs rcvd    : 0                  TCN BPDUs tx      : 0
TC bit BPDUs rcvd : 0                  TC bit BPDUs tx   : 0
RST BPDUs rcvd    : 0                  RST BPDUs tx      : 0

-----
Number of SDPs : 1
-----
=====

A:Dut-A# show service id 1 sdp detail
=====
Services: Service Destination Points Details
=====
Sdp Id 1:1  -(10.20.1.2)
```

```

-----
Description      : Default sdp description

SDP Id           : 1:1                               Type           : Spoke
VC Type          : Ether                             VC Tag          : n/a
Admin Path MTU    : 0                               Oper Path MTU    : 9186
Far End          : 10.20.1.2                         Delivery         : MPLS

Admin State       : Up                               Oper State       : Up
Acct. Pol        : None                             Collect Stats    : Disabled
Ingress Label     : 2048                             Egress Label     : 2048
Ing mac Fltr      : n/a                             Egr mac Fltr     : n/a
Ing ip Fltr       : n/a                             Egr ip Fltr      : n/a
Ing ipv6 Fltr     : n/a                             Egr ipv6 Fltr    : n/a
Admin ControlWord : Not Preferred                    Oper ControlWord  : False
Last Status Change : 05/31/2007 00:45:43             Signaling        : None
Last Mgmt Change  : 05/31/2007 00:45:43

Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr : No Limit                       Total MAC Addr   : 0
Learned MAC Addr   : 0                               Static MAC Addr   : 0

MAC Learning       : Enabled                         Discard Unkwn Srce : Disabled
MAC Aging          : Enabled                         BPDU Translation  : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled

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

Statistics         :
I. Fwd. Pkts.      : 0                               I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                               I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                               E. Fwd. Octets    : 0
MCAC Policy Name   :
MCAC Max Unconst BW : no limit                       MCAC Max Mand BW  : no limit
MCAC In use Mand BW : 0                             MCAC Avail Mand BW : unlimited
MCAC In use Opnl BW : 0                             MCAC Avail Opnl BW : unlimited

Associated LSP LIST :
Lsp Name           : A_B_1
Admin State        : Up                               Oper State        : Up
Time Since Last Tr* : 00h26m35s

Lsp Name           : A_B_2
Admin State        : Up                               Oper State        : Up
Time Since Last Tr* : 00h26m35s

Lsp Name           : A_B_3
Admin State        : Up                               Oper State        : Up
Time Since Last Tr* : 00h26m34s

Lsp Name           : A_B_4
Admin State        : Up                               Oper State        : Up
Time Since Last Tr* : 00h26m34s

```

Show, Clear, Debug Commands

```

Lsp Name          : A_B_5
Admin State       : Up
Time Since Last Tr*: 00h26m34s
Oper State        : Up

Lsp Name          : A_B_6
Admin State       : Up
Time Since Last Tr*: 00h26m34s
Oper State        : Up

Lsp Name          : A_B_7
Admin State       : Up
Time Since Last Tr*: 00h26m34s
Oper State        : Up

Lsp Name          : A_B_8
Admin State       : Up
Time Since Last Tr*: 00h26m35s
Oper State        : Up

Lsp Name          : A_B_9
Admin State       : Up
Time Since Last Tr*: 00h26m34s
Oper State        : Up

Lsp Name          : A_B_10
Admin State       : Up
Time Since Last Tr*: 00h26m34s
Oper State        : Up
-----
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
hl                   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
Core Connectivity     : Down
Port Role             : N/A
Port Number           : 2049
Port Path Cost        : 10
Admin Edge            : Disabled
Link Type             : Pt-pt
Root Guard            : Disabled
Last BPDU from        : N/A
Designated 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 Id: 0

Fwd Transitions       : 0
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0
Bad BPDUs rcvd       : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 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)

```

Show, Clear, Debug Commands

```

-----
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
Learned MAC Addr   : 0                               Total MAC Addr    : 0
OAM MAC Addr       : 0                               Static MAC Addr   : 0
Host MAC Addr      : 0                               DHCP MAC Addr     : 0
SPB MAC Addr       : 0                               Intf MAC Addr     : 0
Cond MAC Addr      : 0                               Cond MAC Addr     : 0

MAC Learning       : Enabled                          Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                            Block On Mesh Fail: False
Oper Group         : (none)                           Monitor Oper Grp  : (none)
Rest Prot Src Mac   : Disabled
Auto Learn Mac Prot: Disabled                         RestProtSrcMacAct : Disable

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

-----
ETH-CFM SDP-Bind specifics
-----

```

```

V-MEP Filtering      : Disabled

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

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

Squelch Levels       : 0 1 2 3 4 5 6 7

```

site

| | |
|--------------------|---|
| Syntax | site [detail] site <i>name</i> |
| Context | show>service>id |
| Description | This command displays sites configures for the service. |
| Parameters | <i>name</i> — Specifies the site name. |
| Values | 32 chars max |

split-horizon-group

| | |
|--------------------|--|
| Syntax | split-horizon-group [<i>group-name</i>] |
| Context | show>service>id |
| Description | <p>This command displays service split horizon groups.</p> <pre> *A:ALA-1# show service id 700 split-horizon-group ===== Service: Split Horizon Group ===== Name Description ----- R DSL-group1 Split horizon group for DSL ----- R = Residential Split Horizon Group No. of Split Horizon Groups: 1 ===== *A:ALA-1# *A:ALA-1# show service id 700 split-horizon-group DSL-group1 ===== Service: Split Horizon Group ===== Name Description ----- R DSL-group1 Split horizon group for DSL ----- </pre> |

Show, Clear, Debug Commands

```
Associations
-----
SAP                1/1/3:1
SDP                108:1
SDP                109:1
-----
R = Residential Split Horizon Group
SAPs Associated : 1          SDPs Associated : 2
=====
*A:ALA-1#
```

stp

| | |
|--------------------|--|
| Syntax | stp [detail] stp mst-instance <i>mst-inst-number</i> |
| Context | show>service>id |
| Description | This command displays information for the spanning tree protocol instance for the service. |
| Parameters | detail — Displays detailed information. <i>mst-inst-number</i> — 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 |
|-------------------|--|
| RSTP Admin State | Indicates the administrative state of the Rapid Spanning Tree Protocol instance associated with this service. |
| Core Connectivity | Indicates the connectivity status to the core. |
| RSTP Oper State | Indicates the operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router. |
| Bridge-id | Specifies the MAC address used to identify this bridge in the network. |
| Hold Time | Specifies the interval length during which no more than two Configuration BPDUs shall be transmitted by this bridge. |
| Bridge fwd delay | Specifies how fast a bridge changes its state when moving toward the forwarding state. |
| Bridge Hello time | Specifies the amount of time between the transmission of Configuration BPDUs. |
| Bridge max age | Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. This is the actual value that this bridge is currently using. |
| Bridge priority | Defines the priority of the Spanning Tree Protocol instance associated with this service. |
| Topology change | Specifies whether a topology change is currently in progress. |
| Last Top. change | Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service. |
| Top. change count | Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management entity was last reset or initialized. |

| Label | Description (Continued) |
|-----------------------|---|
| 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. |
| Root hello time | Specifies the amount of time between the transmission of configuration BPDUs. |
| Root max age | Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. |
| Root priority | This object specifies the priority of the bridge that is currently selected as root-bridge for the network. |
| Root port | Specifies the port number of the port which offers the lowest cost path from this bridge to the root bridge. |
| SAP Identifier | The ID of the access port where this SAP is defined. |
| RSTP State | The operational state of RSTP. |
| STP Port State | Specifies the port identifier of the port on the designated bridge for this port's segment. |
| BPDU encap | Specifies the type of encapsulation used on BPDUs sent out and received on this SAP. |
| Port Number | Specifies the value of the port number field which is contained in the least significant 12 bits of the 16-bit port ID associated with this SAP. |
| Priority | Specifies the value of the port priority field which is contained in the most significant 4 bits of the 16-bit port ID associated with this SAP. |
| Cost | Specifies the contribution of this port to the path cost of paths towards the spanning tree root which include this port. |
| Fast Start | Specifies whether Fast Start is enabled on this SAP. |
| Designated Port | Specifies the port identifier of the port on the designated bridge for this port's segment. |
| Designated Bridge | Specifies the bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment. |
| Service Access Points | |
| Managed by Service | Specifies the service ID of the management VPLS managing this SAP or spoke SDP. |

| Label | Description (Continued) |
|----------------------|---|
| Managed by SAP/spoke | Specifies the SAP ID or SDP ID inside the management VPLS managing this SAP or spoke SDP. |
| Prune state | Specifies the STP state inherited from the management VPLS. |

Sample Output

```
*A:ALA-12# show service id 11 stp
=====
Stp info, Service 11
=====
Bridge Id       : 80:00.22:68:ff:00:00:00  Top. Change Count : 1
Root Bridge     : 00:00.22:69:ff:00:00:00  Stp Oper State    : Syncing Vcp
Primary Bridge  : N/A                      Topology Change    : Inactive
Mode            : Mstp                      Last Top. Change   : 0d 19:12:58
Vcp Active Prot. : N/A
Root Port       : 2048                      External RPC        : 10
=====
MSTP specific info for CIST
=====
Regional Root   : This Bridge                Root Port          : 2048
Internal RPC    : 0                          Remaining Hopcount : 20
=====
Stp port info for CIST
=====
Sap/Sdp Id      Oper-   Port-   Port-   Port-   Oper-   Link-   Active
                  State   Role    State   Num     Edge    Type    Prot.
-----
1/1/1:0         Up      Root    Forward 2048    False   Pt-pt   Mstp
1/1/3:0         Up      N/A     Forward 2049    N/A     Pt-pt   N/A
1/1/4:*         Up      Designated Forward 2050    False   Pt-pt   Mstp
=====
MSTP specific info for MSTI 111
=====
Regional Root   : 80:6f.1c:65:ff:00:00:00  Root Port          : 2050
Internal RPC    : 10                          Remaining Hopcount : 19
=====
MSTP port info for MSTI 111
=====
Sap/Sdp Id      Oper-   Port-   Port-   Port-   Same
                  State   Role    State   Num     Region
-----
1/1/1:0         Up      Master  Forward 2048    False
1/1/3:0         Up      N/A     Forward 2049    N/A
1/1/4:*         Up      Root    Forward 2050    True
=====
*A:ALA-12#

*A:ALA-12# show service id stp detail
=====
Spanning Tree Information
=====
VPLS Spanning Tree Information
```

Show, Clear, Debug Commands

```

-----
VPLS oper state      : Up                      Core Connectivity : Down
Stp Admin State      : Up                      Stp Oper State       : Up
Mode                 : Mstp                    Vcp Active Prot.    : N/A

Bridge Id            : 80:00.22:68:ff:00:00:00 Bridge Instance Id: 0
Bridge Priority       : 32768                  Tx Hold Count        : 6
Topology Change      : Inactive                Bridge Hello Time     : 2
Last Top. Change     : 0d 19:14:34             Bridge Max Age        : 20
Top. Change Count    : 1                      Bridge Fwd Delay      : 15
MST region revision  : 0                      Bridge max hops       : 20
MST region name      : abc

Root Bridge          : 00:00.22:69:ff:00:00:00
Primary Bridge       : N/A

Root Path Cost       : 10                      Root Forward Delay    : 15
Rcvd Hello Time      : 2                      Root Max Age          : 20
Root Priority         : 0                      Root Port              : 2048

MSTP info for CIST :
Regional Root        : This Bridge              Root Port              : 2048
Internal RPC         : 0                      Remaining Hopcount     : 20
MSTP info for MSTI 111 :
Regional Root        : 80:6f.1c:65:ff:00:00:00 Root Port              : 2050
Internal RPC         : 10                     Remaining Hopcount     : 19
-----
Spanning Tree Virtual Core Port (VCP) Specifics
-----
Mesh Sdp Id          Sdp          Sdp Bind    Mesh Sdp    HoldDown    Awaiting
                    Oper-state   Oper-state  Port-state  Timer       Agreement
-----
3:11                 Down        Down        Discard     Inactive    N/A
4:11                 Down        Down        Discard     Inactive    N/A
-----
Spanning Tree Sap/Spoke SDP Specifics
-----
SAP Identifier       : 1/1/1:0                      Stp Admin State       : Up
Port Role            : Root                      Port State            : Forwarding
Port Number          : 2048                      Port Priority         : 128
Port Path Cost       : 10                      Auto Edge             : Enabled
Admin Edge           : Disabled                  Oper Edge             : False
Link Type            : Pt-pt                    BPDU Encap           : Dot1d
Root Guard           : Disabled                  Active Protocol       : Mstp
Last BPDU from       : 00:00.22:69:ff:00:00:00 Inside Mst Region     : False
CIST Desig Bridge    : 00:00.22:69:ff:00:00:00 Designated Port       : 34816
MSTI 111 Port Prio   : 128                      Port Path Cost        : 10
MSTI 111 Desig Brid  : This Bridge                Designated Port       : 34816
Forward transitions   : 1                      Bad BPDUs rcvd        : 0
Cfg BPDUs rcvd       : 0                      Cfg BPDUs tx          : 0
TCN BPDUs rcvd       : 0                      TCN BPDUs tx          : 0
RST BPDUs rcvd       : 0                      RST BPDUs tx          : 0
MST BPDUs rcvd       : 34638                     MST BPDUs tx          : 3

SAP Identifier       : 1/1/3:0                      Stp Admin State       : Down
Port Role            : N/A                      Port State            : Forwarding
Port Number          : 2049                      Port Priority         : 128
Port Path Cost       : 10                      Auto Edge             : Enabled
Admin Edge           : Disabled                  Oper Edge             : N/A
Link Type            : Pt-pt                    BPDU Encap           : Dot1d
Root Guard           : Disabled                  Active Protocol       : N/A

```



```

Last BPDUs from      : N/A
CIST Desig Bridge    : N/A
MSTI 111 Port Prio   : 128
MSTI 111 Desig Brid  : N/A
Forward transitions: 1
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0
MST BPDUs rcvd       : 0

Designated Port      : 0
Port Path Cost       : 10
Designated Port      : 0
Bad BPDUs rcvd       : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 0
RST BPDUs tx         : 0
MST BPDUs tx         : 0

SAP Identifier       : 1/1/4:*
Port Role            : Designated
Port Number          : 2050
Port Path Cost       : 10
Admin Edge           : Disabled
Link Type            : Pt-pt
Root Guard           : Disabled
Last BPDUs from      : 50:00.1c:65:ff:00:00:00
CIST Desig Bridge    : This Bridge
MSTI 111 Port Prio   : 128
MSTI 111 Desig Brid  : 80:6f.1c:65:ff:00:00:00
Forward transitions: 1
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0
MST BPDUs rcvd       : 34636

Stp Admin State      : Up
Port State           : Forwarding
Port Priority         : 128
Auto Edge            : Enabled
Oper Edge            : False
BPDU Encap           : Dot1d
Active Protocol      : Mstp
Inside Mst Region    : True
Designated Port      : 34818
Port Path Cost       : 10
Designated Port      : 34819
Bad BPDUs rcvd       : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 0
RST BPDUs tx         : 0
MST BPDUs tx         : 34640

=====
*A:ALA-12#
*A:SetupCLI# show service id 2001 stp
=====
Stp info, Service 2001
=====
Bridge Id           : 80:00.70:ec:ff:00:00:00
Root Bridge         : N/A
Primary Bridge      : N/A
Mode                : Rstp
Vcp Active Prot.    : N/A
Root Port           : N/A
Top. Change Count   : 0
Stp Oper State      : Down
Topology Change     : Inactive
Last Top. Change    : 0d 00:00:00
External RPC        : 0
=====
Stp port info
=====
Sap/    PIP Id      Oper-   Port-   Port-   Port-   Oper-   Link-   Active
          State     Role    State   Num     Edge    Type    Prot.
-----
Backbone VPLS      Down    N/A     Discard 2048    N/A     N/A     N/A
1/1/12:2001.2001  Down    N/A     Disabled 2049    N/A     Pt-pt  N/A
=====
*A:SetupCLI#

*A:SetupCLI# show service id 2001 stp detail
=====
Spanning Tree Information
-----
VPLS Spanning Tree Information
-----
VPLS oper state     : Down
Stp Admin State      : Down
Mode                : Rstp
Core Connectivity    : Down
Stp Oper State       : Down
Vcp Active Prot.     : N/A

Bridge Id           : 80:00.70:ec:ff:00:00:00
Bridge Instance Id   : 0

```

Show, Clear, Debug Commands

```
Bridge Priority      : 32768                      Tx Hold Count      : 6
Topology Change     : Inactive                    Bridge Hello Time   : 2
Last Top. Change    : 0d 00:00:00                 Bridge Max Age      : 20
Top. Change Count   : 0                          Bridge Fwd Delay    : 15
MST region revision : 0                          Bridge max hops     : 20
MST region name     :

Root Bridge         : N/A
Primary Bridge      : N/A

Root Path Cost      : 0                          Root Forward Delay  : 15
Rcvd Hello Time     : 2                          Root Max Age        : 20
Root Priority        : 32768                      Root Port           : N/A
-----
Spanning Tree Sap/Spoke SDP Specifics
-----
SAP Identifier      : 1/1/12:2001.2001           Stp Admin State     : Up
Port Role           : N/A                       Port State          : Unknown
Port Number         : 2049                      Port Priority        : 128
Port Path Cost      : 10                       Auto Edge           : Enabled
Admin Edge          : Disabled                  Oper Edge           : N/A
Link Type           : Pt-pt                     BPDU Encap          : Dot1d
Root Guard          : Disabled                  Active Protocol      : N/A
Last BPDU from      : N/A                      Designated Port      : N/A
CIST Desig Bridge   : N/A                      Bad BPDUs rcvd      : 0
Forward transitions : 0                        Cfg BPDUs tx        : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx        : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx        : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx        : 0
-----
Spanning Tree PIP (Provider Internal Port) Specifics
-----
Oper Status         : Down                      mVPLS Prune State   : N/A
Port Num            : 2048                      Oper Protocol        : N/A
Port Role           : N/A                       Port State           : Discarding
CIST Desig Bridge   : N/A                      Designated Port      : N/A
b-Vpls STP state    : Disabled
Forward transitions : 0                        Bad BPDUs rcvd      : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx        : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx        : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx        : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx        : 0
=====
*A:SetupCLI#
```

subscriber-hosts

| | |
|--------------------|--|
| Syntax | subscriber-hosts [sap <i>sap-id</i>] [ip <i>ip-address[/mask]</i>] [mac <i>ieee-address</i>] [sub-profile <i>sub-profile-name</i>] [sla-profile <i>sla-profile-name</i>] [detail] |
| Context | show>service>id |
| Description | This command displays subscriber host information. |
| Parameters | sap <i>sap-id</i> — Displays the specified subscriber host SAP information. See Common CLI Command Descriptions on page 1319 for command syntax. |

ip-address/mask — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values Allowed values are IP addresses in the range 1.0.0.0 — 223.255.255.255 (with support of /31 subnets).
mask: 1 — 32

mac *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sub-profile *sub-profile-name* — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile *sla-profile-name* — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

detail — Displays detailed information.

Sample Output

```
A:ALA#-SR12# show service id 20 subscriber-hosts
=====
Subscriber Host table
=====
Sap Id                IP Address          MAC Address          Origin(*) Subscriber
-----
1/2/6:0              101.1.1.10          00:bb:bb:00:00:00 S/-/-
    Eval-20-static
-----
Number of subscriber hosts : 1
=====
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:ALA#

A:ALA# show service id 10 subscriber-hosts
=====
Subscriber Host table
=====
Sap Id                IP Address          MAC Address          Origin(*) Subscriber
-----
1/2/5:0              100.1.1.10          00:aa:aa:00:00:01 -/D/-
    SUB-10-00aaaa000001
-----
Number of subscriber hosts : 1
=====
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:ALA-SR12#
```

statistics

| | |
|--------------------|---|
| Syntax | statistics [<i>policy name</i>] [sap <i>sap-id</i>] |
| Context | show>service>id>authentication |
| Description | This command displays session authentication statistics for this service. |
| Parameters | sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1319 for command syntax. |

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
=====
Authentication statistics
=====
Interface / SAP                Authentication    Authentication
                               Successful         Failed
-----
vpls-11-90.1.0.254             1582             3
-----
Number of entries: 1
=====
*A:ALA-1#
```

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. |
| MVR From VPLS | Specifies MVR from VPLS. |
| Num Groups | Specifies the actual number of multicast groups that can be joined on this SAP. |

Sample Output

```
A:ALA-42>show>service>id>igmp-snooping>snooping# all
```

```
=====
```

Show, Clear, Debug Commands

```

IGMP Snooping info for service 100
=====
IGMP Snooping Base info
-----
Admin State : Up
Querier      : 10.20.1.6 on SAP 2/2/1:100
-----
Sap/Sdp      Oper   MRtr  Send   Max Num  MVR      Num
Id           State  Port  Queries Groups  From-VPLS Groups
-----
sap:1/1/4    Up     No    Disabled No Limit Local     1
sap:2/2/1:100 Up     No    Disabled No Limit Local     0
-----
IGMP Snooping Querier info
-----
Sap Id       : 2/2/1:100
IP Address   : 10.20.1.6
Expires      : 3s
Up Time      : 0d 00:15:23
Version      : 3

General Query Interval : 2s
Query Response Interval : 1.0s
Robust Count           : 2
-----
IGMP Snooping Multicast Routers
-----
MRouter      Sap/Sdp Id      Up Time      Expires      Version
-----
10.20.1.6    2/2/1:100      0d 00:15:24  2s           3
-----
Number of mrouter: 1
-----
IGMP Snooping Proxy-reporting DB
-----
Group Address Mode    Up Time      Num Sources
-----
225.0.0.0     exclude 0d 00:00:04  0
-----
Number of groups: 1
-----
IGMP Snooping SAP 1/1/4 Port-DB
-----
Group Address Mode    Type    From-VPLS Up Time      Expires      Num Src
-----
225.0.0.0     exclude dynamic local      0d 00:00:05  4s           0
-----
Number of groups: 1
-----
IGMP Snooping SAP 2/2/1:100 Port-DB
-----
Group Address Mode    Type    From-VPLS Up Time      Expires      Num Src
-----
Number of groups: 0
-----
IGMP Snooping Static Source Groups
-----
IGMP Snooping Statistics
-----
Message Type      Received      Transmitted    Forwarded
-----

```

| | | | |
|----------------------|-----|-----|-----|
| General Queries | 463 | 0 | 463 |
| Group Queries | 0 | 0 | 0 |
| Group-Source Queries | 0 | 0 | 0 |
| V1 Reports | 0 | 0 | 0 |
| V2 Reports | 0 | 0 | 0 |
| V3 Reports | 4 | 4 | 0 |
| V2 Leaves | 0 | 0 | 0 |
| Unknown Type | 0 | N/A | 0 |

Drop Statistics

| | |
|-------------------|-----|
| Bad Length | : 0 |
| Bad IP Checksum | : 0 |
| Bad IGMP Checksum | : 0 |
| Bad Encoding | : 0 |
| No Router Alert | : 0 |
| Zero Source IP | : 0 |

| | |
|-------------------------|-----|
| Send Query Cfg Drops | : 0 |
| Import Policy Drops | : 0 |
| Exceeded Max Num Groups | : 0 |

| | |
|-------------------------|-----|
| MVR From VPLS Cfg Drops | : 0 |
| MVR To SAP Cfg Drops | : 0 |

IGMP Snooping Multicast VPLS Registration info

IGMP Snooping Admin State : Up

| | |
|-----------------|--------|
| MVR Admin State | : Down |
| MVR Policy | : None |

Local SAPs/SDPs

| Svc Id | Sap/Sdp Id | Oper State | From VPLS | Num Local Groups |
|--------|---------------|------------|-----------|------------------|
| 100 | sap:1/1/4 | Up | Local | 1 |
| 100 | sap:2/2/1:100 | Up | Local | 0 |

MVR SAPs (from-vpls=100)

| Svc Id | Sap/Sdp Id | Oper State | From VPLS | Num MVR Groups |
|--------|------------|------------|-----------|----------------|
|--------|------------|------------|-----------|----------------|

No MVR SAPs found.

A:ALA-42>show>service>id>snooping#

mfib

| | |
|--------------------|--|
| Syntax | mfib [brief statistics] [ip mac] brief mfib [group <i>grp-address</i> *] [statistics] |
| Context | show>service>id |
| Description | This command displays the multicast FIB on the VPLS service. |
| Parameters | brief — Displays a brief output. statistics — Displays statistics on the multicast FIB. ip — Displays IP address information. mac — Displays MAC address information. group <i>grp grp-address</i> — Displays the multicast FIB for a specific multicast group address. |
| Output | Show Output — The following table describes the command output fields: |

| Label | Description |
|---------------------|--|
| Source Address | IPv4 unicast source address. |
| Group Address | IPv4 multicast group address. |
| SAP ID | Indicates the SAP/SDP to which the corresponding multicast stream will be forwarded/blocked. |
| Forwarding/Blocking | Indicates whether the corresponding multicast stream will be blocked/forwarded. |
| Number of Entries | Specifies the number of entries in the MFIB. |
| Forwarded Packets | Indicates the number of multicast packets forwarded for the corresponding source/group. |
| Forwarded Octets | Indicates the number of octets forwarded for the corresponding source/group. |
| Svc ID | Indicates the service to which the corresponding multicast stream will be forwarded/blocked. Local means that the multicast stream will be forwarded/blocked to a SAP or SDP local to the service. |

Sample Output

```
*A:ALA-SR12-D# mfib
=====
Multicast FIB, Service 10
=====
Source Address  Group Address  Sap/Sdp Id          Svc Id  Fwd/Blk
-----
*                226.1.1.1    sap:1/1/1:10        Local   Fwd
                  sap:1/1/2         20                  Fwd
                  sap:1/1/4:100     20                  Fwd
```



```

*                226.1.1.2      sap:1/1/4:200      20      Fwd
*                226.1.1.2      sap:1/1/4:200      20      Fwd
-----
Number of entries: 2
=====
*A:ALA-SR12-D#
*A:ALA-SR12-D# show service id 10 mfib statistics
=====
IGMP Snooping MFIB for service 10
=====
Source Address  Group Address  Fwd Pkts      Fwd Octets
-----
1.1.1.1        225.0.0.1      291           9281
1.1.1.2        225.0.0.1      0             0
-----
Number of entries: 2
=====
*A:ALA-SR12-D#

*A:PE-A# show service id 10 mfib
=====
Multicast FIB, Service 10
=====
Source Address  Group Address  Sap/Sdp Id      Svc Id  Fwd/Blk
-----
*                01:1E:83:00:00:65  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:66  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:67  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:68  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:69  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:6A  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:6B  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:6C  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:6D  sap:2/2/5:10    Local   Fwd
*                01:1E:83:00:00:6E  sap:2/2/5:10    Local   Fwd
-----
Number of entries: 10
=====
*A:PE-A#

```

To show which I-SIDs are local, the MFIB command will display ISIDs that are local and advertised. Static I-SIDs are included in this display. However, I-SID policy can override the I-SIDS that are designated to use the default multicast tree and these do not show up in the mfib. This is displayed on a B-VPLS control service.

```

*A:cses-B0102>show>service>id# mfib
=====
Multicast FIB, Service 510
=====
Source Address  Group Address  Sap/Sdp Id      Svc Id  Fwd/Blk
-----
*                01:1E:83:00:01:F4  b-sap:1/1/22:510  Local   Fwd
*                01:1E:83:00:01:F5  b-sap:1/1/22:510  Local   Fwd
*                01:1E:83:00:01:F6  b-sap:1/1/22:510  Local   Fwd
*                01:1E:83:00:01:F7  b-sap:1/1/22:510  Local   Fwd
*                01:1E:83:00:01:F8  b-sap:1/1/22:510  Local   Fwd
*                01:1E:83:00:01:F9  b-sap:1/1/22:510  Local   Fwd

```

```
*          01:1E:83:00:01:FA      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:01:FB      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:01:FC      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:01:FD      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:01:FE      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:01:FF      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:02:00      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:02:01      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:02:02      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:02:03      b-sap:1/1/22:510      Local      Fwd
*          01:1E:83:00:02:04      b-sap:1/1/22:510      Local      Fwd
-----
Number of entries: 21
=====
```

To show the I-SID policy under a B-VPLS, the I-SID policy is used.

```
*A:cses-B07>show>service>id# isid-policy

=====
Isid Policy Range
=====
Entry      Range              AdvLocal  UseDefMCTree
-----
2          1500-1600          Disabled  Enabled
=====
```

The following example shows the MFIB for an EVPN-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
=====
```

mroute

| | |
|-------------|--|
| Syntax | mroute [detail] |
| Context | show>service>id>igmp-snooping |
| Description | This command displays all multicast routers. |
| Parameters | detail — Displays detailed information. |

Sample Output

```
*A:ala-427# show service id 1 igmp-snooping mroute
=====
```

```

IGMP Snooping Multicast Routers for service 1
=====
MRouter          Sap/Sdp Id          Up Time          Expires          Version
-----
10.10.1.1         1/1/5:1          0d 00:00:26      14s              3
10.20.1.6         1/1/2:1          0d 00:10:16      2s              3
-----

Number of mrouter: 2
=====
*A:ala-427#

*A:ala-427# show service id 1 igmp-snooping mrouter detail
=====
IGMP Snooping Multicast Routers for service 1
=====
MRouter 10.10.1.1
-----
Sap Id           : 1/1/5:1
Expires          : 17s
Up Time          : 0d 00:00:32
Version          : 3

General Query Interval : 10s
Query Response Interval : 1.0s
Robust Count          : 2
-----
MRouter 10.20.1.6
-----
Sap Id           : 1/1/2:1
Expires          : 3s
Up Time          : 0d 00:10:22
Version          : 3

General Query Interval : 2s
Query Response Interval : 1.0s
Robust Count          : 2
-----

Number of mrouter: 2
=====
*A:ala-427#

```

mvr

| | |
|--------------------|---|
| Syntax | mvr |
| Context | show>service>id>igmp-snooping |
| Description | This command displays Multicast VPLS Registration (MVR) information. |
| Output | Show All Service-ID — The following table describes the show all service-id command output fields: |

| Label | Description |
|-----------------|-----------------------|
| MVR Admin State | Administrative state. |

| Label | Description (Continued) |
|------------|---|
| MVR Policy | Policy name. |
| Svc ID | The service identifier. |
| Sap/Sdp Id | Displays the SAP and SDP IDs of the service ID. |
| Oper State | Displays the operational state of the SAP and SDP IDs of the svcid. |
| Mrtr Port | Specifies if the port is a multicast router port. |
| From VPLS | Specifies from which VPLS the multicast streams corresponding to the groups learned via this SAP will be copied. If local, it is from its own VPLS. |
| Num Groups | Specifies the number of groups learned via this local SAP. |

Sample Output

```
*A:ALA-1>show>service>id>snooping# mvr
=====
IGMP Snooping Multicast VPLS Registration info for service 10
=====
IGMP Snooping Admin State : Up
MVR Admin State           : Up
MVR Policy                 : mvr-policy
-----
Local SAPs/SDPs
-----
Svc Id      Sap/Sdp      Oper      From      Num Local
            Id           State     VPLS      Groups
-----
100         sap:1/1/10:10      Up        Local     100
100         sap:1/1/10:20      Up        Local     100
-----
MVR SAPs (from-vpls=10)
-----
Svc Id      Sap/Sdp      Oper      From      Num MVR
            Id           State     VPLS      Groups
-----
20          sap:1/1/4:100      Up        10        100
30          sap:1/1/31:10.10  Up        10        100
=====
*A:ALA-1>show>service>id>snooping#
```

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

group *grp-ip-address* — Displays the IGMP snooping port database for a specific multicast group address.

sap *sap-id* — Displays the IGMP snooping port database for a specific SAP. See [Common CLI Command Descriptions on page 1319](#) for command syntax.

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.

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

group *grp-address* — Displays IGMP snooping statistics matching the specified group address.

source *ip-address* — Displays IGMP snooping statistics matching one particular source within the multicast group.

vxlan vtep *ip-address vni* <1..16777215> — Displays the IGMP snooping entries associated with a specific VXLAN binding, given by the VXLAN Termination Endpoint (VTEP) and VXLAN Network Identifier (VNI).

vni *vni* — The VXLAN Network Identifier (VNI) for which to display information.

Values 1 — 16777215

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 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. |

| Label | Description |
|---------------------|--|
| 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. |

Sample Output

```
*A:ALA-1>show>service>id>snooping# port-db sap 1/1/2
=====
IGMP Snooping SAP 1/1/2 Port-DB for service 10
=====
Group Address      Mode      Type      From-VPLS  Up Time      Expires  Num Src
-----
225.0.0.1          include   dynamic   Local      0d 00:04:44  0s       2
-----
Number of groups: 1
=====
*A:ALA-1>show>service>id>snooping#
```

```

*A:ALA-1>show>service>id>snooping# port-db sap 1/1/2 detail
=====
IGMP Snooping SAP 1/1/2 Port-DB for service 10
=====
IGMP Group 225.0.0.1
-----
Mode           : include           Type           : dynamic
Up Time        : 0d 00:04:57       Expires        : 0s
Compat Mode    : IGMP Version 3
V1 Host Expires : 0s               V2 Host Expires : 0s
-----
Source Address  Up Time      Expires  Type      Fwd/Blk
-----
1.1.1.1         0d 00:04:57  20s     dynamic   Fwd
1.1.1.2         0d 00:04:57  20s     dynamic   Fwd
-----
Number of groups: 1
=====
*A:ALA-1>show>service>id>snooping#

```

proxy-db

| | |
|--------------------|--|
| Syntax | proxy-db [detail] proxy-db group <i>grp-address</i> |
| Context | show>service>id>igmp-snooping |
| Description | This command displays information on the IGMP snooping proxy reporting database for the VPLS service. |
| Parameters | group <i>grp-ip-address</i> — 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. |

| Label | Description (Continued) |
|------------------|--|
| 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. |

Sample Output

```
*A:ALA-1>show>service>id>snooping# proxy-db
=====
IGMP Snooping Proxy-reporting DB for service 10
=====
Group Address      Mode      Up Time      Num Sources
-----
225.0.0.1          include   0d 00:05:40   2
-----
Number of groups: 1
=====
*A:ALA-1>show>service>id>snooping#

*A:ALA-1>show>service>id>snooping# proxy-db detail
=====
IGMP Snooping Proxy-reporting DB for service 10
-----
IGMP Group 225.0.0.1
-----
Up Time : 0d 00:05:54          Mode : include
-----
Source Address  Up Time
-----
1.1.1.1         0d 00:05:54
1.1.1.2         0d 00:05:54
-----
Number of groups: 1
=====
*A:ALA-1>show>service>id>snooping#
```

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. |
| 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. |

Sample Output

```
*A:ALA-1>show>service>id>snooping# querier
=====
IGMP Snooping Querier info for service 10
=====
Sap Id           : 1/1/1
IP Address       : 10.10.10.1
Expires          : 6s
Up Time          : 0d 00:56:50
Version          : 3

General Query Interval : 5s
Query Response Interval : 2.0s
Robust Count          : 2
=====
*A:ALA-1>show>service>id>snooping#
```

static

| | |
|--------------------|---|
| Syntax | static [sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i>] |
| Context | show>service>id>igmp-snooping |
| Description | This command displays information on static IGMP snooping source groups for the VPLS service. |
| Parameters | <p>sap <i>sap-id</i> — Displays static IGMP snooping source groups for a specific SAP. See Common CLI Command Descriptions on page 1319 for command syntax.</p> <p>sdp <i>sdp-id</i> — Displays the IGMP snooping source groups 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> |
| 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 | sap <i>sap-id</i> — Displays IGMP snooping statistics for a specific SAP. See Common CLI Command Descriptions on page 1319 for command syntax. sdp <i>sdp-id</i> — Displays the IGMP snooping statistics for a specific spoke or mesh SDP. Values 1 — 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to display information. Default For mesh SDPs only, all VC IDs. Values 1 — 4294967295 vxlan vtep <i>ip-address vni</i> <1..16777215> — Displays the IGMP snooping entries associated with a specific VXLAN binding, given by the VXLAN Termination Endpoint (VTEP) and VXLAN Network Identifier (VNI). vni <i>vni</i> — The VXLAN Network Identifier (VNI) for which to display information. Values 1 — 16777215 |

Sample Output

```
*A:ALA-1>show>service>id>snooping# statistics
=====
IGMP Snooping Statistics for service 1
=====
Message Type           Received      Transmitted   Forwarded
-----
General Queries        4             0             4
Group Queries          0             0             0
Group-Source Queries   0             0             0
V1 Reports             0             0             0
V2 Reports             0             0             0
V3 Reports             0             0             0
V2 Leaves              0             0             0
Unknown Type           0             N/A           0
-----
Drop Statistics
-----
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0

MVR From VPLS Cfg Drops : 0
```

```
MVR To SAP Cfg Drops      : 0
=====
*A:ALA-1>show>service>id>snooping#
```

egress-replication

| | |
|--------------------|--|
| Syntax | egress-replication |
| Context | show |
| Description | This command enables the context to display egress flooding information for a VPLS service context on a given MDA. A VPLS service context supports both Layer 2 and Layer 3 flooding modes. The Layer 2 flooding mode is used for broadcast, Layer 2 multicast and unknown destination MAC addressed packets. All available interfaces (SAP, spoke SDP and mesh-SDP) that reside on an egress forwarding complex are included in the egress list except for SAPs that are defined in a residential split horizon group (Layer 2 flooding is not permitted on residential SAPs). The Layer 3 flooding mode is used for VPLS interfaces participating in IGMP snooping and is represented by an IP multicast [s,g] record. |

vpls

| | |
|--------------------|---|
| Syntax | vpls <i>vpls-service-id</i> mda <i>card/slot</i> vpls <i>vpls-service-id</i> mda <i>card/slot</i> [igmp-record <i>group ip-address</i> { source <i>ip-address</i> starg }] |
| Context | show>egress-replication |
| Description | <p>The vpls <i>vpls-service-id</i> mda <i>slot/mda</i> command displays the flooding list used by the Layer 2 flooding mode for the VPLS service on the specified MDA. The Layer 2 flooding list is limited to SAPs, spoke SDP and mesh-SDP bindings that exist on the egress forwarding complex serviced by the specified MDA. For the 10G IOM, two MDAs share the same egress forwarding plane. In this case the Layer 2 flooding list will contain destinations for both MDAs (if entries exist). The only VPLS interfaces that will not be included in the list are residential SAPs because Layer 2 replication is not permitted to a residential SAP. A packet processed by the egress Layer 2 flooding list may not be replicated to each destination. A packet will not be replicated to an interface on the Layer 2 flooding list because of the following:</p> <ul style="list-style-type: none"> • The ingress interface split horizon group is the same as the egress interface (residential bridging rule). • The egress interface is down or blocking. • The packet matches a discard event while processing that destination interface. • An egress MTU violation occurs for the destination interface. <p>Destination SAPs in the list may be displayed in a chain context representing common replication behavior. All SAPs in a single chain are processed a single time through the egress forwarding plane. If a discard decision is made for the first SAP in the chain, no replication processing is done for any of the chain members. If the forwarding plane decides to replicate the first SAP in the chain, it will replicate to all SAPs in the chain.</p> |

The **vpls vpls-service-id mda card/slot igmp-record grp-address {source source-ip-address | starg}** command displays the IGMP record based flooding list for the *vpls-service-id* on the specified MDA. Unlike the Layer 2 flooding list for the VPLS context, an IGMP record list may contain interfaces from other VPLS contexts due to MVR (Multicast VPLS Registration) events on the individual VPLS interfaces. VPLS interfaces in other VPLS contexts become associated with the specified vpls-service-id based on the MVR from-vpls definition. Another difference between the VPLS Layer 2 flooding list and IGMP lists is that many IGMP lists may exist (each associated with a different [s,g] record) and the lists may contain residential SAPs. The SAP chaining and replication behavior is similar to the VPLS Layer 2 flooding list.

IP multicast packets ingressing the vpls-service-id must match either a [*g] or [s,g] record to be associated with the record's egress IP multicast IGMP flooding list. A [*g] record will match any ingress IP multicast packet destined to the class D destination IP address represented by "g". An [s,g] record will match any ingress IP multicast packet with a source IP address matching "s" and a destination IP address matching "g". In the case that a packet could match both a [*g] and [s,g] record, the [s,g] record takes precedence. Each [*g] and [s,g] record has its own IGMP flooding list. The list will only appear on an egress forwarding plane (MDA) when a member of the list (VPLS interface) exists on the forwarding plane.

Parameters

service-id — Displays information about the specified service ID or service name.

Values

service-id: 1 — 214748364

svc-name: A string up to 64 characters in length.

slot/mda — Specifies a chassis and MDA slot.

grp-ip-address — Specifies a multicast group address.

src-ip-address — Specifies a source IP address.

starg — Specifies a (*, G) record.

mRouter — Specifies the (*,*) record

ipv6 — Displays IPv6 information.

grp-ipv6-address — ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0..FFFF]H

d - [0..255]D

multicast group IPv6 address

src-ipv6-address — ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0..FFFF]H

d - [0..255]D

IGMP Commands

group

- Syntax** **group** [*grp-ip-address*]
- Context** show>router>igmp
- Description** This command displays the multicast group and (s, g) addresses. If no *grp-ip-address* parameters are specified then all IGMP group, (*, g) and (s, g) addresses are displayed.
- Parameters** *grp-ip-address* — Displays specific multicast group addresses.
- Output** **IGMP Group Output** — The following table describes the output fields for IGMP group information.

| Label | Description |
|-------------|---|
| IGMP Groups | Displays the IP multicast sources corresponding to the IP multicast groups which are statically configured. |
| Fwd List | Displays the list of interfaces in the forward list. |
| Blk List | Displays the list of interfaces in the block list. |

Sample Output

```
A:NYC# show router igmp group
=====
IGMP Groups
=====
(*,224.24.24.24)                               Up Time : 0d 05:21:38
    Fwd List  : nyc-vlc

(*,239.255.255.250)                             Up Time : 0d 05:21:38
    Fwd List  : nyc-vlc
-----
(*,G)/(S,G) Entries : 2
=====
A:NYC#

A:NYC# show router igmp group 224.24.24.24
=====
IGMP Groups
=====
(*,224.24.24.24)                               Up Time : 0d 05:23:23
    Fwd List  : nyc-vlc
-----
(*,G)/(S,G) Entries : 1
=====
A:NYC#
```

ssm-translate

- Syntax** **ssm-translate**
- Context** show>router>igmp
- Description** This command displays IGMP SSM translate configuration information.
- Output** **GMP Interface Output** — The following table provides IGMP field descriptions

| Label | Description |
|-----------------------|---|
| Group Range | Displays the address ranges of the multicast groups for which this router can be an RP. |
| Source | Displays the unicast address that sends data on an interface. |
| SSM Translate Entries | Displays the total number of SSM translate entries. |

```
A:ALA-48>config>router>igmp# show router igmp ssm-translate
=====
IGMP SSM Translate Entries
=====
Group Range                               Source
-----
<224.0.1.0 - 224.0.1.255>                 1.1.1.1
<225.1.0.0 - 225.240.3.57>               2.2.2.2
<239.255.255.0 - 239.255.255.255>       3.3.3.3
-----
SSM Translate Entries : 3
=====
A:ALA-48>config>router>igmp#
```

interface

- Syntax** **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-address*] [**detail**]
- Context** show>router>igmp
- Description** This command displays IGMP interface information.
- Parameters**
- ip-int-name* — Only displays the information associated with the specified IP interface name.
 - ip-address* — Only displays the information associated with the specified IP address.
 - group** *grp-address* — Only displays IP multicast group address for which this entry contains information.
 - detail** — Displays detailed IP interface information along with the source group information learned on that interface.
- Output** **IGMP Interface Output** — The following table provides IGMP field descriptions.

| Label | Description |
|---------------------------------------|--|
| Interface | Specifies the interfaces that participates in the IGMP protocol. |
| Adm Admin Status | Displays the administrative state for the IGMP protocol on this interface. |
| Oper Oper Status | Displays the current operational state of IGMP protocol on the interface. |
| Querier | Displays the address of the IGMP querier on the IP subnet to which the interface is attached. |
| Querier Up Time | Displays the time since the querier was last elected as querier. |
| Querier Expiry Timer | Displays the time remaining before the querier ages out. If the querier is the local interface address, the value will be zero. |
| Cfg/Opr Version Admin/Oper version | Cfg — The configured version of IGMP running on this interface. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN. Opr . The operational version of IGMP running on this interface. If the cfg value is 3 but all of the routers in the local subnet of this interface use IGMP version v1 or v2, the operational version will be v1 or v2. |
| Num Groups | The number of multicast groups which have been learned by the router on the interface. |
| Policy | Specifies the policy that is to be applied on the interface. |
| Group Address | Specifies the IP multicast group address for which this entry contains information. |
| Up Time | Specifies the time since this source group entry got created. |
| Last Reporter | Specifies the IP address of the source of the last membership report received for this IP Multicast group address on this interface. If no membership report has been received, this object has the value 0.0.0.0. |
| Mode | The mode is based on the type of membership report(s) received on the interface for the group. In the 'include' mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In 'exclude' mode, reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter. |
| V1 Host Timer | The time remaining until the local router will assume that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface. |

| Label | Description (Continued) |
|---------------|---|
| V2 Host Timer | The time remaining until the local router will assume that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 Leave messages for this group that it receives on this interface. |
| Type | Indicates how this group entry was learned. If this group entry was learned by IGMP, it will be set to 'dynamic'. For statically configured groups, the value will be set to 'static'. |
| Compat Mode | Used in order for routers to be compatible with older version routers. IGMPv3 hosts MUST operate in version 1 and version 2 compatibility modes. IGMPv3 hosts MUST keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the Host Compatibility Mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of General Queries heard on that interface as well as the Older Version Querier Present timers for the interface. |

Sample Output

```
A:BA# show router igmp interface
=====
IGMP Interfaces
=====
Interface           Adm  Oper  Querier      Cfg/Opr Num  Policy
                   Version Groups
-----
IGMP_to_CE          Up   Up    11.1.1.1      1/1    3    igmppol
-----
Interfaces : 1
=====
A:BA#

A:BA# show router 100 igmp interface IGMP_to_CE
=====
IGMP Interface IGMP_to_CE
=====
Interface           Adm  Oper  Querier      Cfg/Opr Num  Policy
                   Version Groups
-----
IGMP_to_CE          Up   Up    11.1.1.1      1/1    3    igmppol
-----
Interfaces : 1
=====
A:BA#

A:BA# show router 100 igmp interface 11.1.1.1
=====
IGMP Interface 11.1.1.1
=====
```

Show, Clear, Debug Commands

```
Interface                Adm  Oper Querier          Cfg/Opr Num    Policy
                        Version Groups
-----
IGMP_to_CE              Up   Up   11.1.1.1          1/1    3      igmppol
-----
Interfaces : 1
=====
A:BA#

A:BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1
=====
IGMP Interface IGMP_to_CE
=====
Interface                Adm  Oper Querier          Cfg/Opr Num    Policy
                        Version Groups
-----
IGMP_to_CE              Up   Up   11.1.1.1          1/1    3      igmppol
-----
IGMP Group
-----
Group Address : 227.1.1.1          Up Time       : 0d 00:03:52
Interface      : IGMP_to_CE        Expires       : never
Last Reporter  : 0.0.0.0           Mode          : exclude
V1 Host Timer  : Not running        Type          : static
V2 Host Timer  : Not running        Compat Mode   : IGMP Version 3
-----
Interfaces : 1
=====

A:BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1 detail
=====
IGMP Interface IGMP_to_CE
=====
Interface      : IGMP_to_CE
Admin Status   : Up                Oper Status    : Up
Querier        : 11.1.1.1          Querier Up Time : 0d 00:04:01
Querier Expiry Time: N/A           Time for next query: 0d 00:13:42
Admin/Oper version : 1/1           Num Groups     : 3
Policy         : igmppol           Subnet Check    : Disabled
Max Groups Allowed : 16000         Max Groups Till Now: 3
MCAC Policy Name :                  MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit      MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0             MCAC Avail Mand BW : unlimited
MCAC In use Opnl BW: 0             MCAC Avail Opnl BW : unlimited
-----
IGMP Group
-----
Group Address : 227.1.1.1          Up Time       : 0d 00:04:02
Interface      : IGMP_to_CE        Expires       : never
Last Reporter  : 0.0.0.0           Mode          : exclude
V1 Host Timer  : Not running        Type          : static
V2 Host Timer  : Not running        Compat Mode   : IGMP Version 3
-----
Interfaces : 1
=====
A:BA#
```

static

- Syntax** **static** [*ip-int-name* | *ip-addr*]
- Context** show>router>igmp
- Description** This command displays static IGMP, (*, g) (s, g) information.
- Parameters** *ip-int-name* — Only displays the information associated with the specified IP interface name.
ip-addr — Only displays the information associated with the specified IP address.
- Output** **Static IGMP Output** — The following table provides static IGMP field descriptions

| Label | Description |
|-----------|---|
| Source | Displays entries which represents a source address from which receivers are interested/not interested in receiving multicast traffic. |
| Group | Displays the IP multicast group address for which this entry contains information. |
| Interface | Displays the interface name. |

Sample Output

```
A:BA# show router 100 igmp static
=====
IGMP Static Group Source
=====
Source          Group          Interface
-----
11.11.11.11     226.136.22.3   IGMP_to_CE
*               227.1.1.1      IGMP_to_CE
22.22.22.22     239.255.255.255 IGMP_to_CE
-----
Static (*,G)/(S,G) Entries : 3
=====
A:BA#
```

statistics

- Syntax** **statistics** [*ip-int-name* | *ip-address*]
- Context** show>router>igmp
- Description** This command displays IGMP statistics information.
- Parameters** *ip-int-name* — Only displays the information associated with the specified IP interface name.
ip-addr — Only displays the information associated with the specified IP address.

Output **IGMP Statistics Output** — The following table provides statistical IGMP field descriptions

| Label | Description |
|------------------------------|---|
| IGMP Interface Statistics | The section listing the IGMP statistics for a particular interface. |
| Message Type | <p>Queries — The number of IGMP general queries transmitted or received on this interface.</p> <p>Report — The total number of IGMP V1, V2, or V3 reports transmitted or received on this interface.</p> <p>Leaves — The total number of IGMP leaves transmitted on this interface.</p> |
| Received | Column that displays the total number of IGMP packets received on this interface. |
| Transmitted | Column that displays the total number of IGMP packets transmitted from this interface. |
| General Interface Statistics | The section listing the general IGMP statistics. |
| Bad Length | Displays the total number of IGMP packets with bad length received on this interface. |
| Bad Checksum | Displays the total number of IGMP packets with bad checksum received on this interface. |
| Unknown Type | Displays the total number of IGMP packets with unknown type received on this interface. |
| Bad Receive If | Displays the total number of IGMP packets incorrectly received on this interface. |
| Rx Non Local | Displays the total number of IGMP packets received from a non-local sender. |
| Rx Wrong Version | Displays the total number of IGMP packets with wrong versions received on this interface. |
| Policy Drops | Displays the number of times IGMP protocol instance matched the host IP address or group/source addresses in the import policy. |
| No Router Alert | Displays the total number of IGMPv3 packets received on this interface which did not have the router alert flag set. |

Sample Output

```
A:BA# show router 100 igmp statistics
=====
IGMP Interface Statistics
=====
Message Type      Received      Transmitted
-----
Queries           0             5
```

```

Report V1          0          0
Report V2          0          0
Report V3          0          0
Leaves             0          0
-----
General Interface Statistics
-----
Bad Length         : 0
Bad Checksum       : 0
Unknown Type       : 0
Bad Receive If     : 0
Rx Non Local       : 0
Rx Wrong Version   : 0
Policy Drops       : 0
No Router Alert    : 0
Rx Bad Encodings   : 0
Rx Pkt Drops       : 0
-----
Source Group Statistics
-----
(S,G)              : 2
(*,G)              : 1
=====
A:BA#

```

status

Syntax **status**

Context show>router>igmp

Description This command displays IGMP status information.
If IGMP is not enabled, the following message appears:

```

A:NYC# show router igmp status
MINOR: CLI IGMP is not configured.
A:NYC#

```

Output **IGMP Status Output** — The following table provides IGMP status field descriptions

| Label | Description |
|----------------------------|--|
| Admin State | Displays the administrative status of IGMP. |
| Oper State | Displays the current operating state of this IGMP protocol instance on this router. |
| Query Interval | The frequency at which IGMP query packets are transmitted. |
| Last Member Query Interval | The maximum response time inserted into Group-Specific Queries sent in response to Leave Group messages, and is also the amount of time between Group-Specific Query messages. |
| Query Response Interval | The maximum query response time advertised in IGMPv2 queries. |
| Robust Count | Displays the number of times the router will retry a query. |

Sample Output

```
A:BA# show router 100 igmp status
=====
IGMP Status
=====
Admin State           : Up
Oper State            : Up
Query Interval        : 1024
Last Member Query Interval : 1024
Query Response Interval : 1023
Robust Count          : 10
=====
A:BA
```

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 | Sample Output |

```
*A:DutA# show service id 1 bgp-evpn
=====
BGP EVPN Table
=====
MAC Advertisement   : Enabled           Unknown MAC Route   : Disabled
VXLAN Admin Status  : Enabled           Creation Origin      : manual
MAC Dup Detn Moves  : 5                 MAC Dup Detn Window : 3
MAC Dup Detn Retry  : 9                 Number of Dup MACs  : 1

-----
Detected Duplicate MAC Addresses           Time Detected
-----
00:12:12:12:12:00                        01/17/2014 16:01:02
-----
=====
```

dhcp

| | |
|--------------------|---|
| Syntax | dhcp |
| Context | show>service>id |
| Description | This command enables the context to display DHCP information for the specified service. |

lease-state

| | |
|--------------------|--|
| Syntax | lease-state [[sap <i>sap-id</i>] [sdp <i>sdp-id:vc-id</i>] [interface <i>interface-name</i>] [ip-address <i>ip-address</i>]] [detail] |
| Context | show>service>id>dhcp |
| Description | This command displays DHCP lease state related information. |
| Parameters | <p>sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1319 for command syntax.</p> <p><i>sdp-id</i> — The SDP identifier.</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>Values 1 — 4294967295</p> <p>interface <i>interface-name</i> — Displays information for the specified IP interface.</p> <p>ip-address <i>ip-address</i> — Displays information associated with the specified IP address.</p> <p>detail — Displays detailed information.</p> |

Sample Output

```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
=====
DHCP lease state table, service 13
=====
IP Address           Mac Address           Sap/Sdp Id           Remaining   Lease   MC
                    LifeTime             Origin              Stdbby
-----
13.13.40.1           00:00:00:00:00:13    1/1/1:13            00h00m58s   Radius
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#

A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
=====
DHCP lease states for service 13
=====
Service ID           : 13
IP Address           : 13.13.40.1
Mac Address          : 00:00:00:00:00:13
Interface            : ies-13-13.13.1.1
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h00m58s
Persistence Key      : N/A

Sub-Ident            : "TEST"
Sub-Profile-String    : "ADSL GO"
SLA-Profile-String    : "BE-Video"
Lease ANCP-String     : ""

Sub-Ident origin      : Radius
```

Show, Clear, Debug Commands

```
Strings origin      : Radius
Lease Info origin   : Radius

Ip-Netmask          : 255.255.0.0
Broadcast-Ip-Addr   : 13.13.255.255
Default-Router      : N/A
Primary-Dns         : 13.13.254.254
Secondary-Dns       : 13.13.254.253

ServerLeaseStart    : 12/24/2006 23:44:07
ServerLastRenew     : 12/24/2006 23:44:07
ServerLeaseEnd      : 12/24/2006 23:45:07
Session-Timeout     : 0d 00:01:00
DHCP Server Addr    : N/A

Persistent Relay Agent Information
  Circuit Id        : ancstb6_Dut-A|13|ies-13-13.13.1.1|0|13
  Remote Id         : stringtest
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#
```


Routed CO Output Example

```

A:ALA-_Dut-A# show service id 13 dhcp lease-state
=====
DHCP lease state table, service 13
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining      Lease      MC
                  LifeTime      Origin      Stdbby
-----
13.13.40.1      00:00:00:00:00:13  1/1/1:13      00h00m58s     Radius
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#

```

```

A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
=====
DHCP lease states for service 13
=====
Service ID      : 13
IP Address      : 13.13.40.1
Mac Address      : 00:00:00:00:00:13
Subscriber-interface : ies-13-13.13.1.1
Group-interface  : intf-13
SAP              : 1/1/1:13
Remaining Lifetime : 00h00m58s
Persistence Key   : N/A

Sub-Ident       : "TEST"
Sub-Profile-String : "ADSL GO"
SLA-Profile-String : "BE-Video"
Lease ANCP-String : ""

Sub-Ident origin : Radius
Strings origin   : Radius
Lease Info origin : Radius

Ip-Netmask       : 255.255.0.0
Broadcast-Ip-Addr : 13.13.255.255
Default-Router    : N/A
Primary-Dns       : 13.13.254.254
Secondary-Dns     : 13.13.254.253

ServerLeaseStart  : 12/24/2006 23:48:23
ServerLastRenew   : 12/24/2006 23:48:23
ServerLeaseEnd    : 12/24/2006 23:49:23
Session-Timeout   : 0d 00:01:00
DHCP Server Addr  : N/A

Persistent Relay Agent Information
  Circuit Id      : ancstb6_Dut-A|13|intf-13|0|13
  Remote Id      : stringtest
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#

```

Wholesaler/Retailer Output Example

Show, Clear, Debug Commands

```
A:ALA-_Dut-A# show service id 2000 dhcp lease-state detail
=====
DHCP lease states for service 2000
-----
Wholesaler 1000 Leases
-----
Service ID           : 1000
IP Address           : 13.13.1.254
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : whole-sub
Group-interface      : intf-13
Retailer             : 2000
Retailer If          : retail-sub
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h09m59s
Persistence Key       : N/A

Sub-Ident             : "TEST"
Sub-Profile-String    : "ADSL GO"
SLA-Profile-String    : "BE-Video"
Lease ANCP-String     : ""

Sub-Ident origin      : Retail DHCP
Strings origin        : Retail DHCP
Lease Info origin     : Retail DHCP

Ip-Netmask            : 255.255.0.0
Broadcast-Ip-Addr     : 13.13.255.255
Default-Router        : N/A
Primary-Dns           : N/A
Secondary-Dns         : N/A

ServerLeaseStart      : 12/25/2006 00:29:41
ServerLastRenew       : 12/25/2006 00:29:41
ServerLeaseEnd        : 12/25/2006 00:39:41
Session-Timeout       : 0d 00:10:00
DHCP Server Addr      : 10.232.237.2

Persistent Relay Agent Information
  Circuit Id          : 1/1/1:13
  Remote Id           : stringtest
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#
```

statistics

| | |
|--------------------|---|
| Syntax | statistics [sap <i>sap-id</i> statistics [sdp <i>sdp-id:vc-id</i> statistics [interface <i>interface-name</i>] |
| Context | show>service>id>dhcp |
| Description | Displays DHCP statistics information. |
| Parameters | sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1319 for command syntax. sdp-id — The SDP identifier. Values 1 — 17407 vc-id — The virtual circuit ID on the SDP ID for which to display information. Values 1 — 4294967295 interface <i>interface-name</i> — Displays information for the specified IP interface. |

summary

| | |
|--------------------|---|
| Syntax | summary |
| Context | show>service>id>dhcp |
| Description | Displays DHCP configuration summary information. |
| Output | Show DHCP Summary Output — The following table describes the output fields for DHCP summary. |

| Label | Description |
|----------------|--|
| Interface Name | Name of the router interface. |
| Arp Populate | Specifies whether or not ARP populate is enabled. |
| Used/Provided | <p>Used — The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the 'Provided' field.</p> <p>Provided — The lease-populate value that is configured for a specific interface.</p> |
| Info Option | Indicates whether Option 82 processing is enabled on the interface. |
| Admin State | Indicates the administrative state. |

Sample Output

```
A:ALA-49# show service id 88 dhcp summary
```

Show, Clear, Debug Commands

```
=====
DHCP Summary, service 88
=====
Interface Name      Arp      Used/      Info      Admin
  SapId/Sdp        Populate Provided      Option    State
-----
Sector A            No        0/0              Keep      Up
  sap:7/1/1.2.2              0/0
  sap:2/2/2:0                0/1
test                 No        0/0              Keep      Up
  sap:10/1/2:0              0/0
-----
Interfaces: 3
=====
A:ALA-49#
```

Show Multi-Chassis Endpoint Commands

endpoint

| | |
|--------------------|--|
| Syntax | endpoint [<i>endpoint-name</i>] |
| Context | show>service>id |
| Description | This command displays service endpoint information. |
| Parameters | <i>endpoint-name</i> — Specifies an endpoint name created in the config>service>vpls context. |

Sample Output

```
*A:Dut-B# show service id 1 endpoint
=====
Service 1 endpoints
=====
Endpoint name           : mcep-t1
Description             : (Not Specified)
Revert time             : 0
Act Hold Delay          : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
Block On Mesh Fail      : true
Multi-Chassis Endpoint  : 1
MC Endpoint Peer Addr   : 3.1.1.3
Psv Mode Active         : No
Tx Active               : 231:1
Tx Active Up Time       : 0d 00:06:57
Revert Time Count Down  : N/A
Tx Active Change Count  : 5
Last Tx Active Change   : 02/13/2009 22:08:33
-----
Members
-----
Spoke-sdp: 221:1 Prec:1                      Oper Status: Up
Spoke-sdp: 231:1 Prec:2                      Oper Status: Up
=====
*A:Dut-B#
```

etree

| | |
|--------------------|--|
| Syntax | etree |
| Context | show>service>id |
| Description | This command displays the same information shown in the show service ID base context, with the addition of the role of each object in the VPLS E-Tree service. |

The following labels identify the configuration of the SAPs and SDP bindings:

- (L) indicates leaf-ac
- (RL) indicates root-leaf-tag

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
=====
```

multi-chassis

| | |
|--------------------|--|
| Syntax | multi-chassis |
| Context | show>redundancy |
| Description | This command enables the context to display multi-chassis information. |

mc-endpoint

| | |
|--------------------|--|
| Syntax | mc-endpoint statistics mc-endpoint peer <i>[ip-address]</i> statistics mc-endpoint endpoint <i>[mcep-id]</i> statistics mc-endpoint peer <i>[ip-address]</i> |
| Context | show>redundancy>multi-chassis |
| Description | This command displays multi-chassis endpoint information. |
| Parameters | statistics — Displays the global statistics for the MC endpoint. peer <i>ip-address</i> — Specifies the IP address of multi-chassis end-point peer. endpoint <i>mcep-id</i> — Specifies the multi-chassis endpoint. Values 1 — 4294967295 |

Sample Output

```
*A:Dut-B# show redundancy multi-chassis mc-endpoint statistics
=====
Multi-Chassis Endpoint Global Statistics
=====
Packets Rx                               : 533
Packets Rx Keepalive                     : 522
Packets Rx Config                         : 3
Packets Rx Peer Config                   : 1
Packets Rx State                         : 7
Packets Dropped Keep-Alive Task          : 7
Packets Dropped Too Short                 : 0
Packets Dropped Verify Failed            : 0
Packets Dropped Tlv Invalid Size         : 0
Packets Dropped Out Of Seq               : 0
Packets Dropped Unknown Tlv              : 0
Packets Dropped Tlv Invalid MC-Endpoint Id : 0
Packets Dropped MD5                      : 0
Packets Dropped Unknown Peer             : 0
Packets Dropped MC Endpoint No Peer      : 0
Packets Tx                               : 26099
Packets Tx Keepalive                     : 8221
Packets Tx Config                         : 2
Packets Tx Peer Config                   : 17872
Packets Tx State                         : 4
Packets Tx Failed                        : 0
=====
*A:Dut-B#

*A:Dut-B# show redundancy multi-chassis mc-endpoint peer 3.1.1.3 statistics
=====
Multi-Chassis MC-Endpoint Statistics
=====
Peer Addr                               : 3.1.1.3
-----
Packets Rx                               : 597
Packets Rx Keepalive                     : 586
Packets Rx Config                         : 3
Packets Rx Peer Config                   : 1
```

Show, Clear, Debug Commands

```
Packets Rx State : 7
Packets Dropped State Disabled : 0
Packets Dropped Packets Too Short : 0
Packets Dropped Tlv Invalid Size : 0
Packets Dropped Tlv Invalid LagId : 0
Packets Dropped Out of Seq : 0
Packets Dropped Unknown Tlv : 0
Packets Dropped MD5 : 0
Packets Tx : 636
Packets Tx Keepalive : 600
Packets Tx Peer Config : 30
Packets Tx Failed : 0
Packets Dropped No Peer : 0
=====
*A:Dut-B#

*A:Dut-B# show redundancy multi-chassis mc-endpoint endpoint 1 statistics
=====
Multi-Chassis Endpoint Statistics
=====
MC-Endpoint Id 1
=====
Packets Rx Config : 3
Packets Rx State : 7
Packets Tx Config : 2
Packets Tx State : 4
Packets Tx Failed : 0
=====
Number of Entries 1
=====
*A:Dut-B#

*A:Dut-B# tools dump redundancy multi-chassis mc-endpoint peer 3.1.1.3
=====
MC Endpoint Peer Info
  peer addr : 3.1.1.3
  peer name : Dut-C
  peer name refs : 1
  src addr conf : Yes
  source addr : 2.1.1.2
  num of mcep : 1
  num of non-mcep : 0
  own sess num : 58ba0d39
  mc admin state : Up
  tlv own mc admin state : Up
  tlv peer mc admin state : Up
  reachable : Yes

  own sys priority : 50
  own sys id : 00:03:fa:72:c3:c0
  peer sys priority : 21
  peer sys id : 00:03:fa:c6:31:f8
  master : No

  conf boot timer : 300
  boot timer active : No
  conf ka intv : 10
  conf hold on num of fail : 3
  tlv own ka intv : 10
  tlv peer ka intv : 10
```



```

ka timeout tmr active      : Yes
ka timeout tmr intvl      : 20
ka timeout tmr time left   : 4
peer ka intv              : 10
mc peer timed out         : No

initial peer conf rx       : Yes
peer-mc disabled          : No
initial peer conf sync     : Yes
peer conf sync            : Yes

own passive mode          : Disable
peer passive mode         : No

retransmit pending        : No
non-mcep retransmit pending : No
retransmit intvl          : 5
last tx time              : 1437130
last rx time              : 1437156

own bfd                   : Enable
peer bfd                  : Enable
bfd vrtr if              : 2
bfd handle                : 1
bfd state                 : 3
bfd code                  : 0
=====
*A:Dut-B#

*A:Dut-B#  tools dump service mc-endpoint 1
=====
MC Endpoint Info
mc-endpoint id           : 1
endpoint                 : mcep-t1
service                  : 1
peer ref type            : peer-name
peer                     : Dut-C
mc sel logic             : peer selected active
selection master         : No
retransmit pending       : No
initial config sync      : Yes
config sync              : Yes
peer not mcep            : No
peer acked non-mcep      : No
config mismatch          : No
initial state rx         : Yes
initial state sync       : Yes
state sync               : Yes
can aggregate            : Yes
sel peer active          : No
peer sel active          : Yes
passive mode active      : No
own eligible force       : No
own eligible double active : Yes
own eligible pw status bits : 0
own eligible precedence  : 2
own eligible conf chg    : No
own eligible revert wait : No
peer eligible force      : No
peer eligible double active : Yes

```

Show, Clear, Debug Commands

```
peer eligible pw status bits : 0
peer eligible precedence     : 3
peer eligible conf chg       : No
peer eligible revert wait    : No
=====
*A:Dut-B#

*A:Dut-B# tools perform service id 1 endpoint mcep-t1 force-switchover 221:1
*A:Dut-B>show#
*A:Dut-B# show service id 1 endpoint
=====
Service 1 endpoints
=====
Endpoint name           : mcep-t1
Description              : (Not Specified)
Revert time              : 0
Act Hold Delay           : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
Block On Mesh Fail       : true
Multi-Chassis Endpoint   : 1
MC Endpoint Peer Addr    : 3.1.1.3
Psv Mode Active          : No
Tx Active                 : 221:1(forced)
Tx Active Up Time        : 0d 00:00:17
Revert Time Count Down   : N/A
Tx Active Change Count    : 6
Last Tx Active Change    : 02/14/2009 00:17:32
-----
Members
-----
Spoke-sdp: 221:1 Prec:1                               Oper Status: Up
Spoke-sdp: 231:1 Prec:2                               Oper Status: Up
=====
*A:Dut-B#
```

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. |

arp-host

| | |
|--------------------|---|
| Syntax | arp-host arp-host { mac <i>ieee-address</i> sap <i>sap-id</i> ip-address <i>ip-address</i> [/ <i>mask</i>] } arp-host [port <i>port-id</i>] [inter-dest-id <i>intermediate-destination-id</i> no-inter-dest-id] arp-host statistics [sap <i>sap-id</i> interface <i>interface-name</i>] |
| Context | clear>service>id |
| Description | This command clears ARP host data. |

authentication

| | |
|--------------------|---|
| Syntax | authentication |
| Context | clear>service>id |
| Description | This command enables the context to clear session authentication information. |

capture-sap

| | |
|--------------------|--|
| Syntax | capture-sap <i>sap-id</i> [<i>trigger</i>] |
| Context | clear>service>id |
| Description | This command clears the statistics for a particular capture SAP. |

cem

| | |
|--------------------|--|
| Syntax | cem |
| Context | clear>service>id |
| Description | This command clears CEM statistics for this service. |

statistics

| | |
|--------------------|--|
| Syntax | statistics |
| Context | clear>service>stats clear>service>id>authentication |
| Description | This command clears session statistics for this service. |

fdb

| | | | | | | | | | | | | | | | | | | | |
|--------------------|---|---------------|----------------|---------------|-----------------------|---------------|-----------|--|--|--------------|----------------|--|---------------------|---------------|-----------|--|--|--------------|----------------|
| Syntax | fdb { all mac <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> <i>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 1319 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. <i>sdp-id</i> — The SDP ID for which to clear associated FDB entries. <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to clear associated FDB entries. <table><tr><td>Values</td><td><i>sdp-id[:vc-id]</i></td><td><i>sdp-id</i></td><td>1 — 17407</td></tr><tr><td></td><td></td><td><i>vc-id</i></td><td>1 — 4294967295</td></tr><tr><td></td><td><i>sdp-id:vc-id</i></td><td><i>sdp-id</i></td><td>1 — 17407</td></tr><tr><td></td><td></td><td><i>vc-id</i></td><td>1 — 4294967295</td></tr></table> | | | Values | <i>sdp-id[:vc-id]</i> | <i>sdp-id</i> | 1 — 17407 | | | <i>vc-id</i> | 1 — 4294967295 | | <i>sdp-id:vc-id</i> | <i>sdp-id</i> | 1 — 17407 | | | <i>vc-id</i> | 1 — 4294967295 |
| Values | <i>sdp-id[:vc-id]</i> | <i>sdp-id</i> | 1 — 17407 | | | | | | | | | | | | | | | | |
| | | <i>vc-id</i> | 1 — 4294967295 | | | | | | | | | | | | | | | | |
| | <i>sdp-id:vc-id</i> | <i>sdp-id</i> | 1 — 17407 | | | | | | | | | | | | | | | | |
| | | <i>vc-id</i> | 1 — 4294967295 | | | | | | | | | | | | | | | | |

mld-snooping

| | |
|--------------------|--|
| Syntax | mld-snooping |
| Context | clear>service>id |
| Description | This command enables the context to clear MLD snooping-related data. |

port-db

| | |
|--------------------|--|
| Syntax | port-db sap <i>sap-id</i> [group <i>grp-ipv6-address</i>] port-db sap <i>sap-id</i> group <i>grp-ipv6-address</i> source <i>src-ipv6-address</i> port-db sdp <i>sdp-id:vc-id</i> [group <i>grp-ipv6-address</i>] port-db sdp <i>sdp-id:vc-id</i> group <i>grp-ipv6-address</i> source <i>src-ipv6-address</i> |
| Context | clear>service>id>mld-snooping |
| Description | This command clears MLD snooping port-db group data. |

querier

| | |
|--------------------|---|
| Syntax | querier |
| Context | clear>service>id>mld-snooping |
| Description | This command clears MLD snooping querier information. |

statistics

| | |
|--------------------|---|
| Syntax | statistics all statistics sap <i>sap-id</i> statistics sdp <i>sdp-id:vc-id</i> |
| Context | clear>service>id>mld-snooping |
| Description | This command clears MLD snooping statistics. |

msap

| | |
|--------------------|---|
| Syntax | msap <i>msap-id</i> |
| Context | clear>service>id |
| Description | This command clears the managed SAP (MSAP). |
| Parameters | <i>msap-id</i> — Specifies the MSAP ID. |

| | | |
|---------------|-------|--|
| Values | dot1q | <i>port-id</i> <i>lag-id</i> :qtag1 |
| | qinq | <i>port-id</i> <i>lag-id</i> ::qtag1.qtag2 |
| | qtag1 | 0 — 4094 |
| | qtag2 | 0 — 4094 |

mesh-sdp

| | | | | | | | |
|--------------------|--|---------------|-----------|----------------|---------------------------|---------------|----------------|
| Syntax | mesh-sdp <i>sdp-id</i> [: <i>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. <table><tr><td>Values</td><td>1 — 17407</td></tr></table> <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset. <table><tr><td>Default</td><td>All VC IDs on the SDP ID.</td></tr><tr><td>Values</td><td>1 — 4294967295</td></tr></table> | Values | 1 — 17407 | Default | All VC IDs on the SDP ID. | Values | 1 — 4294967295 |
| Values | 1 — 17407 | | | | | | |
| 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.

sap

| | |
|--------------------|--|
| Syntax | sap <i>sap-id</i> { all cem counters l2pt stp mrp } |
| Context | clear>service>statistics |
| Description | This command clears statistics for the SAP bound to the service. |
| Parameters | <i>sap-id</i> — See Common CLI Command Descriptions on page 1319 for command syntax. all — Clears all queue statistics and STP statistics associated with the SAP. cem — Clears all CEM statistics associated with the SAP. counters — Clears all queue statistics associated with the SAP. l2pt — Clears all L2PT statistics associated with the SAP. stp — Clears all STP statistics associated with the SAP. mrp — Clears all MRP statistics associated with the SAP. |

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 | <i>sdp-id</i> — The SDP ID for which to clear statistics. Values 1 — 17407 keep-alive — Clears the keep-alive history associated with this SDP ID. |

counters

| | |
|--------------------|--|
| Syntax | counters |
| Context | clear>service>statistics>id |
| Description | This command clears all traffic queue counters associated with the service ID. |

l2pt

| | |
|--------------------|---|
| Syntax | l2pt |
| Context | clear>service>statistics>id |
| Description | This command clears the l2pt statistics for this service. |

mesh-sdp

| | |
|--------------------|--|
| Syntax | mesh-sdp <i>sdp-id[:vc-id]</i> { all counters stp mrp } |
| Context | clear>service>statistics>id |
| Description | This command clears the statistics for a particular mesh SDP bind. |
| Parameters | <i>sdp-id[:vc-id]</i> — sdp-id - [1..17407] vc-id - [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. mrp — Clears all MRP statistics associated with the SDP. |

mrp

| | |
|--------------------|--|
| Syntax | mrp |
| Context | clear>service>statistics>id |
| Description | This command clears all MRP statistics for the service ID. |

pip

| | |
|--------------------|--|
| Syntax | pip |
| Context | clear>service>statistics>id |
| Description | This command clears the Provider Internal Port statistics for this service |

spoke-sdp

| | |
|--------------------|--|
| Syntax | spoke-sdp <i>sdp-id[: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.

mrp — Clears all MRP 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> spoke-sdp <i>sdp-id[:vc-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 or spoke SDP. |
| Parameters | <p>all — Clears all detected protocol statistics.</p> <p><i>sap-id</i> — Clears the specified lease state SAP information. See Common CLI Command Descriptions on page 1319 for command syntax.</p> <p><i>sdp-id</i> — The SDP ID to be cleared.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID to be cleared.</p> <p>Values 1 — 4294967295</p> |

lease-state

| | |
|--------------------|--|
| Syntax | lease-state [no-dhcp-release] lease-state ip-address <i>ip-address</i> [no-dhcp-release] lease-state mac <i>ieee-address</i> no-dhcp-release lease-state sap <i>sap-id</i> [no-dhcp-release] lease-state sdp <i>sdp-id:vc-id</i> [no-dhcp-release] |
| Context | clear>service>id>dhcp |
| Description | This command clears DHCP lease state information. |
| Parameters | no-dhcp-release — Specifies that the node will clear the state without sending the DHCP release message. |

ip-address *ip-address* — Clears the DHCP IP address lease state information. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mac *ieee-address* — Clears DHCP MAC address lease state information. The 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap *sap-id* — Clears DHCP SAP lease state information. See [Common CLI Command Descriptions on page 1319](#) for command syntax.

sdp-id — The SDP ID to be cleared.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID to be cleared.

Values 1 — 4294967295

statistics

| | |
|--------------------|--|
| Syntax | statistics [sap <i>sap-id</i> sdp [<i>sdp-id</i> [: <i>vc-id</i>]] interface [<i>ip-address</i> <i>ip-int-name</i>]] |
| Context | clear>service>id>dhcp |
| Description | Clears DHCP statistics for this service. |
| Parameters | <p><i>sap-id</i> — Clears the specified SAP statistics. See Common CLI Command Descriptions on page 1319 for command syntax.</p> <p><i>sdp-id</i> — The SDP ID to be cleared.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID to be cleared.</p> <p>Values 1 — 4294967295</p> <p>interface <i>ip-int-name</i> — Clears the statistics for the IP interface with the specified name.</p> <p>interface <i>ip-addr</i> — Clears the statistics for the IP interface with the specified IP address.</p> |

statistics

| | |
|--------------------|--|
| Syntax | statistics { all sap <i>sap-id</i> sdp <i>sdp-id</i> [: <i>vc-id</i>]}] |
| Context | clear>service>id>igmp-snooping |
| Description | Clears IGMP snooping statistics for the VPLS service. |
| Parameters | sap <i>sap-id</i> — Clears the IGMP snooping information on the specified SAP. See Common CLI Command Descriptions on page 1319 for command syntax. |

sdp *sdp-id* — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to clear statistics.

Default For mesh SDPs only, all VC IDs.

Values 1 — 4294967295

port-db

Syntax **port-db** [**sap** *sap-id*] [**group** *grp-address* [**source** *ip-address*]]
port-db **sdp** *sdp-id:vc-id* [**group** *grp-address* [**source** *ip-address*]]

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping port database for the VPLS service.

Parameters **sap** *sap-id* — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See [Common CLI Command Descriptions on page 1319](#) for command syntax.

sdp-id — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to clear information.

Default For mesh SDPs only, all VC IDs.

Values 1 — 4294967295

group *grp-address* — Clears IGMP snooping statistics matching the specified group address.

source *ip-address* — Clears IGMP snooping statistics matching the specified particular source.

querier

Syntax **querier**

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping queriers for the VPLS service.

mfib

Syntax **mfib**

Context clear>service>id>

Description This command enables the context to clear multicast FIB info for the VPLS service.

statistics

| | |
|--------------------|--|
| Syntax | statistics {all ip mac} statistics group <i>grp-address</i> statistics [all sap <i>sap-id</i>] |
| Context | clear>service>id>mfib |
| Description | This command clears multicast FIB statistics for the VPLS service. |
| Parameters | <i>grp-address</i> — Specifies an IGMP multicast group address that receives data on an interface. all — Clears all statistics for the service ID. sap <i>sap-id</i> — Clears statistics for the specified SAP ID. |

statistics

| | | | | | | | | | |
|--------------------|---|----------------|--|---------------|----------------|-----------|--|---------------|----------------|
| Syntax | statistics { all sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i> } | | | | | | | | |
| Context | clear>service>id>snooping | | | | | | | | |
| Description | This command clears IGMP snooping statistics. | | | | | | | | |
| Parameters | all — Clears all statistics for the service ID. sap <i>sap-id</i> — Clears statistics for the specified SAP ID. sdp <i>sdp-id:vc-id</i> — <table><tr><td>Values</td><td><i>sdp-id:</i></td><td>1 — 17407</td></tr><tr><td></td><td><i>vc-id:</i></td><td>1 — 4294967295</td></tr></table> | | | Values | <i>sdp-id:</i> | 1 — 17407 | | <i>vc-id:</i> | 1 — 4294967295 |
| Values | <i>sdp-id:</i> | 1 — 17407 | | | | | | | |
| | <i>vc-id:</i> | 1 — 4294967295 | | | | | | | |

dhcp

| | |
|--------------------|--|
| Syntax | dhcp |
| Context | clear>router |
| Description | This command enables the context to clear and reset DHCP entities. |

statistics

| | |
|--------------------|---|
| Syntax | statistics [interface <i>ip-int-name</i> <i>ip-address</i>] |
| Context | clear>router>dhcp |
| Description | Clears DHCP statistics. interface <i>ip-int-name</i> — Clears the statistics for the IP interface with the specified name. interface <i>ip-addr</i> — Clears the statistics for the IP interface with the specified IP address. |

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. |

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 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 | <p><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.</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>Values 1 — 4294967295</p> |

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 |

Description This command enables debugging of the applicant state machine.
The **no** form of the command disables debugging of the applicant state machine.

leave-all-sm

Syntax **[no] leave-all-sm**

Context debug>service>id>mrp

Description This command enables debugging of the leave all state machine.
The **no** form of the command disables debugging of the leave all state machine.

mmrp-mac

Syntax **[no] mmrp-mac** *ieee-address*

Context debug>service>id>mrp

Description This command filters debug events and only shows events related to the MAC address specified.
The **no** form of the command removes the debug filter.

Parameters *ieee-address* — xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all 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 <i>sap-id</i> |
| 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 1319 for command syntax. |

sdp

| | |
|--------------------|--|
| Syntax | [no] sdp <i>sdp-id:vc-id</i> |
| 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. |

host-connectivity-verify

| | |
|--------------------|--|
| Syntax | [no] host-connectivity-verify |
| Context | debug>service>id |
| Description | This command enables Subscriber Host Connectivity Verification (SHCV) debugging. The no form of the command disables the SHCV debugging. |

ip

| | |
|--------------------|--|
| Syntax | [no] ip ip-address |
| Context | debug>service>id>host-connectivity-verify |
| Description | This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular IP address. |
| Parameters | <i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> 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). |

mac

| | |
|--------------------|--|
| Syntax | [no] mac <i>ieee-address</i> |
| Context | debug>service>id>host-connectivity-verify |
| Description | This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular MAC address. |
| Parameters | <i>mac-address</i> — 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. |

proxy-arp

| | |
|--------------------|--|
| Syntax | proxy-arp [mac [<i>ieee-address</i>]] [ip [<i>ipaddr</i>] all]] |
| Context | debug>service>id |
| Description | This command 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. |

proxy-nd

| | |
|--------------------|---|
| Syntax | proxy-nd [mac [<i>ieee-address</i>]] [ip [<i>ipaddr</i>] all]] |
| Context | debug>service>id |
| Description | This command 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 <i>sap-id</i> |
| Context | debug>service>id>host-connectivity-verify |
| Description | This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular SAP. |
| Parameters | <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1319 for command syntax. |

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. |

stp

| | |
|--------------------|---|
| Syntax | stp |
| Context | debug>service>id |
| Description | This command enables the context for debugging STP. |

all-events

| | |
|--------------------|--|
| Syntax | all-events |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for all events. |

bpdu

| | |
|--------------------|--|
| Syntax | [no] bpdu |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for received and transmitted BPDUs. |

core-connectivity

| | |
|--------------------|---|
| Syntax | [no] core-connectivity |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for core connectivity. |

exception

| | |
|--------------------|--|
| Syntax | [no] exception |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for exceptions. |

fsm-state-changes

| | |
|--------------------|---|
| Syntax | [no] fsm-state-changes |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for FSM state changes. |

fsm-timers

| | |
|--------------------|---|
| Syntax | [no] fsm-timers |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for FSM timer changes. |

port-role

| | |
|--------------------|---|
| Syntax | [no] port-role |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for changes in port roles. |

port-state

| | |
|--------------------|---|
| Syntax | [no] port-state |
| Context | debug>service>id>stp |
| Description | This command enables STP debugging for port states. |

sap

| | |
|----------------|-------------------------------|
| Syntax | [no] sap <i>sap-id</i> |
| Context | debug>service>id>stp |

| | |
|--------------------|--|
| Description | This command enables STP debugging for a specific SAP. |
| Parameters | <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 1319 for command syntax. |

sdp

| | |
|--------------------|--|
| Syntax | [no] sdp <i>sdp-id:vc-id</i> |
| Context | debug>service>stp |
| Description | This command enables STP debugging for a specific SDP. |

interface

| | |
|--------------------|--|
| Syntax | [no] interface [<i>ip-int-name</i> <i>ip-address</i>] |
| Context | debug>router>igmp |
| Description | This command enables debugging on the IGMP interface. |
| Parameters | <i>ip-int-name</i> — Only displays the information associated with the specified IP interface name. <i>ip-address</i> — Only displays the information associated with the specified IP address. |

mcs

| | |
|--------------------|---|
| Syntax | [no] mcs [<i>ip-int-name</i>] |
| Context | debug>router>igmp |
| Description | This command enables debugging for IGMP MCS. |
| Parameters | <i>ip-int-name</i> — Only displays the information associated with the specified IP interface name. |

misc

| | |
|--------------------|--|
| Syntax | [no] misc |
| Context | debug>router>igmp |
| Description | This command enables debugging for IGMP miscellaneous. |

packet

| | |
|--------------------|--|
| Syntax | [no] packet [<i>query</i> <i>v1-report</i> <i>v2-report</i> <i>v3-report</i> <i>v2-leave</i>] [<i>ip-int-name</i> <i>ip-address</i>] |
| Context | debug>router>igmp |
| Description | This command enables debugging for IGMP packets. |
| Parameters | <i>query v1/v2/v3-report, v2-leave</i> — Select the type of packet to debug. <i>ip-int-name</i> — Only displays the information associated with the specified IP interface name. <i>ip-address</i> — Only displays the information associated with the specified IP address. |

provider-tunnels

| | |
|--------------------|---|
| Syntax | provider-tunnels type |
| Context | tools>dump>service>vpls |
| Description | This command dumps the inclusive provider tunnels based on type. |
| Output | <pre>*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</pre> |

```

-----

=====

VPLS 1001 Inclusive Provider Tunnels Terminating

=====

ipmsi (RSVP)                                P2MP-ID  Tunl-ID  Ext-Tunl-ID
-----
                                1001      61440   10.20.1.1
                                1001      64944   10.20.1.2
-----

*A:Dut-C>tools# dump service vpls 1001 provider-tunnels type terminating
=====
VPLS 1001 Inclusive Provider Tunnels Terminating

=====
ipmsi (RSVP)                                P2MP-ID  Tunl-ID  Ext-Tunl-ID
-----
                                1001      61440   10.20.1.1
                                1001      64944   10.20.1.2
-----

*A:Dut-C>tools# dump service vpls 1001 provider-tunnels type originating
=====
VPLS 1001 Inclusive Provider Tunnels Originating

=====
ipmsi (RSVP)                                P2MP-ID  Tunl-ID  Ext-Tunl-ID
-----

ipmsi-1001-73728                          1001      61440   10.20.1.3
-----

*A:Dut-C>tools# dump service vpls 1001 provider-tunnels
=====

VPLS 1001 Inclusive Provider Tunnels Originating
=====
ipmsi (RSVP)                                P2MP-ID  Tunl-ID  Ext-Tunl-ID
-----

ipmsi-1001-73728                          1001      61440   10.20.1.3
-----

```

```
=====
VPLS 1001 Inclusive Provider Tunnels Terminating
=====

ipmsi (RSVP)                                P2MP-ID Tunl-ID Ext-Tunl-ID
-----
                                1001      61440   10.20.1.1
                                1001      64944   10.20.1.2
-----
```

proxy-arp

| | |
|--------------------|---|
| Syntax | proxy-arp usage |
| Context | tools>dump>service |
| Description | This command provides information about the usage and limit of the system-wide proxy-arp table for all the services. The command also shows if the limit has been exceeded and a trap raised. |
| Output | <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 |
| Description | <p>This command dumps the <VTEP, VNI> bindings that have been detected as duplicate attempts, that is, an attempt to add the same binding to more than one service. The commands provides a 'clear' option.</p> |
| Output | <pre>*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</pre> |

usage

| | |
|--------------------|--|
| Syntax | usage |
| Context | tools>dump>service>vxlan |
| Description | <p>This command displays the consumed VXLAN resources in the system.</p> |
| Output | <pre>*A:PE71# tools dump service vxlan usage VXLAN usage statistics at 001 17:46:11.170:</pre> |

Show, Clear, Debug Commands

```
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 1060](#)
- [PBB Features on page 1061](#)
 - [Integrated PBB-VPLS Solution on page 1061](#)
 - [PBB Technology on page 1063](#)
 - [PBB Mapping to Existing VPLS Configurations on page 1064](#)
 - [SAP and SDP Support on page 1066](#)
 - [PBB Packet Walkthrough on page 1068](#)
 - [IEEE 802.1ak MMRP for Service Aggregation and Zero Touch Provisioning on page 1093](#)
 - [MMRP Support Over B-VPLS SAPs and SDPs on page 1095](#)
 - [PBB and BGP-AD on page 1100](#)
 - [PBB ELINE Service on page 1100](#)
 - [PBB Using G.8031-Protected Ethernet Tunnels on page 1101](#)
 - [MAC Flush on page 1110](#)
 - [Access Multi-Homing for Native PBB \(B-VPLS over SAP Infrastructure\) on page 1115](#)
 - [PBB and IGMP/MLD Snooping on page 1128](#)
 - [PBB QoS on page 1129](#)
 - [PBB OAM on page 1145](#)
- [Configuration Examples on page 1147](#)

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 102](#).

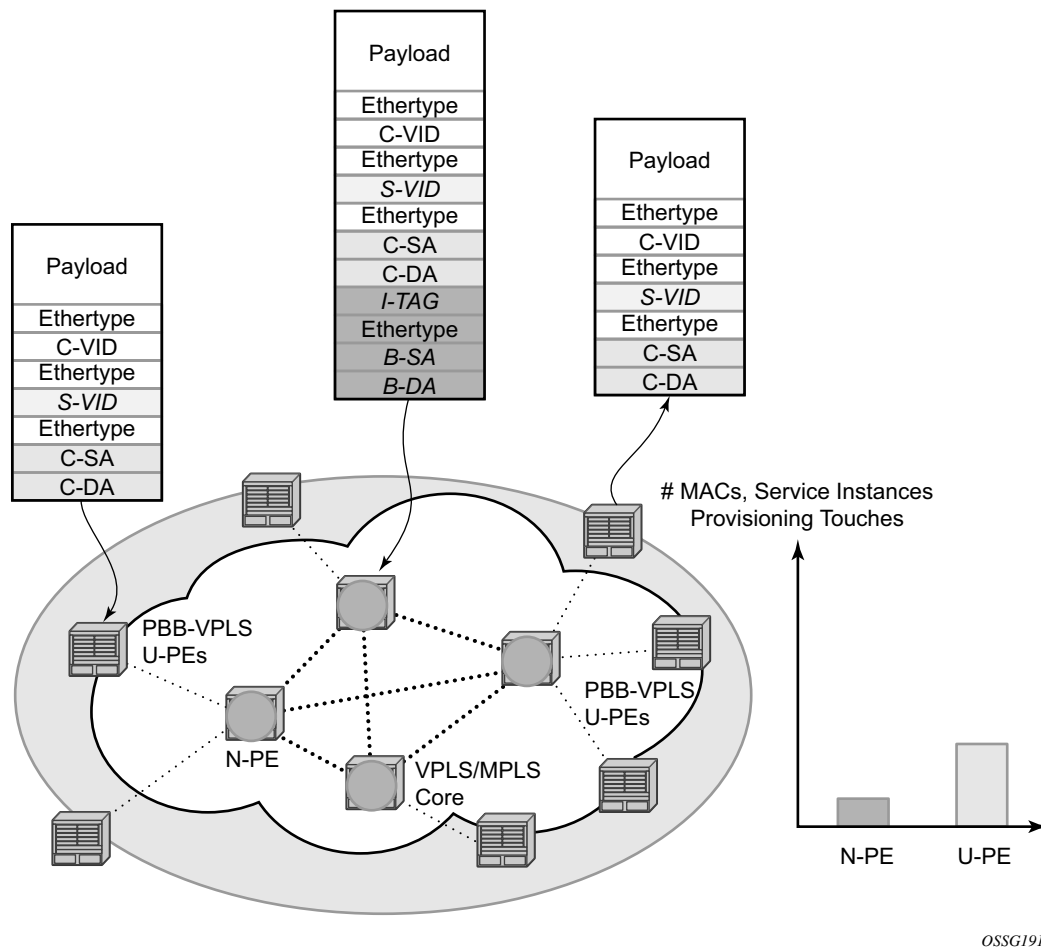


OSSG190

Figure 102: 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 103](#).



OSSG191

Figure 103: 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 104](#).

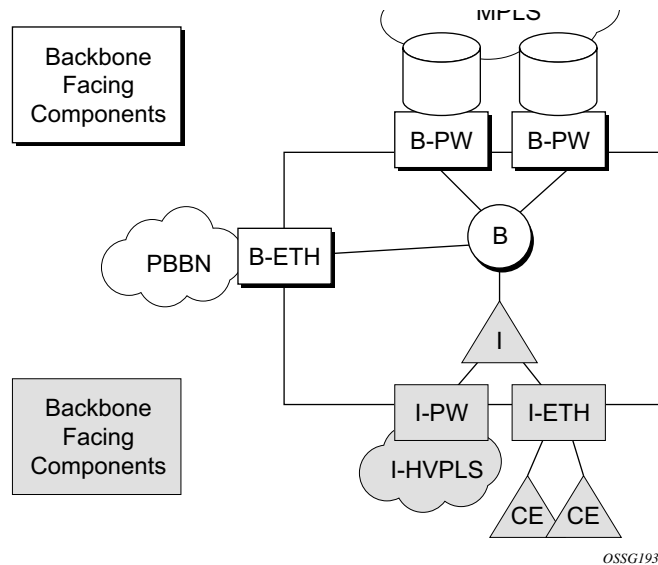


Figure 104: QinQ Payload in Provider Header Example

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 105](#).

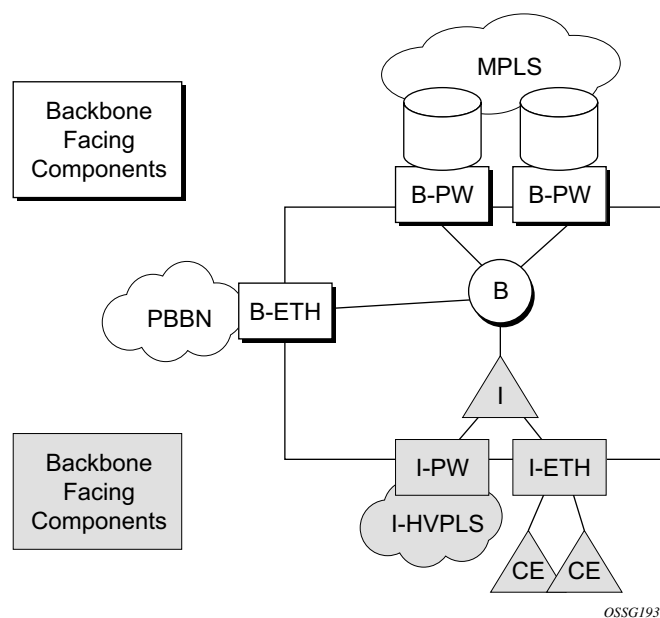


Figure 105: PBB Mapping to VPLS Constructs

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 105](#)). 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 106](#).

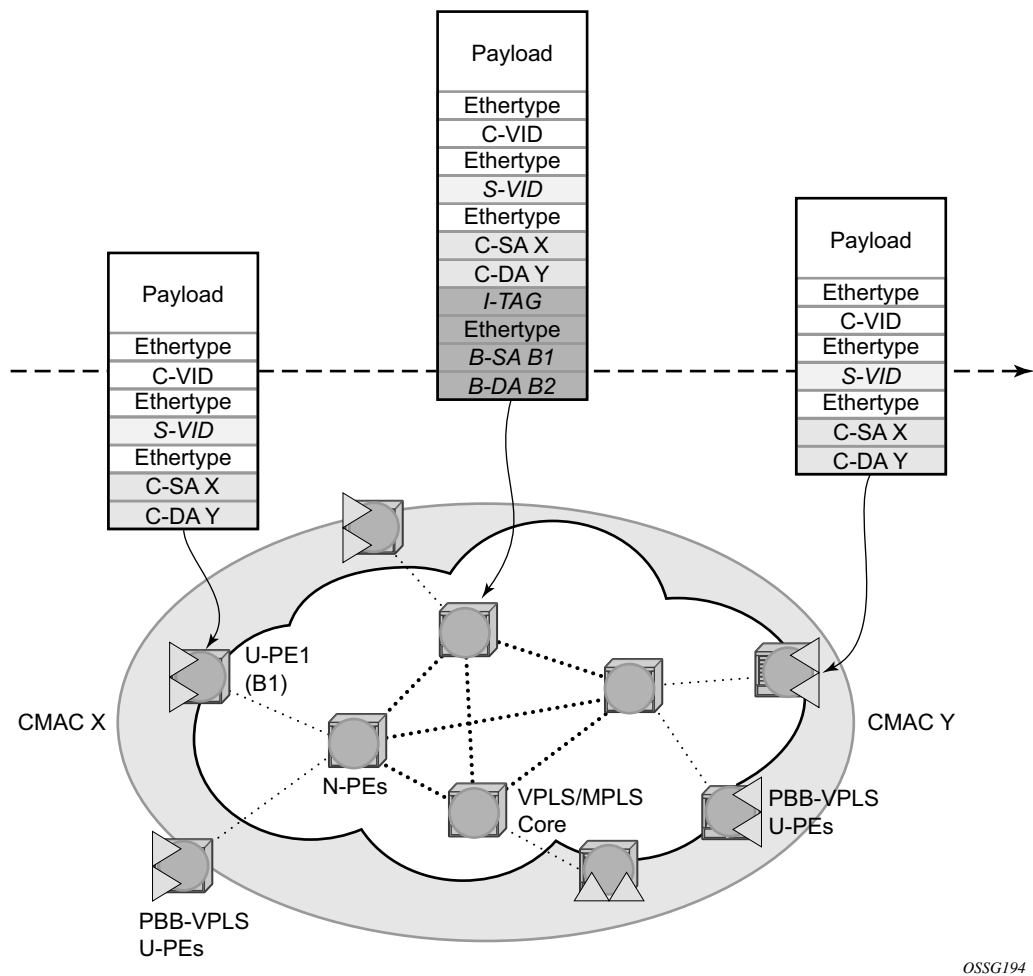


Figure 106: PBB Packet Walkthrough

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 BMAC of

the egress U-PE for CMAC Y. In our example, B2 is assumed to be known as the B-DA for Y. If CMAC Y is not present in the U-PE1 forwarding database, the PBB packet is sent in the B-VPLS using the standard group MAC address for the ISID associated with the customer VPN. If the uplink to the N-PE is a spoke pseudowire, the related PWE3 encapsulation is added in front of the B-DA.

Next, only the Backbone Header in green is used to switch the frame through the green B-VPLS/VPLS instances in the N-PEs. At the receiving U-PE2, the CMAC X is learned as being behind BMAC B1; then the PBB encapsulation is removed and the lookup for CMAC Y is performed. In the case where a pseudowire is used between N-PE and U-PE2, the pseudowire encapsulation is removed first.

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 1070](#)
- Rapid Spanning Tree Protocol (RSTP) optionally with MMRP (SAPs and spoke SDPs); see [MMRP Support Over B-VPLS SAPs and SDPs on page 1095](#).
- 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 1093](#).

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 21: 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 21](#)). 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 107 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 107 illustrates a control VPLS and two user B-VPLS. The User B-VPLS must share the same topology and are required to have interfaces on SAPs/Spoke SDPs on the same links or LAG groups as the B-VPLS. To allow services on different B-VPLS to use a path when there are multiple paths a different ECT algorithm can be configured on a B-VPLS instance. In this case, the user B-VPLS still fate shared the same topology but they may use different paths for data traffic; see [Shortest Path and Single Tree on page 1074](#).




Figure 107: 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 108](#).





Figure 108: 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 109](#) 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).



al_0023

Figure 109: 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 110](#) illustrates the multicast paths for low-path-id and high-path-id algorithm.



al_0024

Figure 110: 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 22: 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

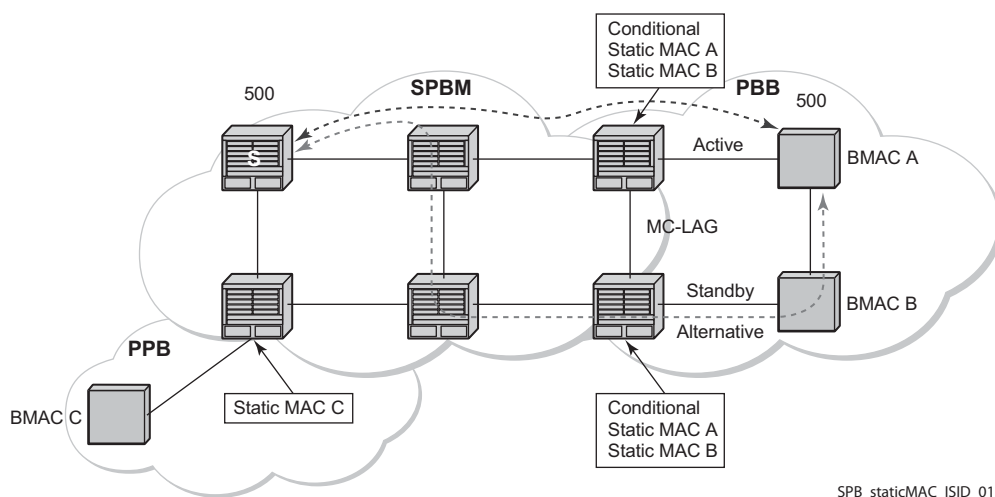
To extend SPBM networks to other PBB networks, static MACs and ISIDs can be defined under SPBM SAPs/SDPs. The declaration of a static MAC in an SPBM context allows a non-SPBM PBB system to receive frames from an SPBM system. These static MACs are conditional on the SAP/SDP operational state. (Currently 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.



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.

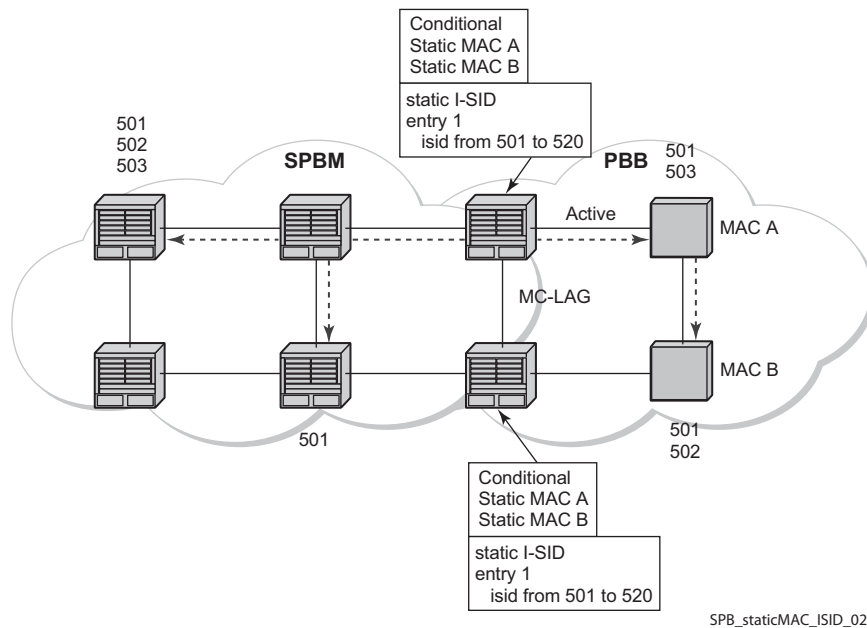


Figure 112: Static ISIDs Example

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 23](#).

Table 23: 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 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.

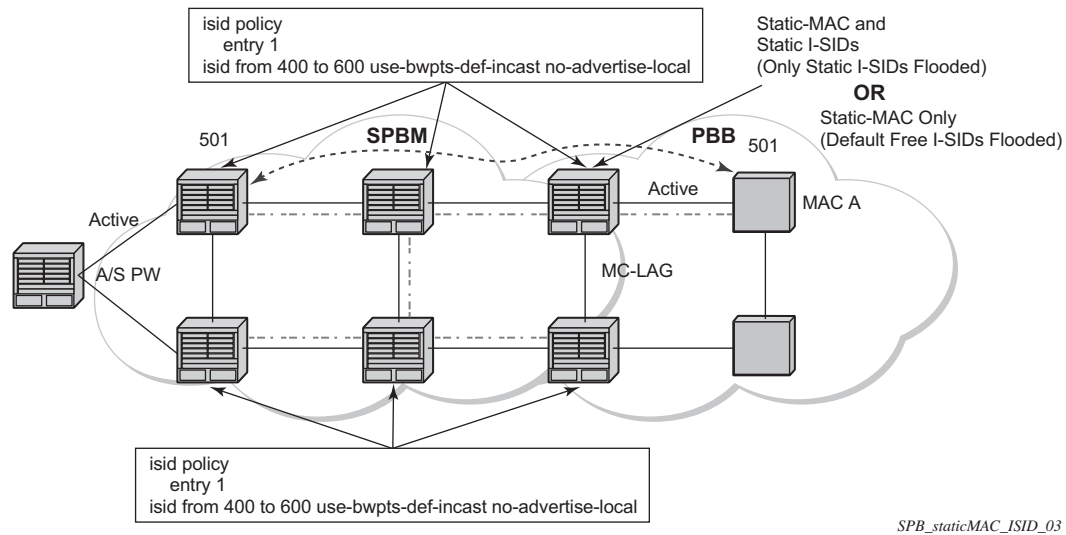


Figure 113: ISID Policy Example

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

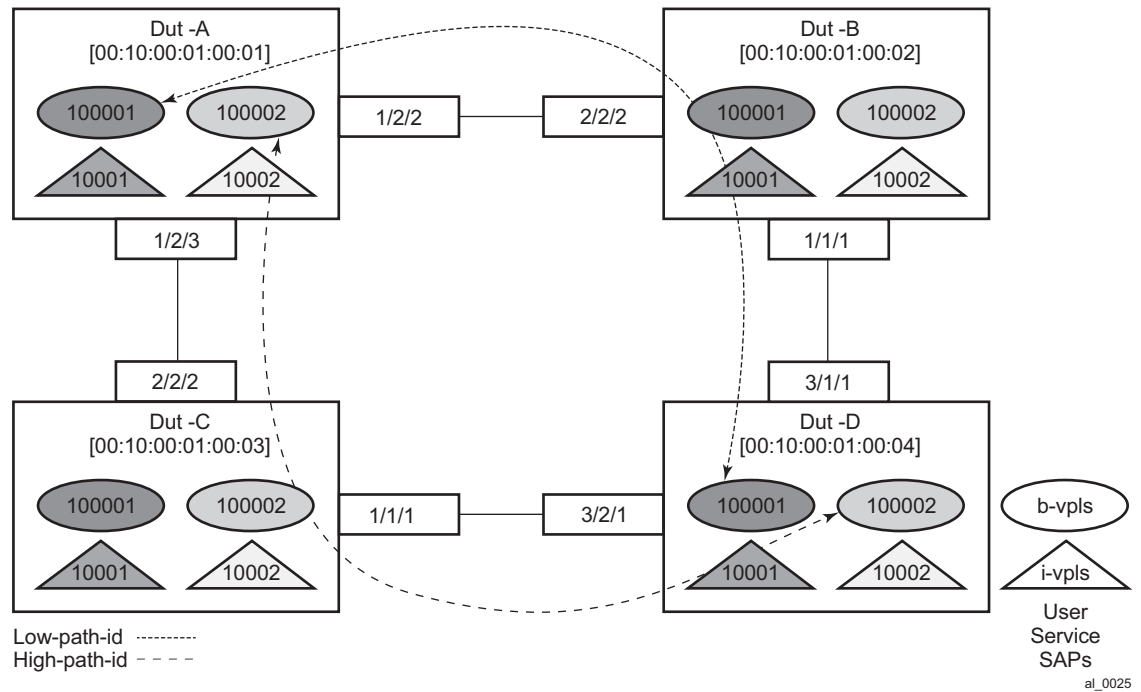


Figure 114: Sample Network

Figure 114 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 114, 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                               Ver.
    NextHop If                      SysID
-----
1    10001                             10
    sap:1/2/2:1.1                   Dut-B
    sap:1/2/3:1.1                   Dut-C
100  10002                             10
    sap:1/2/2:1.1                   Dut-B
    sap:1/2/3:1.1                   Dut-C
-----

No. of ISID Routes: 2
=====
```

The **show service spb fdb** command shows the programmed unicast and multicast source MACs in SPB-managed B-VPLS service.

```
*A:Dut-A# show service id 100001 spb fdb

=====
User service FDB information
=====
MacAddr          UCast Source          State  MCast Source          State
-----
00:10:00:01:00:02 1/2/2:1.1             ok     1/2/2:1.1             ok
00:10:00:01:00:03 1/2/3:1.1             ok     1/2/3:1.1             ok
00:10:00:01:00:04 1/2/2:1.1             ok     1/2/2:1.1             ok
-----
Entries found: 3
=====

*A:Dut-A# show service id 100002 spb fdb

=====
User service FDB information
=====
MacAddr          UCast Source          State  MCast Source          State
-----
00:10:00:02:00:02 1/2/2:1.2             ok     1/2/2:1.2             ok
00:10:00:02:00:03 1/2/3:1.2             ok     1/2/3:1.2             ok
00:10:00:02:00:04 1/2/3:1.2             ok     1/2/3:1.2             ok
-----
Entries found: 3
=====
```

The **show service spb mfib** command shows the programmed multicast ISID addresses. MACs in SPB-managed B-VPLS service shows the multicast ISID pbb group mac addresses in SPB-managed B-VPLS. 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 115](#).

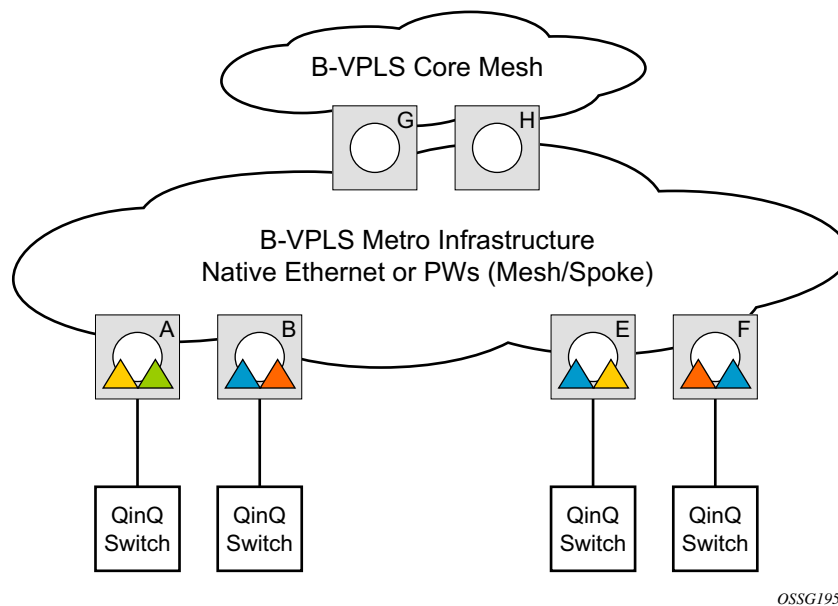


Figure 115: Customer Services Transported in 1 B-VPLS (M:1 Model)

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 116](#) 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 116](#).

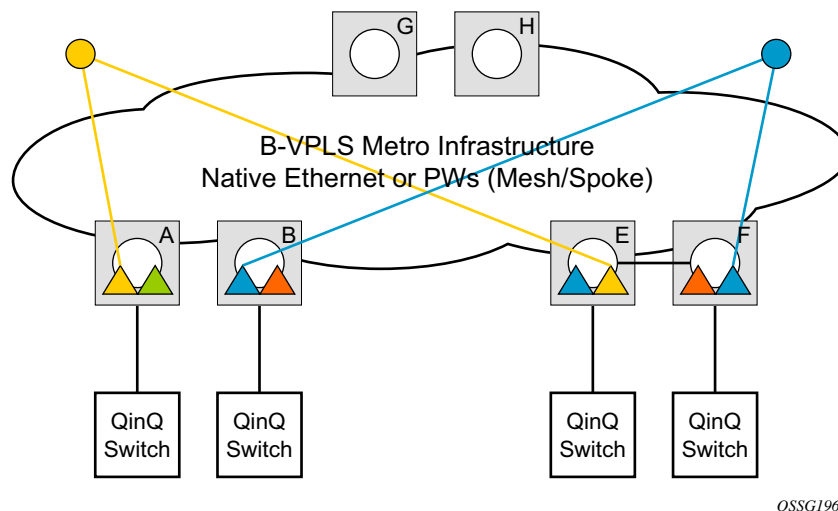


Figure 116: Flood Containment Requirement in M:1 Model

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 1155](#) 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 117: Inter-Domain Topology

Figure 117 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 118. The case of native PBB metro domains inter-connected via a MPLS core is used in this example. Other technology combinations are possible.




Figure 118: 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.

PBB Using G.8031-Protected Ethernet Tunnels

IEEE 802.1ah Provider Backbone Bridging (PBB) specification employs provider MSTP (PMSTP) to ensure loop avoidance in a resilient native Ethernet core. The usage of P-MSTP means failover times depend largely on the size and the connectivity model used in the network. The use of MPLS tunnels provides a way to scale the core while offering fast failover times using MPLS FRR. There are still service provider environments where Ethernet services are deployed using native Ethernet backbones. A solution based on native Ethernet backbone is required to achieve the same fast failover times as in the MPLS FRR case.

The Alcatel-Lucent PBB implementation offers the capability to use core Ethernet tunnels compliant with ITU-T G.8031 specification to achieve 50 ms resiliency for backbone failures. This is required to comply with the stringent SLAs provided by service providers in the current competitive environment. The implementation also allows a LAG-emulating Ethernet Tunnel providing a complimentary native Ethernet ELAN capability. The LAG-emulating Ethernet tunnels and G.8031 protected Ethernet tunnels operate independently.

The next section describes an applicability example where an Ethernet service provider using native PBB offers a carrier of carrier backhaul service for mobile operators.

Solution Overview

A simplified topology example for a PBB network offering a carrier of carrier service for wireless service providers is depicted in [Figure 119](#).




Figure 119: Mobile Backhaul Use Case

The wireless service provider in this example purchases an ELINE service between the ENNI on PBB edge nodes, BEB1 and BEBn. PBB services are employing a type of Ethernet tunneling (Eth-tunnels) between BEBs where primary and backup member paths controlled by G.8031 1:1 protection are used to ensure faster backbone convergence. Ethernet CCMs based on IEEE 802.1ag specification may be used to monitor the liveness for each individual member paths.

The Ethernet paths span a native Ethernet backbone where the BCBs are performing simple Ethernet switching between BEBs using an Epipe or a VPLS service.

Although the network diagram shows just the Epipe case, both PBB ELINE and ELAN services are supported.

Detailed Solution Description

This section discusses the details of the Ethernet tunneling for PBB. The main solution components are depicted in Figure 120.

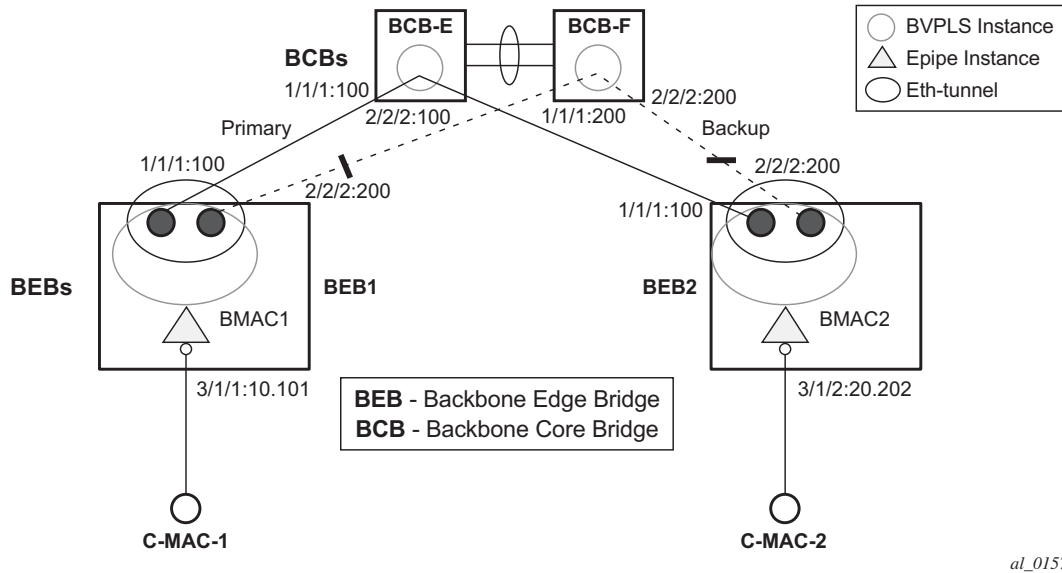


Figure 120: PBB-Epipe with B-VPLS over Ethernet Tunnel

The PBB ELINE service is represented in the BEBs as a combination of an Epipe mapped to a BVPLS instance. A eth-tunnel object is used to group two possible paths defined by specifying a member port and a control tag. In our example, the blue-circle representing the eth-tunnel is associating in a protection group the two paths instantiated as (port, control-tag/bvid): a primary one of port 1/1/1, control-tag 100 and respectively a secondary one of port 2/2/2, control tag 200.

The BCBs devices will stitch each BVID between different BEB-BCB links using either a VPLS or Epipe service. Epipe instances are recommended as the preferred option due to increased tunnel scalability.

Fast failure detection on the primary and backup paths is provided using IEEE 802.1ag CCMs that can be configured to transmit at 10 msec interval. Alternatively, the link layer fault detection mechanisms like LoS/RDI or 802.3ah can be employed.

Path failover is controlled by an Ethernet protection module, based on standard G.8031 Ethernet Protection Switching. The Alcatel-Lucent implementation of Ethernet protection switching supports only the 1:1 model which is common practice for packet based services since it makes

better use of available bandwidth. The following additional functions are provided by the protection module:

- Synchronization between BEBs such that both send and receive on the same Ethernet path in stable state.
- Revertive / non-revertive choices.
- Compliant G.8031 control plane.

The secondary path requires a MEP to exchange the G.8031 APS PDUs. The following Ethernet CFM configuration in the **eth-tunnel>path>eth-cfm>mep** context can be used to enable the G.8031 protection without activating the Ethernet CCMs:

- Create the domain (MD) in CFM.
- Create the association (MA) in CFM. NOTE: Do not put remote MEPs.
- Create the MEP.
- Configure control-mep and no shutdown on the MEP.
- The CCM transmission should stay disabled using the **no ccm-enable** command.

If a MEP is required for troubleshooting issues on the primary path, the configuration described above for the secondary path must be used to enable the use of Link Layer OAM on the primary path.

LAG loadsharing is offered to complement G.8031 protected Ethernet tunnels for situations where unprotected VLAN services are to be offered on some or all of the same native Ethernet links.




Figure 121: G.8031 P2P Tunnels and LAG-Like Loadsharing Co-Existence

In [Figure 121](#), the G.8031 Ethernet tunnels are used by the B-SAP(s) mapped to the green BVPLS entities supporting the ELINE services. A LAG-like loadsharing solution is provided for the Multipoint BVPLS (white circles) supporting the ELAN (IVPLS) services. The green G.8031 tunnels co-exist with LAG-emulating Ethernet tunnels (loadsharing mode) on both BEB-BCB and BCB-BCB physical links.

The G.8031-controlled Ethernet tunnels will select an active tunnel based on G.8031 APS operation, while emulated-LAG Ethernet tunnels will hash traffic within the configured links. Upon failure of one of the links the emulated-LAG tunnels will rehash traffic within the remaining links and fail the tunnel once the number of links breaches the minimum required (independent of G.8031-controlled Ethernet tunnels on the links shared emulated-LAG).

Detailed PBB Emulated LAG Solution Description

This section discusses the details of the emulated LAG Ethernet tunnels for PBB. The main solution components are depicted in Figure 122 which overlays Ethernet Tunnels services on the network from Figure 120.

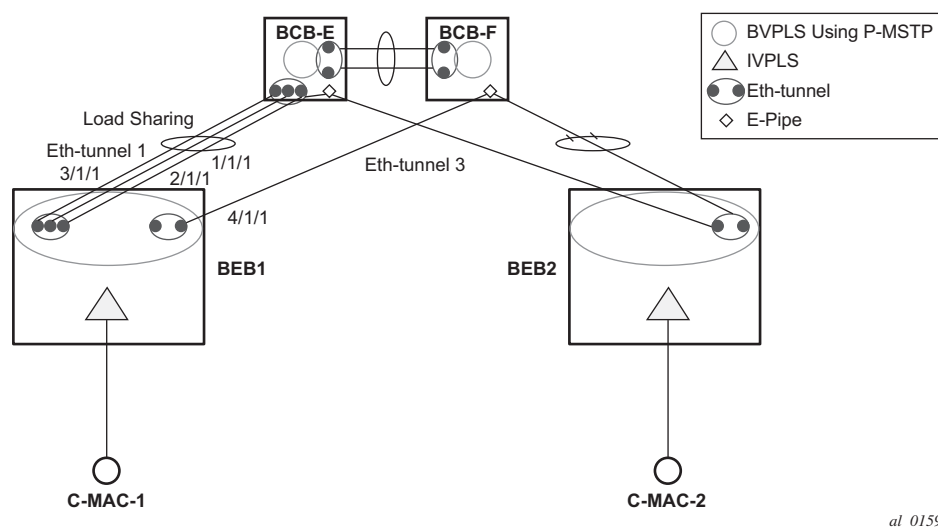


Figure 122: Ethernet Tunnel Overlay

For a PBB Ethernet VLAN to make efficient use of an emulated LAG solution, a Management-VPLS (m-VPLS) is configured enabling Provider Multi-Instance Spanning Tree Protocol (P-MSTP). The m-VPLS is assigned to two SAPs; the eth-tunnels connecting BEB1 to BCB-E and BCB-F respectively reserving a range of VLANs for P-MSTP.

The PBB P-MSTP service is represented in the BEBs as a combination of an Epipe mapped to a BVPLS instance as before but now the PBB service is able to use the Ethernet tunnels under the P-MSTP control and load share traffic on the emulated LAN. In our example, the blue-circle representing the BVPLS is assigned to the SAPs which define two paths each. All paths are specified as primary precedence to load share the traffic.

A Management VPLS (m-VPLS) is first configured with a VLAN-range and assigned to the SAPs containing the path to the BCBs. The load shared eth-tunnel objects are defined by specifying a member ports and a control tag of zero. Then individual B-VPLS services can be assigned to the member paths of the emulated LAGs and defining the path encapsulation. Then individual services such as the IVPLS service can be assigned to the B-VPLS.

At the BCBs the tunnels are terminated the next BVPLS instance controlled by P-MSTP on the BCBs to forward the traffic.

In the event of link failure, the emulated LAG group will automatically adjust the number of paths. A threshold can be set whereby the LAG group is declared down. All emulated LAG operations are independent of any 8031-1to1 operation.

Support Service and Solution Combinations

The following considerations apply when Ethernet tunnels are configured under a VPLS service:

- Only ports in access or hybrid mode can be configured as eth-tunnel path members. The member ports can be located on the same or different IOMs or MDAs.
- 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.
- The following solutions or features are not supported in the current implementation and are blocked:
 - Capture SAP
 - Subscriber management
 - BSX
 - Eth-tunnels usage as a logical port in the **config>redundancy>multi-chassis>peer>sync>port** context

For further information, refer to the Services Overview Guide.

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 123](#) for the case when a Topology Change Notification (TCN) is received on PBB PE 2 from a QinQ access in the I-VPLS domain.

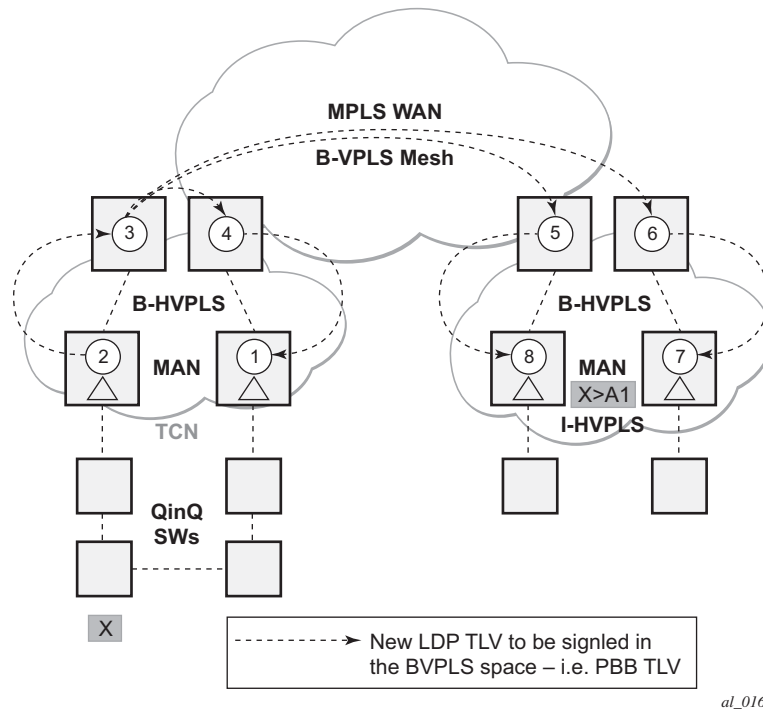


Figure 123: TCN Triggered PBB Flush-ALI-But-Mine Procedure

The received TCN may be related to one or more I-VPLS domains. This will generate a MAC Flush in the local I-VPLS instance(s) and if configured, it will originate a PBB MAC **flush-all-but-mine** throughout the related B-VPLS context(s) represented by the white circles 1-8 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 124](#). MC-LAG is used to multi-home a number of edge switches running QinQ to PBB BEBs.




Figure 124: 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 ES1-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 125](#) 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.

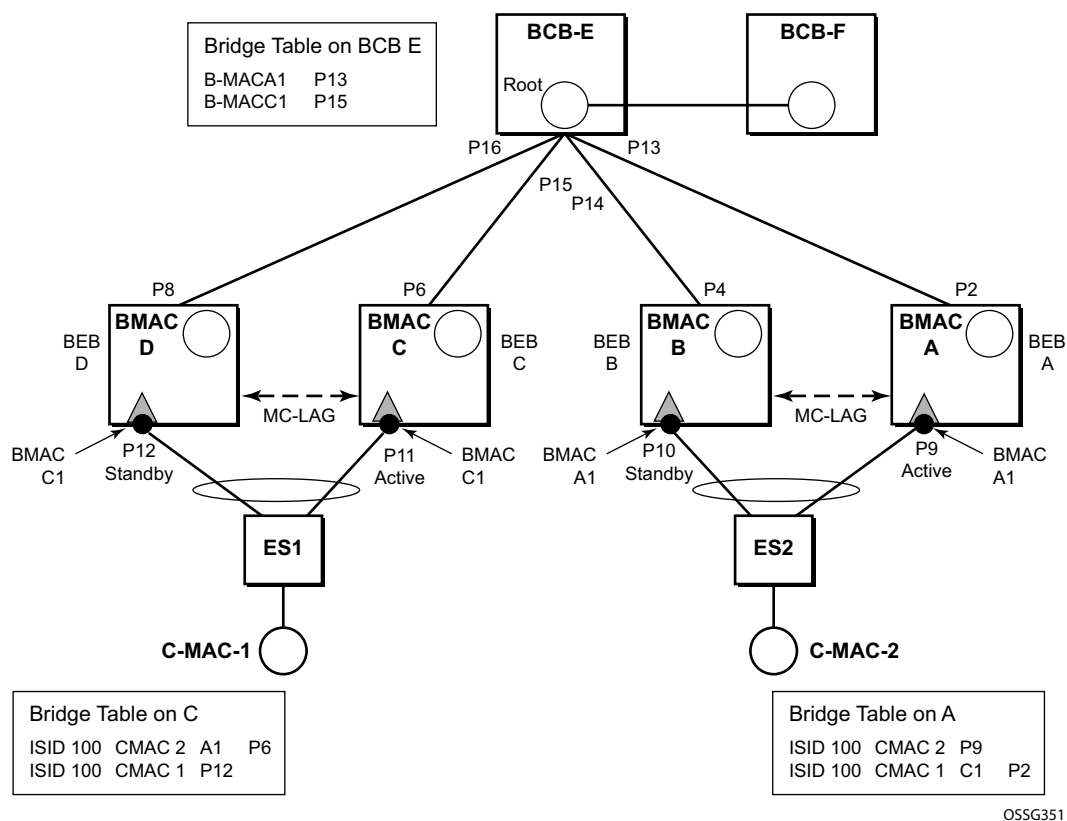


Figure 125: PBB Active Topology and Access Multi-Homing

ES1 and ES2 are dual-homed using MC-LAG into two BEB devices: ES1 to BEB C and BEB D, ES2 to BEB A and BEB B. MC-LAG P11 on BEB C and P9 on BEB A are active on each side.

In the service context, the triangles are I-VPLS instances while the small circles are B-VPLS components with the related, per BVPLS source BMACs indicated next to each BVPLS instances. P-MSTP or RSTP may be used for loop avoidance in the multi-point BVPLS. For simplicity, only the active SAPs (BEB P2, P4, P6 and P8) are shown in the diagram.

In addition to the source BMAC associated with each BVPLS, there is an additional BMAC associated with each MC-LAG supporting multi-homed I-VPLS SAPs. The BEBs that are in a multi-homed MC-LAG configuration share a common B-MAC on the related MC-LAG interfaces. For example, a common BMAC C1 is associated in this example with ports P11 and P12 participating in the MC-LAG between BEB C and BEB D while BMAC A1 is associated with ports P9 and P10 in the MC-LAG between BEB A and BEB B. While BMAC C1 is associated through the I-VPLS SAPs with both BVPLS instances in BEB C and BEB D, it is actively used for forwarding to I-VPLS SAPs only on BEB C containing the active link P11.

MC-LAG protocol keeps track of which side (port or LAG) is active and which is standby for a 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 125](#).

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 126](#) depicts the case of access link failure.

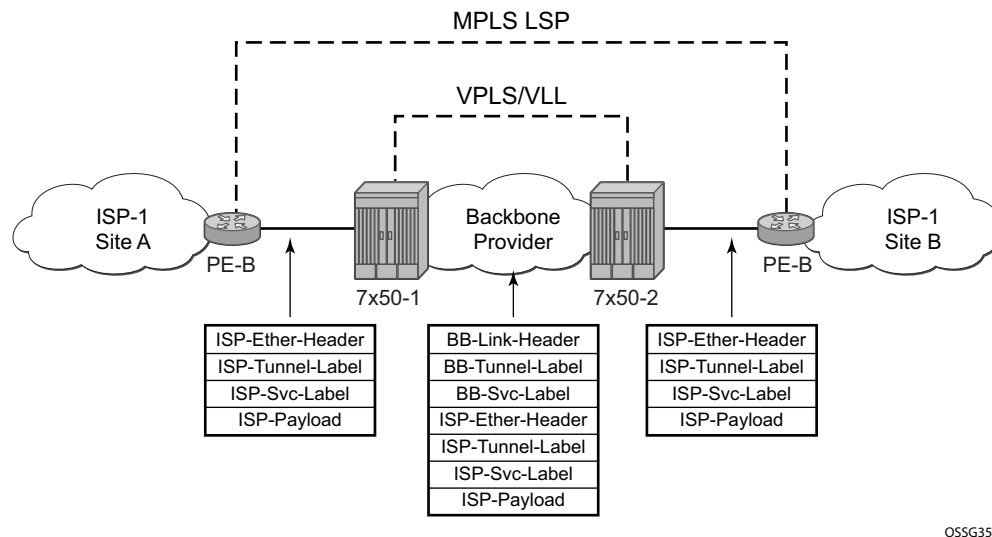


Figure 126: Access Multi-Homing - Link Failure

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.
- B-MAC C1 becomes active on BEB D and any traffic received on BEB D with destination B-MAC 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 B-MAC. 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 B-MAC 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.

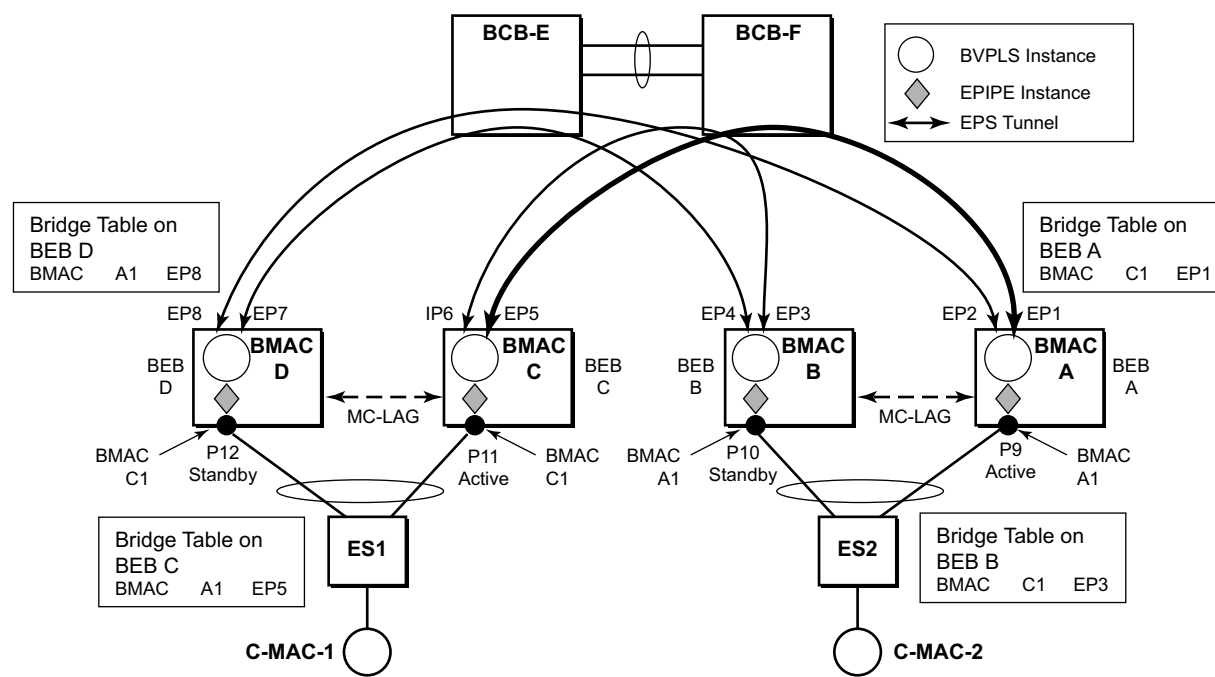
- In other topologies, it is possible that the BMAC C1 FIB entries in the B-VPLS instance on the remote BEBs (like BEB A) will need to move between B-SAPs. This will involve a move of all CMAC using as next hop BMAC C1 and the new egress linecard.

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 127](#).



O5SG353


Figure 127: 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 127. 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 128. The same procedure applies for the link failure case.



OSSG354

Figure 128: 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.
- 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 129](#).




Figure 129: 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 129](#))

- 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-bvpls-flush flush-all-from-me** is configured within the I-VPLS. This causes the system on which the site object fails to send a MAC flush all-from-me message so that customer MACs are flushed on the remote backbone edge bridges using a flush-all-from-me message. The message sent includes a PBB TLV which contains the source BMAC identifying the originator (“mine”/“me”) of the flush indication and the ISID list identifying the I-VPLS instances affected by the flush indication, see section [LDP MAC Flush Solution for PBB Blackholing on page 1111](#).

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 130](#).




Figure 130: 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 130](#), 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 130](#)):

- 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 BMAC address used in the PBB header for the query message. This is the BMAC associated with the B-VPLS instance on the remote PBB PE.

It is also possible to configure that a multicast router exists in a remote I-VPLS service. This can be achieved using the `mrouter-dest` CLI command to specify the mac-name of the destination BMAC to be used to reach the remote I-VPLS service. This command is available in the VPLS service PBB IGMP and MLD snooping contexts.

The following are not supported in a PBB I-VPLS context with IGMP snooping or MLD snooping:

- Multicast VPLS Registration (MVR)
- Multicast CAC
- configuration under a default SAP

The following are not supported in a PBB I-VPLS context with MLD snooping:

- configuration of the maximum number of multicast group sources allowed per group 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 131](#).

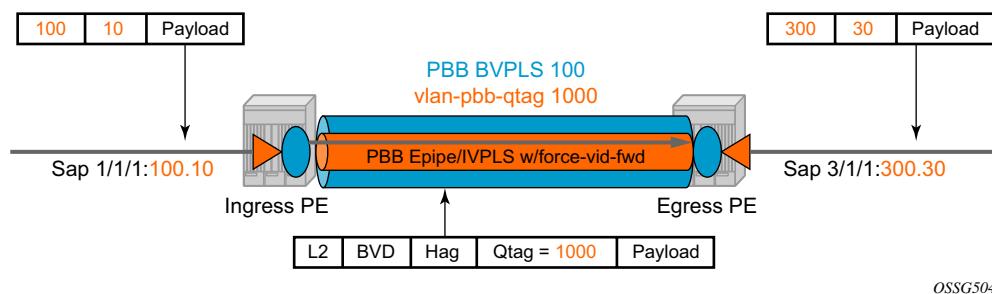


Figure 131: PCP, DE Bits Transparency in PBB

The PBB BVPLS is represented by the blue pipe in the middle with its associated COS represented through both the service (I-tag) and tunnel COS (BVID dot1p+DE or PW EXP bits).

The customer COS is contained in the orange dot1q VLAN tags managed in the customer domains. There may be one (CVID) or two (CVID, SVID) tags used to provide service classification at the SAP. IVPLS or PBB Epipe instances (orange triangles) are used to provide a Carrier-of-Carrier service.

As the VLAN tags are stripped at the ingress SAP and added back at the egress SAP, the PBB implementation must provide a way to maintain the customer QoS marking. This is done using a force-qtag-forwarding configuration on a per IVPLS/Epipe basis under the node specifying the uplink to the related BVPLS. When force-qtag-forwarding is enabled, a new VLAN tag is added

right after the CMAC addresses using the configured QTAG. The dot1p, DE bits from the specified outer/inner customer QTAG will be copied in the newly added tag.

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 131 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.

Note: This feature is supported on IOM-3 and IMM on the SR7/12 and on 7750-C12 and 7750-C4. This feature is not supported on 7710 and ESS1/SR1.

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. Conversely, the user will not be allowed to assign a FC to a policer or a queue-group queue 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 IOM 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 132 displays an example of egress queue scheduling.

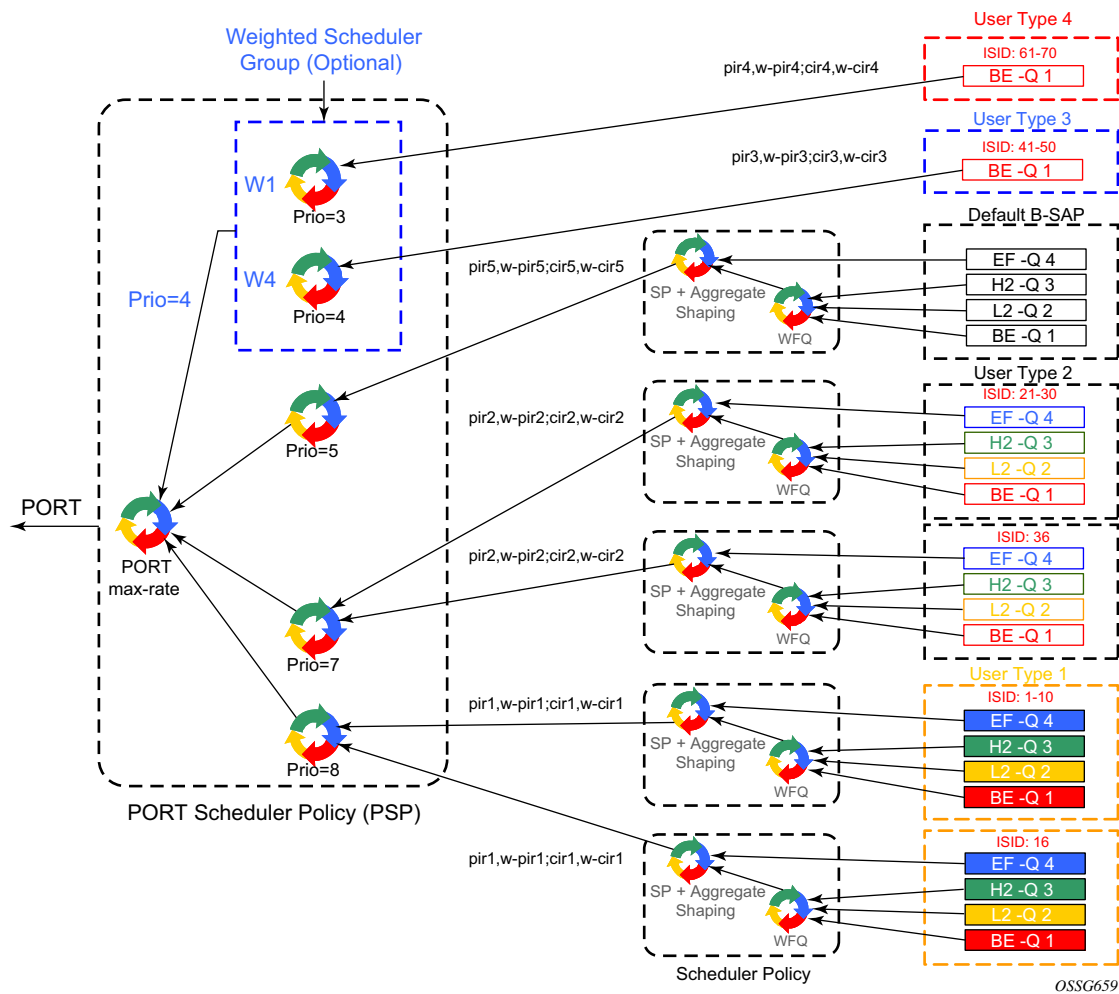


Figure 132: Egress Queue Scheduling

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 132 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 132 on page 1140](#).

```

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
```

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 133](#).



OSSG200

Figure 133: 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

The following displays PBB configuration examples:

Ethernet links on BEB1:

BEB1 to BEB1 L1:

BEB1 to BCB1 L1: 1/1/1 – Member port of LAG-emulation ET1, terminate ET3

BEB1 to BCB1 L2: 2/1/1 – Member port of LAG-emulation ET1

BEB1 to BCB1 L3: 3/1/1 - Member port of LAG-emulation ET1

BEB1 to BCB2:4/1/1 – terminate ET3

```
*A:7750_ALU>config>eth-tunnel 1
  description "LAG-emulation to BCB1 ET1"
  protection-type loadsharing
  ethernet
    mac 00:11:11:11:11:12
    encaps-type dot1q
  exit
  ccm-hold-time down 5 up 10 // 50 ms down, 1 sec up
  lag-emulation
    access adapt-qos distribute
    path-threshold 1
  exit
  path 1
    member 1/1/1
    control-tag 0
    eth-cfm
    ...
    exit
    no shutdown
  exit
  path 2
    member 2/1/1
    control-tag 0
    eth-cfm
    ...
    exit
    no shutdown
  exit
  path 3
    member 3/1/1
    control-tag 0
    eth-cfm
    ...
    exit
    no shutdown
  exit
```

Configuration Examples

```
no shutdown
-----
*A:7750_ALU>config>eth-tunnel 3
description "G.8031 tunnel ET3"
protection-type 8031_ltol
ethernet
    mac 00:11:11:11:11:11
    encap-type dot1q
exit
ccm-hold-time down 5 // 50 ms down, no up hold-down
path 1
    member 1/1/1
    control-tag 5
    precedence primary
    eth-cfm
        mep 2 domain 1 association 1
        ccm-enable
        control-mep
        no shutdown
    exit
exit
no shutdown
exit
path 2
    member 4/1/1
    control-tag 5
    eth-cfm
        mep 2 domain 1 association 2
        ccm-enable
        control-mep
        no shutdown
    exit
exit
no shutdown
exit
no shutdown
-----
# Service config
-----
*A:7750_ALU>config>service vpls 1 customer 1 m-vpls b-vpls create
description "m-VPLS for multipoint traffic"
stp
    mst-name "BVPLS"
    mode p-mstp
    mst-instance 10
        mst-priority 4096
        vlan-range 100-199
    exit
    mst-instance 20
        mst-priority 8192
        vlan-range 200-299
    exit
    no shutdown
exit

sap eth-tunnel-1 create // BSAP0 to BCB E
sap 4/1/1:0 create // physical link to BCB F (NOTE 0 or 0.*)
                    // indicate untagged for m-VPLS)
exit
```



```

no shutdown
-----
# Service config: one of the same-fate SAP over
# loadsharing tunnel
-----
A:7750_ALU>config service vpls 100 customer 1 b-vpls create
  sap eth-tunnel-1:1 create //to BCB E
    // must specify tags for each path for loadsharing
    eth-tunnel
      path 1 tag 100
      path 2 tag 100
      path 3 tag 100
    exit
  no shutdown ...
  sap 3/1/1:200 // to BCBF
  ...

A:7750_ALU>config service vpls 1000 customer 1 i-vpls create
  pbb backbone-vpls 100 isid 1000
  sap 4/1/1:200 // access SAP to QinQ
  ...
-----
# Service config: one of epipes into b-VPLS protected tunnel
# as per R7.0 R4
-----
A:7750_ALU>config service service vpls 3 customer 1 b-vpls create
  sap eth-tunnel-3 create
  ...
service epipe 2000
  pbb-tunnel 100 backbone-dest-mac to-AS20 isid 2000
  sap 3/1/1:400 create

```

CLI Syntax:

```

port 1/1/1
  ethernet
    encaps-type dot1q
port 2/2/2
  ethernet
    encaps-type dot1q
config eth-tunnel 1
  path 1
    member 1/1/1
    control-tag 100
    precedence primary
    eth-cfm
      mep 51 domain 1 association 1 direction down
      ccm-enable
      low-priority-defect allDef
      mac-address 00:AE:AE:AE:AE:AE
      control-mep
      no shutdown
    no shutdown
  path 2

```

```
member 2/2/2
control-tag 200
eth-cfm
  mep
    mep 52 domain 1 association 2 direction down
    ccm-enable
    low-priority-defect allDef
    mac-address 00:BE:BE:BE:BE:BE
    control-mep
    no shutdown
  no shutdown

config service vpls 1 b-vpls
  sap eth-tunnel-1
config service epipe 1000
  pbb-tunnel 1 backbone-dest-mac remote-beb
  sap 3/1/1:400.10
```

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 130](#).

```
*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 1155](#)
- [SAP Commands on page 1157](#)
- [Mesh SDP Commands on page 1157](#)
- [Spoke SDP Commands on page 1158](#)
- [Show Commands on page 1160](#)
- [Clear Commands on page 1160](#)
- [Debug Commands on page 1161](#)

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
    — mrp
      — [no] mrp-policy policy-name
        — description description-string
        — no description
        — scope {exclusive | template}
        — no scope
        — default-action {block | allow}

```

```

— [no] entry entry-id
— description description-string
— no description
— [no] match
— [no] isid value | from value to higher-value
— action { block | allow | end-station }
— no action
— copy mrp-policy source-name to dest-name
— renum old-entry-id to new-entry-id

config
— service
— pbb
— mac-name name ieee-address
— no mac-name
— 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] [capture-sap]
— no sap sap-id
— mrp
— [no] join-time value
— [no] leave-all-time value
— [no] leave-time value
— [no] mrp-policy policy-name
— [no] periodic-time value
— [no] periodic-timer
— [no] spb create
— [no] shutdown
— lsp-pacing-interval 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 [port port-id [vlan vlan-id]] | sdp sdp-id[:vc-id] [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
— 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] mmrp-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 }
          — [no] detail
          — [no] spf { system-id }

```


PBB Service Commands

VPLS Service Commands

vpls

| | |
|--------------------|--|
| Syntax | vpls <i>service-id</i> customer <i>customer-id</i> vpn <i>vpn-id</i> [m-vpls] [b-vpls i-vpls] [create] vpls <i>service-id</i> 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> <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>Values 1 — 2147483648</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> |

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

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 7750 SR, 7450 ESS and 7710 SR 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).

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 <i>instance</i> [fid <i>value</i>] [create] | | | | | | | | |
| Context | config>service>vpls b-vpls config>service>vpls b-vpls>sap>spb config>service>vpls b-vpls>spoke-sdp>spb | | | | | | | | |
| Description | <p>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.</p> <p>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.</p> <p>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.</p> | | | | | | | | |
| Default | no spb | | | | | | | | |
| Parameters | <p><i>instance-id</i> — Specifies the instance ID for an SPB IS-IS instance.</p> <table> <tr> <td>Values</td><td>1024–2047 (4 available)</td></tr> <tr> <td>Default</td><td>1024</td></tr> </table> <p><i>FID</i> — Specifies FID value.</p> <table> <tr> <td>Values</td><td>1-4095</td></tr> <tr> <td>Default</td><td>1</td></tr> </table> <p>Note: SPB operates with disable-learning, disable aging and discard-unknown. The state of these commands is ignored when SPB is configured.</p> | Values | 1024–2047 (4 available) | Default | 1024 | Values | 1-4095 | Default | 1 |
| Values | 1024–2047 (4 available) | | | | | | | | |
| Default | 1024 | | | | | | | | |
| Values | 1-4095 | | | | | | | | |
| Default | 1 | | | | | | | | |

spb

| | |
|--------------------|---|
| Syntax | [no] spb [create] |
| Context | config>service>vpls b-vpls>sap>spb> config>service>vpls b-vpls>spoke-sdp>spb> |
| Description | This command enables Shortest Path Bridging (SPB) on SAP or Spoke SDP. The B-VPLS may be a control B-VPLS or user B-VPLS. Since SPB uses IS-IS that supports multiple instances, SPB inherits the instance from the control B-VPLS. |

SPB at this context level is enabled immediately. SPB enables an instance of IS-IS protocol with the no shutdown command. Alternatively, the IS-IS protocol instance under SPB is disabled with the shutdown command in the **config>service>vpls b-vpls>spb** context.

Default no spb

spbm-control-vpls

| | |
|--------------------|---|
| Syntax | spbm-control-vpls <i>service-id fid fid</i> no spbm-control-vpls |
| Context | config>service>vpls <i>service-id</i> b-vpls> |
| Description | <p>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.</p> <p>A Forwarding Identifier (FID) is specified which is an abstraction of the B-VID used for forwarding in SPB. The ect-algorithm is associated with the FID and can be changed only when there are no VPLS, SAPs or SDP bindings associated with the FID. The FID must be independent from the FID assigned to other services.</p> |
| 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 | <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.</p> <p>The no form of this command administratively enables an entity.</p> <p>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.</p> |

lsp-lifetime

| | |
|--------------------|---|
| Syntax | lsp-lifetime <i>seconds</i> 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 | <i>seconds</i> — 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 <i>lsp-wait</i> [<i>lsp-initial-wait</i> [<i>lsp-second-wait</i>]] |
| Context | config>service>vpls b-vpls>spb |
| Description | 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 <code>lsp-wait</code> timer until a maximum value is reached. This is a control B-VPLS command. |
| Parameters | <p><i>lsp-max-wait</i> — Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated.</p> <p>Values 1 — 120</p> <p>Default 5</p> <p><i>lsp-initial-wait</i> — Specifies the initial LSP generation delay in seconds.</p> <p>Values 0 — 100</p> <p>Default 0</p> <p><i>lsp-second-wait</i> — Specifies the hold time in seconds between the first and second LSP generation.</p> <p>Values 1 — 100</p> <p>Default 1</p> |

overload

| | |
|--------------------|--|
| Syntax | overload [timeout <i>seconds</i>] no overload |
| Context | config>service>vpls b-vpls>spb |
| Description | This command administratively sets the SPB to operate in the overload state for a specific time period, in seconds, or indefinitely. During normal operation, the router may be forced to enter an |

overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by SPB and will not 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 SPB is overloaded or used prior to executing a shutdown command to divert traffic around the switch.

The **no** form of the command causes the router to exit the overload state.

| | |
|-------------------|---|
| Default | no overload |
| Parameters | <i>seconds</i> — The time, in seconds, that this router must operate in overload state. |
| Values | 60 — 1800 |
| Default | Infinity (overload state maintained indefinitely) |

overload-on-boot

| | |
|--------------------|---|
| Syntax | overload-on-boot [<i>timeout seconds</i>] no overload-on-boot |
| Context | config>service>vpls b-vpls>spb> |
| Description | <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"> • The timeout timer expires. • A manual override of the current overload state is entered with the config>service>vpls instance>b-vpls>spb>no overload command. <p>The no 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 config>service>vpls instance>b-vpls>spb>no overload 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 config>service>vpls instance>b-vpls>spb>no overload command.</p> <p>The no form of the command removes the overload-on-boot functionality from the configuration.</p> |
| Default | no overload-on-boot |
| Parameters | <i>seconds</i> — The time, in seconds, that this router must operate in overload state. |

| | |
|----------------|---|
| Values | 60 — 1800 |
| Default | Infinity (overload state maintained indefinitely) |

spf-wait

| | | | | | | | | | | | | | |
|--------------------|---|---------------|---------|----------------|----|---------------|-------------|----------------|------|---------------|------------|----------------|------|
| Syntax | [no] spf-wait <i>spf-wait</i> [<i>spf-initial-wait</i> [<i>spf-second-wait</i>]] | | | | | | | | | | | | |
| Context | config>service>vpls b-vpls>spb> | | | | | | | | | | | | |
| Description | <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> | | | | | | | | | | | | |
| Default | no spf-wait | | | | | | | | | | | | |
| Parameters | <p><i>spf-wait</i> — Specifies the maximum interval in seconds between two consecutive spf calculations.</p> <table> <tr> <td>Values</td><td>1 — 120</td></tr> <tr> <td>Default</td><td>10</td></tr> </table> <p><i>spf-initial-wait</i> — Specifies the initial SPF calculation delay in milliseconds after a topology change.</p> <table> <tr> <td>Values</td><td>10 — 100000</td></tr> <tr> <td>Default</td><td>1000</td></tr> </table> <p><i>spf-second-wait</i> — Specifies the hold time in milliseconds between the first and second SPF calculation.</p> <table> <tr> <td>Values</td><td>1 — 100000</td></tr> <tr> <td>Default</td><td>1000</td></tr> </table> | Values | 1 — 120 | Default | 10 | Values | 10 — 100000 | Default | 1000 | Values | 1 — 100000 | Default | 1000 |
| Values | 1 — 120 | | | | | | | | | | | | |
| Default | 10 | | | | | | | | | | | | |
| Values | 10 — 100000 | | | | | | | | | | | | |
| Default | 1000 | | | | | | | | | | | | |
| Values | 1 — 100000 | | | | | | | | | | | | |
| Default | 1000 | | | | | | | | | | | | |

level

| | |
|--------------------|--|
| Syntax | level <i>level-number</i> |
| Context | config>service>vpls b-vpls>spb> |
| Description | <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> |

VPLS Service Commands

| | |
|-------------------|---|
| Default | level 1 |
| Parameters | <i>level-number</i> — The SPB level number. |
| Values | 1, 2 |

bridge-priority

| | |
|--------------------|--|
| Syntax | bridge-priority <i>value</i> |
| Context | config>service>vpls b-vpls>spb>level level-number |
| Description | <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> <p>Values 0 — 15</p> |
| Default | 8 |

ect-algorithm

| | |
|--------------------|---|
| Syntax | ect-algorithm <i>name fid-range fid-range</i> |
| Context | config>service>vpls b-vpls>spb>level level-number |
| Description | <p>This command configures the ect-algorithm associated with a FID. Names are:</p> <ul style="list-style-type: none">• low-path-id• high-path-id <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> |

The ect-algorithm is associated with the FID and can only be changed only when there are no VPLS, SAPs or SDP bindings associated with the FID. The FID must be independent from the FID assigned to other services.

| | |
|-------------------|--|
| Default | low-path-id |
| Parameters | <i>name</i> — low-path-id, high-path-id <i>fid-range</i> — Range of Forwarding Identifier values. |
| Values | 1 — 4095 |

forwarding-tree-topology

| | |
|--------------------|---|
| Syntax | forwarding-tree-topology unicast [st spf] |
| Context | config>service>vpls b-vpls>spb>level level-number |
| Description | This command sets the unicast forwarding to follow the shortest path tree defined by the ECT algorithm shortest path forwarding (spf) or to follow a single tree. (st). Shortest path trees make use of more link resources. Multicast traffic is defaulted to follow the single tree topology. A single tree unicast would make Multicast and unicast follow the same path. |
| Default | spf |

lsp-pacing-interval

| | |
|--------------------|--|
| Syntax | lsp-pacing-interval <i>milliseconds</i> no lsp-pacing-interval |
| Context | config>service>vpls b-vpls>sap>spb> config>service>vpls b-vpls>spoke-sdp>spb> |
| Description | This command configures the interval between SPB LSP PDUs sent from this interface. This command is valid only for interfaces on control B-VPLS. 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 SPB IS-IS LSP's can be sent from the interface expressed as a decimal integer. 0 — 65535 |

retransmit-interval

VPLS Service Commands

| | |
|--------------------|--|
| Syntax | retransmit-interval <i>seconds</i> no retransmit-interval |
| Context | config>service>vpls b-vpls>sap>spb> config>service>vpls b-vpls>spoke-sdp>spb> |
| Description | This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface. This command is valid only for interfaces on control B-VPLS. The no form of the command reverts to the default value. |
| Default | 100 |
| Parameters | <i>seconds</i> — The interval in seconds that SPB IS-IS LSPs can be sent on the interface. Values 1 — 65535 |

metric

| | |
|--------------------|--|
| Syntax | metric <i>value</i> No metric |
| Context | config>service>vpls b-vpls>sap>spb>level config>service>vpls b-vpls>spoke-sdp>spb>level |
| Description | This configures metric for this SPB interface SAP/spoke-sdp. This command is valid only for interfaces on control B-VPLS. Values 1 — 16,777,215 Default 1000 |

hello-interval

| | |
|--------------------|--|
| Syntax | hello-interval <i>seconds</i> no hello-interval |
| Context | config>service>vpls b-vpls>sap>spb>level config>service>vpls b-vpls>spoke-sdp>spb>level |
| Description | This command configures the interval in seconds between hello messages issued on this interface at this level. This command is valid only for interfaces on control B-VPLS. The no form of the command to reverts to the default value. |
| Default | 3 — Hello interval default for the designated 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 b-vpls>sap>spb>level config>service>vpls b-vpls>spoke-sdp>spb>level |
| Description | This command configures the number of missing hello PDUs from a neighbor SPB declares the adjacency down. This command is valid only for interfaces on control B-VPLS. The no form of the command reverts to the default value. |
| Default | 3 — SPB can miss up to 3 hello messages before declaring the adjacency down. |
| Parameters | <i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer. Values 2 — 100 |

mrp

| | |
|--------------------|---|
| Syntax | mrp |
| Context | config>service>vpls config>service>vpls>mesh-sdp config>service>vpls>sap config>service>vpls>spoke-sdp |
| Description | This command configures Multiple Registration Protocol (MRP) parameters. MRP is valid only under B-VPLS. |

attribute-table-size

| | |
|--------------------|--|
| Syntax | [no] attribute-table-size <i>value</i> |
| Context | config>service>vpls>mrp |
| Description | This command controls the number of attributes accepted on a per B-VPLS basis. When the limit is reached, no new attributes will be registered. If a new lower limit (smaller than the current number of attributes) from a local or dynamic I-VPLS is being provisioned, a CLI warning will be issued stating that the system is currently beyond the new limit. The value will be accepted, but any creation of new attributes will be blocked under the attribute count drops below the new limit; the software will then start enforcing the new limit. |
| Default | maximum number of attributes |
| Parameters | <i>value</i> — [1-2048] for ESS-6/7/12 or SR-7/SR-12 [1-1023] for ESS1/SR1/7710 |

attribute-table-high-wmark

| | |
|--------------------|--|
| Syntax | [no] attribute-table-high-wmark <i>high-water-mark</i> |
| Context | config>service>vpls>mrp |
| Description | This command specifies the percentage filling level of the MMRP attribute table where logs and traps are sent. |
| Default | 95% |
| Parameters | <i>high-water-mark</i> — 1%-100% |

attribute-table-low-wmark

| | |
|--------------------|---|
| Syntax | [no] attribute-table-low-wmark <i>low-water-mark</i> |
| Context | config>service>vpls>mrp |
| Description | This command specifies the MMRP attribute table low watermark as a percentage. When the percentage filling level of the MMRP attribute table drops below the configured value, the corresponding trap is cleared and/or a log entry is added. |
| Default | 90% |
| Parameters | <i>low-water-mark</i> — 1%-100% |

flood-time

| | |
|--------------------|---|
| Syntax | flood-time <i>flood-time</i> no flood-time |
| Context | config>service>vpls>mrp |
| Description | This command configures the amount of time, in seconds, after a status change in the VPLS service during which traffic is flooded. 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. |
| Default | no flood-time |
| Parameters | <i>flood-time</i> — Specifies the MRP flood time, in seconds. Values 3 — 600 |

mrp

| | |
|--------------------|--|
| Syntax | mrp |
| Context | config>service |
| Description | This command configures a Multi-service Route Processor (MRP). |

mrp-policy

| | |
|--------------------|---|
| Syntax | [no] mrp-policy <i>policy-name</i> |
| Context | config>service>mrp |
| Description | <p>This command enables the context for a MRP policy. The <i>mrp-policy</i> specifies either a forward or a drop action for the Group BMAC attributes associated with the ISIDs specified in the match criteria. The <i>mrp-policy</i> can be applied to multiple BVPLS services as long as the scope of the policy is template.</p> <p>Any changes made to the existing policy, using any of the sub-commands, will be applied immediately to all services where this policy is applied. For this reason, when many changes are required on a <i>mrp-policy</i>, it is recommended that the policy be copied to a work area. That work-in-progress policy can be modified until complete and then written over the original <i>mrp-policy</i>. Use the <i>config mrp-policy copy</i> command to maintain policies in this manner.</p> <p>The no form of the command deletes the <i>mrp-policy</i>. An MRP policy cannot be deleted until it is removed from all the SAPs or SDPs where it is applied.</p> |
| Default | no <i>mrp-policy</i> is defined |
| Parameters | <p><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.</p> <p>create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the create keyword.</p> |

scope

| | |
|--------------------|--|
| Syntax | scope {exclusive template} no scope |
| Context | config>service>mrp>mrp-policy |
| Description | <p>This command configures the filter policy scope as exclusive or template. If the scope of the policy is template and is applied to one or more services, the scope cannot be changed.</p> <p>The no form of the command sets the scope of the policy to the default of template.</p> |
| Default | template |

- Parameters**
- exclusive** — When the scope of a policy is defined as exclusive, the policy can only be applied to a single entity (SAP or SDP). Attempting to assign the policy to a second entity will result in an error message. If the policy is removed from the entity, it will become available for assignment to another entity.
 - template** — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs or network ports.

default-action

- Syntax** **default-action {block | allow}**
- Context** config>service>mrp>mrp-policy
- Description** This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs do not match the specified criteria in all of the entries of the mrp-policy.
- When multiple default-action commands are entered, the last command will overwrite the previous command.
- Default** default-action-allow
- Parameters**
- block** — Specifies that all MMRP attributes will not be declared or registered unless there is a specific mrp-policy entry which causes them to be allowed on this SAP/SDP.
 - allow** — Specifies that all MMRP attributes will be declared and registered unless there is a specific mrp-policy entry which causes them to be blocked on this SAP/SDP.

entry

- Syntax** **[no] entry entry-id**
- Context** config>service>mrp>mrp-policy
- Description** This command creates or edits an mrp-policy entry. Multiple entries can be created using unique entry-id numbers within the policy. The implementation exits the policy on the first match found and executes the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit. An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.
- The no form of the command removes the specified entry from the mrp-policy. Entries removed from the mrp-policy are immediately removed from all services where the policy is applied.
- The no form of the command removes the specified entry-id.
- Default** none
- Parameters**
- entry-id** — An entry-id uniquely identifies a match criteria and the corresponding action. It is recommended that multiple entries be given entry-ids in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

Values 1-65535

create — Keyword; required when first creating the configuration context. Once the context is created, one can navigate into the context without the create keyword.

description

| | |
|--------------------|--|
| Syntax | description <i>description-string</i> no description |
| Context | config>service>mrp>mrp-policy>entry config>service>mrp>mrp-policy |
| Description | This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The no form of this command removes the string from the configuration. |
| Default | No description associated with the configuration context. |
| Parameters | <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. |

match

| | |
|--------------------|--|
| Syntax | [no] match |
| Context | config>service>mrp>mrp-policy>entry |
| Description | This command creates the context for entering/editing match criteria for the mrp-policy entry. When the match criteria have been satisfied the action associated with the match criteria is executed. In the current implementation just one match criteria (ISID based) is possible in the entry associated with the mrp-policy. Only one match statement can be entered per entry. The no form of the command removes the match criteria for the entry-id. |

isid

| | |
|--------------------|---|
| Syntax | [no] isid <i>value</i> from <i>value</i> to <i>higher-value</i> |
| Context | config>service>mrp>mrp-policy>entry>match |
| Description | This command configures an ISID value or a range of ISID values to be matched by the mrp-policy parent when looking at the related MMRP attributes (Group BMACs). The pbb-etype value for the related SAP (inherited from the ethernet port configuration) or for the related SDP binding (inherited from SDP configuration) will be used to identify the ISID tag. |

Multiple isid statements are allowed under a match node. The following rules govern the usage of multiple isid statements:

- overlapping values are allowed:
 - isid from 1 to 10
 - isid from 5 to 15
 - isid 16
- the minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with “isid from 1 to 16” statement.
- there is no consistency check with the content of isid statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation to execute the associated action for that entry and then to exit the mrp-policy.
- If there are no isid statements under a match criteria but the mac-filter type is isid the following behaviors apply for different actions:
 - For end-station – it treats any ISID value as no match and goes to next entry or default action which must be “block” in this case
 - For allow – it treats any ISID value as a match and allows it
 - For block – it treats any ISID value as a match and blocks it

The **no** form of the command can be used in two ways:

no isid - removes all the previous statements under one match node

no isid value | 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” will not work but “no isid 16 to 100 will be successful.

| | |
|-------------------|---|
| Default | no isid |
| Parameters | <i>value or higher-value</i> — Specifies the ISID value in 24 bits. When just one present identifies a particular ISID to be used for matching. |
| Values | 0..16777215 |
| | <i>from value to higher-value</i> — Identifies a range of ISIDs to be used as matching criteria. |

action

| | |
|--------------------|--|
| Syntax | action {block allow end-station} no action |
| Context | config>service>mrp>mrp-policy>entry |
| Description | <p>This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs match the specified ISID criteria in the related entry.</p> <p>The action keyword must be entered for the entry to be active. Any filter entry without the action keyword will be considered incomplete and will be inactive. If neither keyword is specified (no action is used), this is considered a No-Op policy entry used to explicitly set an entry inactive without modifying match criteria or removing the entry itself. Multiple action statements entered will overwrite previous actions parameters when defined. To remove a parameter, use the no form of the</p> |

action command with the specified parameter.

The **no** form of the command removes the specified action statement. The entry is considered incomplete and hence rendered inactive without the action keyword.

| | |
|-------------------|---|
| Default | no action |
| Parameters | <p>block — Specifies that the matching MMRP attributes will not be declared or registered on this SAP/SDP.</p> <p>allow — Specifies that the matching MMRP attributes will be declared and registered on this SAP/SDP.</p> <p>end-station — Specifies that an end-station emulation is present on this SAP/SDP for the MMRP attributes related with matching ISIDs. Equivalent action with the block keyword on that SAP/SDP— the attributes associated with the matching ISIDs do not get declared or registered on the SAP/SDP. The matching attributes on the other hand are mapped as static MMRP entries on the SAP/SDP which implicitly instantiates in the data plane as a MFIB entry associated with that SAP/SDP for the related Group BMAC. For the other SAPs/SDPs in the BVPLS with MRP enabled (no shutdown) this means permanent declaration of the matching attributes, same as in the case when the IVPLS instances associated with these ISIDs were locally configured.</p> <p>If an mrp-policy has end-station action in one entry, the only default action allowed in the policy is block. Also no other actions are allowed to be configured in other entry configured under the policy.</p> <p>This policy will apply even if the MRP is shutdown on the local SAP/SDP or for the whole BVPLS to allow for manual creation of MMRP entries in the data plane. Specifically the following rules apply:</p> <ul style="list-style-type: none"> – If service vpls mrp shutdown then MMRP on all SAP/SDPs is shutdown - MRP PDUs pass-through transparently – If service vpls mrp no shutdown and endstation statement (even with no ISID values in the related match statement) is used in a mrp-policy applied to SAP/SDP - no declaration is sent on SAP/SDP. The provisioned ISIDs in the match statement are registered on that SAP/SDP and are propagated on all the other MRP enabled endpoints. |

copy

| | |
|--------------------|--|
| Syntax | copy mrp-policy <i>source-name to dest-name</i> |
| Context | config>service>mrp |
| Description | <p>This command copies existing mrp-policy list entries for a specific policy name to another policy name. The copy command is a configuration level maintenance tool used to create new mrp-policy using existing mrp-policy.</p> <p>An error will occur if the destination policy name exists.</p> |
| Parameters | <p>mrp-policy — Indicates that source-name and dest-name are MRP policy names.</p> <p><i>source-name</i> — Identifies the source mrp-policy from which the copy command will attempt to copy. The mrp-policy with this name must exist for the command to be successful.</p> |

dest-name — Identifies the destination mrp-policy to which the copy command will attempt to copy. If the mrp-policy with *dest-name* exist within the system an error message is generated.

renum

| | |
|--------------------|---|
| Syntax | renum <i>old-entry-id</i> to <i>new-entry-id</i> |
| Context | config>service>mrp>mrp-policy |
| Description | This command renumbers existing MRP policy entries to properly sequence policy entries. This may be required in some cases since the implementation exits when the first match is found and executes the actions according to the accompanying action command. This requires that entries be sequenced correctly from most to least explicit. |
| Parameters | <i>old-entry-id</i> — Specifies the entry number of an existing entry. Values 1-65535 <i>new-entry-id</i> — Specifies the new entry number to be assigned to the old entry. If the new entry exists, an error message is generated. |

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> — [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 | 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 <i>leave-time</i> 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 LeaveAllTime<T<1.5*leave-all-time when it is started. Refer to IEEE 802.1ak-2007 section 10.7.4.3.

Default 100

Parameters *value* — [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.

Default 10

Parameters *value* — [10-100] tenths of a second

periodic-timer

Syntax [no] periodic-timer

VPLS Service Commands

| | |
|--------------------|--|
| Context | config>service>vpls>sap>mrp config>service>vpls>spoke-sdp>mrp config>service>vpls>mesh-sdp>mrp |
| Description | This command enables or disables the Periodic Transmission Timer. |
| Default | disabled |

send-flush-on-failure

| | |
|--------------------|--|
| Syntax | [no] send-flush-on-failure |
| Context | config>service>vpls |
| Description | <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> |
| Default | no send-flush-on-failure |

pbb

| | |
|--------------------|--|
| Syntax | pbb |
| Context | config>service config>service>vpl config>service>epipe |
| Description | This command configures global PBB parameters. |

mac-name

| | |
|--------------------|---|
| Syntax | mac-name <i>name ieee-address</i> no mac-name <i>name</i> |
| 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 | <i>name</i> — 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 <i>ieee-address</i> 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 | <i>ieee-address</i> — 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 <i>ieee-address</i> |
| 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 | <i>ieee-address</i> — Specifies the backbone source MAC address. |

backbone-vpls

| | |
|--------------------|--|
| Syntax | backbone-vpls <i>vpls-id</i> [: <i>isid</i>] no backbone-vpls |
| Context | config>service>vpls>pbb |
| Description | This command associated the I-VPLS with the B-VPLS service. The ISID value is used to mux/demux packets for the VPLS flowing through the B-VPLS. |
| Parameters | <i>vpls-id</i> — This value represents the VPLS ID value associated with the B-VPLS. <i>isid</i> — 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 | <p>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.</p> <p>The no form of this command disables the command.</p> |

source-bmac

| | |
|--------------------|---|
| Syntax | source-bmac <i>ieee-address</i> |
| Context | config>service>vpls bvpls>pbb |
| Description | <p>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.</p> |

use-sap-bmac

| | |
|--------------------|--|
| Syntax | [no] use-sap-bmac |
| Context | config>service>vpls bvpls>pbb |
| Description | <p>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.</p> |
| Default | no use-sap-bmac |

mac-notification

| | |
|--------------------|---|
| Syntax | mac-notification |
| Context | config>service>vpls bvpls |
| Description | <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> |

1. When enabled in the BVPLS using no shutdown, a MAC notification will be sent for every active MC-LAG link. The following 3 cases assume no shutdown in the BVPLS.
2. Whenever a related MC-LAG link becomes active (related MC-LAG link = has at least 1 SAP associated with the BVPLS) if the MC-LAG peering is initialized and the PE peers are synchronized.
3. 1st SAP on an active MC-LAG is associated (via IVPLS/Epipe) with the BVPLS
4. The link between IVPLS/Epipe and BVPLS is configured and there are I-SAPs configured on an active MC-LAG link.

The MAC notification is not sent for the following events:

1. Change of source-bmac or source-bmac-lsb
2. On changes of use-sap-bmac parameter
3. If MC-LAG peering is not (initialized and in sync).

interval

| | |
|--------------------|---|
| Syntax | [no] interval <i>value</i> |
| Context | config>service>vpls>pbb>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>pbb>mac-notification |
| Description | This command controls the periodic interval at which sets of MAC notification messages are sent. At each expiration of the renotify timer, a new burst of notification messages is sent, specifically <count> frames at <interval> deci-seconds. |
| Default | no renotify |
| Parameters | <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>pbb>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. |
| 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 <i>service-id</i> [isid <i>isid</i>] no backbone-vpls |
| Context | config>service>vpls>pbb |
| Description | This command configures B-VPLS service associated with the I-VPLS. |
| Parameters | <i>service-id</i> — Specifies the service ID. |
| Values | 1..2147483648 |
| | <i>isid</i> — 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 <i>mac-name</i> |
| 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 1182 . |
| Parameters | <i>mac-name</i> — Specifies the MAC name. |
| Values | 32 chars max |

sap

| | |
|--------------------|--|
| Syntax | [no] sap <i>sap-id</i> |
| Context | config>service>vpls config>service>vpls>pbb>backbone-vpls |
| Description | This command configures attributes of a SAP on the B-VPLS service. |

mrouter-port

| | |
|--------------------|--|
| Syntax | [no] mrouter-port |
| Context | config>service>vpls>pbb>bvpls>sap>igmp-snooping config>service>vpls>pbb>bvpls>sdp>igmp-snooping config>service>vpls>pbb>bvpls>sap>mld-snooping config>service>vpls>pbb>bvpls>sdp>mld-snooping |
| Description | <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</p> |

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.

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.

If the **send-queries** command is enabled on this SAP or spoke-SDP, the **mrouter-port** parameter can not be set.

Default no mrouter-port

sdp

Syntax **[no] sdp** *sdp-id:vc-id*

Context config>service>vpls>pbb>backbone-vpls

Description This command configures attributes of a SDP binding on the B-VPLS service.

Parameters *sdp-id* — Specifies the SDP ID.

Values 1..17407

vc-id — Specifies the VC ID.

Values 1..4294967295

stp

Syntax **[no] stp**

Context config>service>Vpls>pbb>backbone-vpls

Description This command enables or disable STP through B-VPLS service.

force-qtag-forwarding

Syntax **[no] force-qtag-forwarding**

Context config>service>vpls 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 declares and registered on the egress SAP/Mesh-SDP/Spoke-SDP. The Group BMACs will be derived from the ISIDs using the procedure used in the PBB solution. The Group MAC = standard OUI with the last 24 bits being the ISID value. If the policy-name refers to a non-existing mrp-policy the command should return error. Changes to a mrp-policy are allowed and applied to the SAP/SDPs under which the policy is referenced.

Default no mrp-policy

send-bvpls-flush

Syntax [no] send-bvpls-flush {[all-from-me] | [all-but-mine]}

Context config>service>vpls

Description This command configures the BVPLS flush. If B-SDPs are used and MAC notification mechanism is turned on in the related BVPLS (MPLS use case), it makes sense to turn off the T-LDP MAC Flush.

mac-notification

Syntax mac-notification

Context config>service>pbb

Description This command controls the settings for the MAC notification messages.

interval

Syntax [no] interval *value*

Context config>service>pbb>mac-notification

Description This command controls the frequency of subsequent MAC notification messages.

Default 100 ms

Parameters *value* — Specifies the frequency of subsequent MAC notification messages.

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.

count

Syntax [no] count *value*

Context config>service>pbb>mac-notification

Description This command configures how often MAC notification messages are sent.

Parameters *value* — Specifies, in seconds, how often MAC notification messages are sent.

Values 1-10

Default 3

epipe

Syntax **epipe** *service-id* **customer** *customer-id* [*vpn* *vpn-id*] [**vc-switching**] [**create**]
epipe *service-id*
no epipe *service-id*

Context config>service

Description This command configures an Epipe service instance. This command is used to configure a point-to-point epipe service. An Epipe connects two endpoints defined as Service Access Points (SAPs). Both SAPs may be defined in one ESS-Series or they may be defined in separate ESS-Series 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 ESS-Series 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.

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 ESS-Series 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

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.

tunnel

| | |
|--------------------|---|
| Syntax | tunnel <i>service-id</i> backbone-dest-mac { <i>mac-name</i> <i>ieee-mac</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 | <p><i>service-id</i> — Specifies the B-VPLS service for the PBB tunnel associated with this service.</p> <p>Values 1 — 2147483648</p> <p>backbone-dest-mac {<i>mac-name</i> <i>ieee-mac</i>} — Specifies the backbone destination MAC-address for PBB packets.</p> <p>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.</p> <p>Values 0 — 16777215</p> |

PBB Show Commands

eth-cfm

| | |
|--------------------|--|
| Syntax | eth-cfm |
| Context | show |
| Description | This command displays 802.1ag CFM information. |

association

| | |
|--------------------|--|
| Syntax | association [<i>ma-index</i>] [detail] |
| Context | show>eth-cfm |
| Description | Shows association information. |
| Parameters | <i>ma-index</i> — Specifies the MA index value. |

Values 1 — 4294967295

detail — Displays all association detail.

| | |
|---------------|--|
| Output | <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 port [<i>port-id</i> > [vlan <i>vlan-id</i>]] [level 0..7] [direction up down] cfm-stack-table sdp [<i>sdp-id[:vc-id]</i> >] [level 0..7] [direction up down] cfm-stack-table virtual [<i>service-id</i>] [level 0..7] | | | | | | | | | | | | | | | | | | | | |
| Context | show>eth-cfm | | | | | | | | | | | | | | | | | | | | |
| Description | Summarizes all MEPs/MIPs. | | | | | | | | | | | | | | | | | | | | |
| Parameters | <p><i>port-id</i> — Displays information about the specified port.</p> <table><tr><td>Values</td><td>port-id</td><td>slot/mda/port[.channel]</td></tr><tr><td></td><td>lag-id</td><td>lag-id</td></tr><tr><td></td><td>lag</td><td>keyword</td></tr><tr><td></td><td>id</td><td>1 — 800</td></tr></table> <p><i>sdp-id[:vc-id]</i> — Specifies an existing SDP and VC ID.</p> <table><tr><td>Values</td><td>1 — 17407</td></tr></table> <p><i>vlan-id</i> — Specifies the VLAN ID.</p> <table><tr><td>Values</td><td>0 — 4094</td></tr></table> <p><i>level</i> — Specifies the level.</p> <table><tr><td>Values</td><td>0 — 7</td></tr></table> <p>direction up down — Indicates the direction in which the maintenance association (MEP or MIP) faces on the bridge port.</p> <p>down — Displays continuity check information configured away from the MAC relay entity.</p> <p>up — Displays continuity check information configured toward the MAC relay entity.</p> <p><i>service-id</i> — Specifies information about the specified service ID.</p> <table><tr><td>Values</td><td>1 — 2147483648</td></tr></table> | Values | port-id | slot/mda/port[.channel] | | lag-id | lag-id | | lag | keyword | | id | 1 — 800 | Values | 1 — 17407 | Values | 0 — 4094 | Values | 0 — 7 | Values | 1 — 2147483648 |
| Values | port-id | slot/mda/port[.channel] | | | | | | | | | | | | | | | | | | | |
| | lag-id | lag-id | | | | | | | | | | | | | | | | | | | |
| | lag | keyword | | | | | | | | | | | | | | | | | | | |
| | id | 1 — 800 | | | | | | | | | | | | | | | | | | | |
| Values | 1 — 17407 | | | | | | | | | | | | | | | | | | | | |
| Values | 0 — 4094 | | | | | | | | | | | | | | | | | | | | |
| Values | 0 — 7 | | | | | | | | | | | | | | | | | | | | |
| Values | 1 — 2147483648 | | | | | | | | | | | | | | | | | | | | |

Sample Output

```
*A:alcag1-R6# show eth-cfm cfm-stack-table
=====
CFM SAP Stack Table
=====
Sap          Level Dir  Md-index  Ma-index  Mep-id  Mac-address
-----
1/2/9:5      4      Up    1         1         51      00:ae:ae:ae:ae:ae
=====
CFM SDP Stack Table
=====
Sdp          Level Dir  Md-index  Ma-index  Mep-id  Mac-address
-----
No Matching Entries
=====
*A:alcag1-R6#
```

domain

| | |
|--------------------|---|
| Syntax | domain [<i>md-index</i>] [association <i>ma-index</i> all-associations [detail]] |
| Context | show>eth-cfm>domain |
| Description | This command displays domain information. |
| Parameters | <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 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
Ma-index      : 1                      Admin          : Enabled
MepId         : 51                   CCM-Enable     : Enabled
IfIndex       : 38043648             PrimaryVid     : 5
FngState      : fngReset
LowestDefectPri : allDef              HighestDefect   : none
Defect Flags  : None
Mac Address   : 00:ae:ae:ae:ae:ae     CcmLtmPriority  : 7
CcmTx         : 775                  CcmSequenceErr : 0
CcmLastFailure Frame:
None
XconCcmFailure Frame:
None
*A:alcag1-R6#
```

mep

| | |
|--------------------|--|
| Syntax | mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [loopback] [linktrace] |
| Context | show>eth-cfm>domain |
| Description | This command displays Maintenance Endpoint (MEP) information. |
| 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 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

```

```

=====
*A:Dut-B# show service pbb mac-name test detail
=====
Services Using MAC name='test' addr='00:03:03:03:03:02'
=====
Svc-Id ISID
-----
501 501
-----
Number of services: 1
=====
*A:Dut-B#

```

id

| | |
|--------------------|---|
| Syntax | id <i>service-id</i> |
| Context | show>service |
| Description | 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 00:03:03:03:02
-----
*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 [<i>mrp-policy</i>] mrp-policy <i>mrp-policy</i> [association] mrp-policy <i>mrp-policy</i> [entry <i>entry-id</i>] |
| Context | show>service |
| Description | This command displays MRP policy information. |
| Parameters | <i>mrp-policy</i> — Specifies the MRP policy. <div> Values 32 chars max </div> <i>entry-id</i> — Specifies the entry ID. <div> Values 1..65535 </div> |

mmrp

| | |
|--------------------|--|
| Syntax | mmrp mac [<i>ieee-address</i>] |
| Context | show>service>id |
| Description | This command displays information on MACs. If a MAC address is specified, information will be displayed relevant to the specific group. No parameter will display information on all group MACs on a server. |

Parameters *ieee-address* — 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

| | |
|--------------------|--|
| Syntax | adjacency [detail] |
| Context | show>service>id>spb |
| Description | This command displays SPB adjacency information. |
| Parameters | <i>detail</i> — Show detailed information. |
| Output | Sample Ouput |

```
=====
ISIS Adjacency
=====
System ID          Usage State Hold Interface          MT Enab
-----
Dut-B              L1    Up    19    sap:1/2/2:1.1          No
Dut-C              L1    Up    21    sap:1/2/3:1.1          No
-----
Adjacencies : 2
=====
```

base

| | |
|--------------------|---|
| Syntax | base |
| Context | show>service>id>spb |
| Description | This command displays SPB base information. |
| Output | Sample Ouput |

```
*A:Dut-A# show service id 100001 spb base
=====
Service SPB Information
=====
Admin State      : Up                Oper State      : Up
ISIS Instance    : 1024              FID             : 1
Bridge Priority   : 8                Fwd Tree Top Ucast : spf
Fwd Tree Top Mcast : st
Bridge Id        : 80:00.00:10:00:01:00:01
Mcast Desig Bridge : 80:00.00:10:00:01:00:01

=====
ISIS Interfaces
=====
Interface          Level CircID Oper State  L1/L2 Metric
-----
sap:1/2/2:1.1      L1    65536  Up         10/-
sap:1/2/3:1.1      L1    65537  Up         10/-
-----
Interfaces : 2
=====
FID ranges using ECT Algorithm
-----
1-99          low-path-id
```

```

100-100    high-path-id
101-4095   low-path-id
=====

```

database

Syntax **database**

Context show>service>id>spb

Description This command displays SPB database information.

Output **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     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               SvcId    SdpBind
-----
```

```

-----
500          1/1/20:500          502          502          1/1/20:502
=====

*A:Dut-A# show service id 100001 spb fid fdb
=====
Control service FDB information
=====
Fid          MacAddr          UCast Source          MCast Source
              Last Update          Last Update
-----
1            00:10:00:01:00:01  local                  local
              04/04/2012 15:11:24  04/04/2012 15:11:24
1            00:10:00:01:00:02  1/2/2:1.1             1/2/2:1.1
              04/04/2012 15:51:45  04/04/2012 15:51:45
1            00:10:00:01:00:03  1/2/3:1.1             1/2/3:1.1
              04/04/2012 15:51:56  04/04/2012 15:51:56
1            00:10:00:01:00:04  1/2/2:1.1             1/2/2:1.1
              04/04/2012 15:52:11  04/04/2012 15:52:11
-----

Entries found: 4
=====
*A:Dut-A# show service id 100001 spb fid mfib
=====
Control service MFIB information
=====
FID   MacAddr          ISID   Source          Last Update
-----
1     01:1E:83:00:27:11 10001  1/2/2:1.1       04/04/2012 15:51:45
              1/2/3:1.1       04/04/2012 15:51:56
              local        04/04/2012 15:42:44
100   01:1E:83:00:27:12 10002  1/2/2:1.1       04/04/2012 15:51:45
              1/2/3:1.1       04/04/2012 15:51:56
              local        04/04/2012 15:43:09
-----

Entries found: 6
=====

```

hostname

| | |
|--------------------|--|
| Syntax | hostname |
| Context | show>service>id>spb |
| Description | This command displays SPB system-id to hostname mapping. |
| Output | Sample Ouput |

```

*A:Dut-A# show service id 100001 spb hostname
=====
Hosts
=====
System Id          Hostname
-----
0000.00AA.AAAA     cses-B02
0000.00BB.BBBB     cses-B07
=====

```


interface

| | |
|--------------------|---------------------------------------|
| Syntax | interface |
| Context | show>service>id>spb |
| Description | This command displays SPB interfaces. |
| Output | Sample Ouput |

```
*A:Dut-A# show service id 100001 spb interface
=====
ISIS Interfaces
=====
Interface                               Level CircID  Oper State   L1/L2 Metric
-----
sap:1/1/20:500                         L1      65536   Up           10/-
-----
Interfaces : 1
=====
```

mfib

| | |
|--------------------|---|
| Syntax | mfib [group-mac <i>ieee-address</i>][isid <i>isid</i>] |
| Context | show service id <svcl> spb |
| Description | This command displays multicast forwarding data-base information. |
| Parameters | <i>group-mac</i> — Optional IEEE group MAC format: mac-address: xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx <i>isid</i> — Optional I-SID. Format: 0..16777215 |
| Output | Sample Ouput |

```
*A:Dut-A# show service id 100001 spb mfib
=====
User service MFIB information
=====
MacAddr          ISID      Status
-----
01:1E:83:00:27:11 10001     Ok
-----
Entries found: 1
=====
```

routes

| | |
|----------------|---------------------|
| Syntax | routes |
| Context | show>service>id>spb |

Description This command displays SPB route information.

Output Sample Output

```
*A:Dut-A# show service id 100001 spb routes
=====
MAC Route Table
=====
Fid  MAC                               Ver.  Metric
    NextHop If                      SysID
-----
Fwd Tree: unicast
-----
1    00:10:00:01:00:02                10    10
    sap:1/2/2:1.1                    Dut-B
1    00:10:00:01:00:03                10    10
    sap:1/2/3:1.1                    Dut-C
1    00:10:00:01:00:04                10    20
    sap:1/2/2:1.1                    Dut-B
100  00:10:00:02:00:02                10    10
    sap:1/2/2:1.1                    Dut-B
100  00:10:00:02:00:03                10    10
    sap:1/2/3:1.1                    Dut-C
100  00:10:00:02:00:04                10    20
    sap:1/2/3:1.1                    Dut-C

Fwd Tree: multicast
-----
1    00:10:00:01:00:02                10    10
    sap:1/2/2:1.1                    Dut-B
1    00:10:00:01:00:03                10    10
    sap:1/2/3:1.1                    Dut-C
1    00:10:00:01:00:04                10    20
    sap:1/2/2:1.1                    Dut-B
100  00:10:00:02:00:02                10    10
    sap:1/2/2:1.1                    Dut-B
100  00:10:00:02:00:03                10    10
    sap:1/2/3:1.1                    Dut-C
100  00:10:00:02:00:04                10    20
    sap:1/2/3:1.1                    Dut-C
-----
No. of MAC Routes: 12
=====

ISID Route Table
=====
Fid  ISID                               Ver.
    NextHop If                      SysID
-----
1    10001                             10
    sap:1/2/2:1.1                    Dut-B
    sap:1/2/3:1.1                    Dut-C
100  10002                             10
    sap:1/2/2:1.1                    Dut-B
    sap:1/2/3:1.1                    Dut-C
-----
No. of ISID Routes: 2
=====
A:Dut-A# show service id spb fate-sharing
```

```

=====
User service fate-shared sap/sdp-bind information
=====
Control   Control Sap/      FID      User      User Sap/
SvcId     SdpBind
-----
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 |
|---------------|----------------|

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

| | |
|--------------------|--|
| Syntax | counters |
| Context | clear>service>statistics>id |
| Description | This command clears all traffic queue counters associated with the service ID. |

mesh-sdp

| | | | | | | | |
|--------------------|--|----------------|-----------|---------------|-----------|---------------|----------------|
| Syntax | mesh-sdp <i>sdp-id[:vc-id]</i> { all counters stp mrp } | | | | | | |
| Context | clear>service>statistics>id | | | | | | |
| Description | This command clears the statistics for a particular mesh SDP bind. | | | | | | |
| Parameters | <i>sdp-id</i> — Specifies the SDP ID for which to display information. <table><tr><td>Default</td><td>All SDPs.</td></tr><tr><td>Values</td><td>1 — 17407</td></tr></table> <i>vc-id</i> — Displays information about the virtual circuit identifier. <table><tr><td>Values</td><td>1 — 4294967295</td></tr></table> all — Clears all queue statistics and STP statistics associated with the SDP. counters — Clears all queue statistics associated with the SDP. stp — Clears all STP statistics associated with the SDP. mrp — Clears all MRP statistics associated with the SDP. | Default | All SDPs. | Values | 1 — 17407 | Values | 1 — 4294967295 |
| Default | All SDPs. | | | | | | |
| Values | 1 — 17407 | | | | | | |
| Values | 1 — 4294967295 | | | | | | |

mrp

| | |
|--------------------|--|
| Syntax | mrp |
| Context | clear>service>statistics>id |
| Description | This command clears all MRP statistics for the service ID. |

spoke-sdp

| | |
|----------------|---|
| Syntax | spoke-sdp <i>sdp-id[: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 | <p><i>sdp-id</i> — The spoke SDP ID for which to clear statistics.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset.</p> <p>Values 1 — 4294967295</p> <p>all — Clears all queue statistics and STP statistics associated with the SDP.</p> <p>counters — Clears all queue statistics associated with the SDP.</p> <p>stp — Clears all STP statistics associated with the SDP.</p> <p>l2pt — Clears all L2PT statistics associated with the SDP.</p> <p>mrp — Clears all MRP statistics associated with the SDP.</p> |

sap

| | |
|--------------------|---|
| Syntax | sap <i>sap-id</i> { all counters stp l2pt mrp } |
| Context | clear>service>statistics>id |
| Description | This command clears statistics for the SAP. |
| Parameters | <p><i>sap-id</i> — The SAP ID for which to clear statistics.</p> <p>all — Clears all queue statistics and STP statistics associated with the SAP.</p> <p>counters — Clears all queue statistics associated with the SAP.</p> <p>stp — Clears all STP statistics associated with the SAP.</p> <p>l2pt — Clears all L2PT statistics associated with the SAP.</p> <p>mrp — Clears all MRP statistics associated with the SAP.</p> |

stp

| | |
|--------------------|---|
| Syntax | stp |
| Context | clear>service>statistics>id |
| Description | Clears all spanning tree statistics for the service ID. |

PBB Debug Commands

mrp

| | |
|--------------------|--|
| Syntax | [no] mrp |
| Context | debug>service>id |
| Description | This command enables and configures MRP debugging. |

all-events

| | |
|--------------------|---|
| Syntax | all-events |
| Context | debug>service>id>mrp |
| Description | This command enables MRP debugging for the applicant, leave all, periodic and registrant state machines and enables debugging of received and transmitted MRP PDUs. |

applicant-sm

| | |
|--------------------|--|
| Syntax | [no] applicant-sm |
| Context | debug>service>id>mrp |
| Description | This command enables debugging of the applicant state machine. The no form of the command disables debugging of the applicant state machine. |

leave-all-sm

| | |
|--------------------|--|
| Syntax | [no] leave-all-sm |
| Context | debug>service>id>mrp |
| Description | This command enables debugging of the leave all state machine. The no form of the command disables debugging of the leave all state machine. |

mmrp-mac

| | |
|----------------|--|
| Syntax | [no] mmrp-mac <i>ieee-address</i> |
| 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 | <i>ieee-address</i> — 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 1469 for command syntax. |

sdp

| | | | | | | | |
|--------------------|---|----------------|-----------|---------------|-----------|---------------|----------------|
| Syntax | [no] sdp <i>sdp-id:vc-id</i> | | | | | | |
| 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> — Specifies the SDP ID for which to display information. <table><tr><td>Default</td><td>All SDPs.</td></tr><tr><td>Values</td><td>1 — 17407</td></tr></table> <i>vc-id</i> — Displays information about the virtual circuit identifier. <table><tr><td>Values</td><td>1 — 4294967295</td></tr></table> | Default | All SDPs. | Values | 1 — 17407 | Values | 1 — 4294967295 |
| Default | All SDPs. | | | | | | |
| Values | 1 — 17407 | | | | | | |
| Values | 1 — 4294967295 | | | | | | |

Ethernet Virtual Private Networks (EVPN)

In This Chapter

This chapter provides information about Ethernet Virtual Private Networks (EVPN), process overview, and implementation notes.

Topics in this chapter include:

- [Overview on page 1216](#)
- [EVPN for VXLAN Tunnels in a Layer-2 DC GW on page 1217](#)
- [EVPN for VXLAN Tunnels in a Layer-2 DC on page 1219](#)
- [EVPN for VXLAN Tunnels in a Layer 3 DC on page 1220](#)
- [EVPN for VXLAN Tunnels in a Layer 3 DC on page 1222](#)
- [VXLAN on page 1224](#)
- [BGP-EVPN Control Plane for VXLAN Overlay Tunnels on page 1235](#)
- [EVPN for VXLAN in VPLS Services on page 1239](#)
- [EVPN for VXLAN in R-VPLS Services on page 1252](#)
- [EVPN for VXLAN in IRB Backhaul R-VPLS Services and IP Prefixes on page 1254](#)
- [EVPN for VXLAN in EVPN Tunnel R-VPLS Services on page 1258](#)
- [Interaction of EVPN and VXLAN with Existing VPLS Features on page 1263](#)
- [Interaction of EVPN and VXLAN with Existing VPRN Features on page 1264](#)
- [Routing Policies for BGP EVPN IP Prefixes on page 1265](#)

Overview

EVPN is an IETF technology per RFC7432 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 setup the flooding trees are distributed by BGP.

EVPN is defined to fill the gaps of other L2VPN technologies like VPLS. EVPN's main objective 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 EVPN for VXLAN overlay tunnels, 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 sdw-bindings or SAPs are expected for the connectivity to the WAN. R-VPLS and VPRN connectivity to the WAN is also supported.

The 7x50 SR/ESS/XRS EVPN for VXLAN implementation is perfectly integrated in the Nuage Data Center architecture, where the 7x50 plays the role of the Data Center Gateway (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

Figure 134 depicts the use of EVPN for VXLAN overlay tunnels on the 7x50 SR/ESS/XRS when it is used as a layer-2 DC GW.

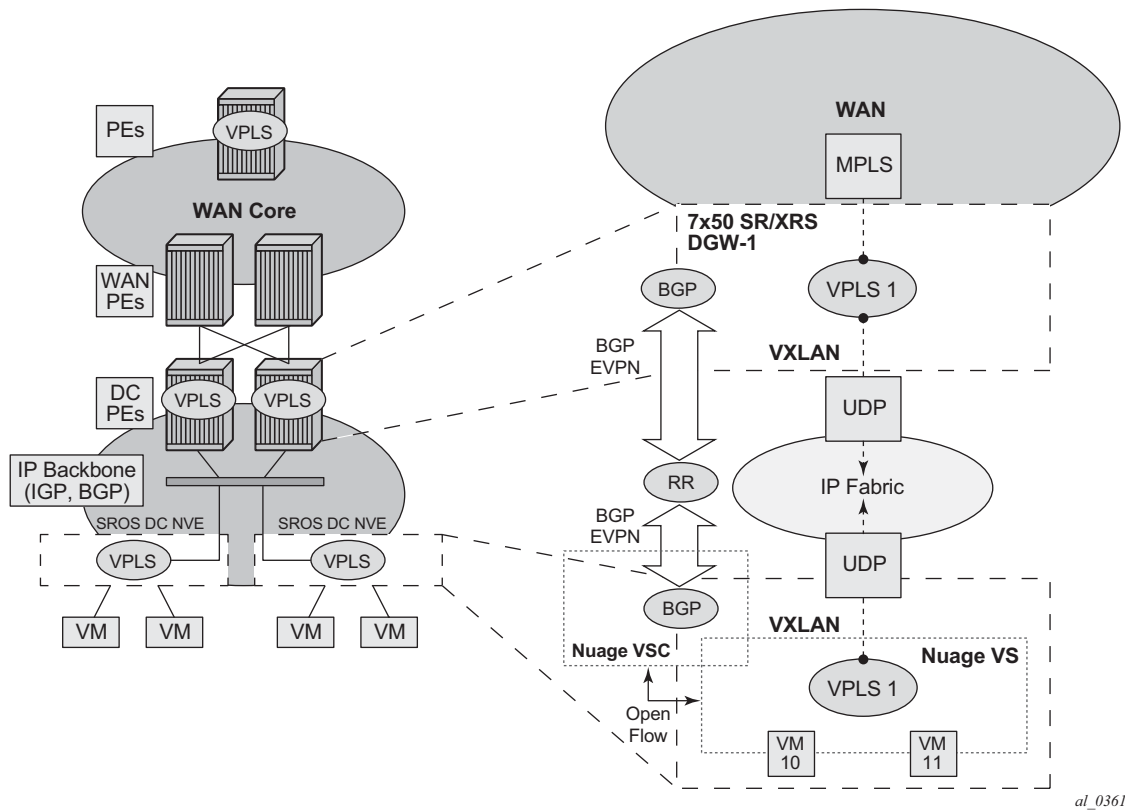


Figure 134: Layer-2 DC PE with VPLS to the WAN

DC providers require a DC GW solution that can extend tenant subnets to the WAN. Due to the shortcomings explained in the previous section, customers will be deploying NVO3-based solutions in the DC, where EVPN is the standard control plane and VXLAN is the data plane encapsulation becoming predominant. The Alcatel-Lucent DC architecture, a.k.a Nuage, uses EVPN and VXLAN as the control and data plane solutions for layer-2 connectivity within the DC and so does the 7x50 SR OS.

While EVPN with VXLAN will be used within the DC, most of the Service Providers use VPLS and H-VPLS as the solution to extend layer-2 VPN connectivity. Figure 134 above illustrates 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

Figure 135 depicts 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.

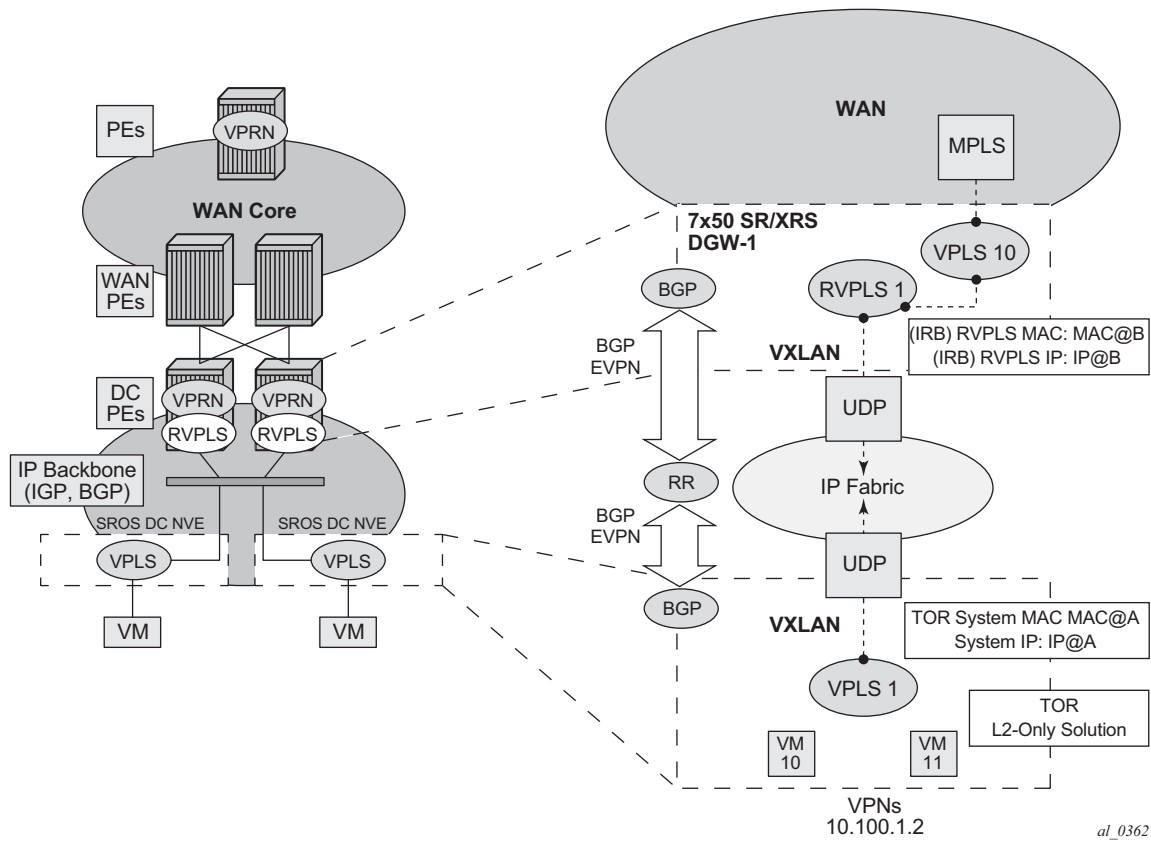


Figure 135: GW IRB on the DC PE for an L2 EVPN/VXLAN DC

In some cases, the DC GW must provide a Layer 3 'default gateway' function to all the hosts in a given 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.

EVPN for VXLAN Tunnels in a Layer 3 DC

Figure 136 depicts the use of EVPN for VXLAN tunnels on the 7x50 SR/XRS, when the DC provides distributed layer-3 connectivity to the DC tenants.

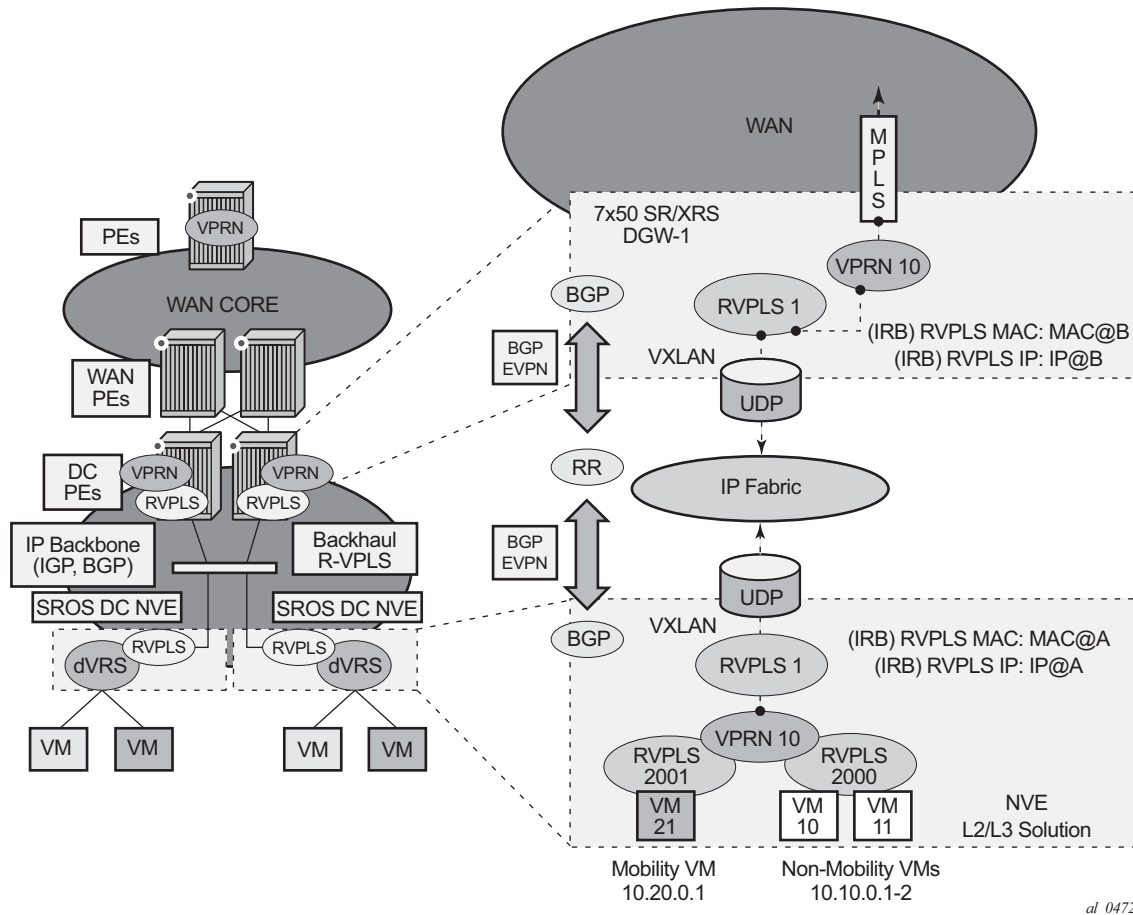


Figure 136: GW IRB on the DC PE for an L3 EVPN/VXLAN DC

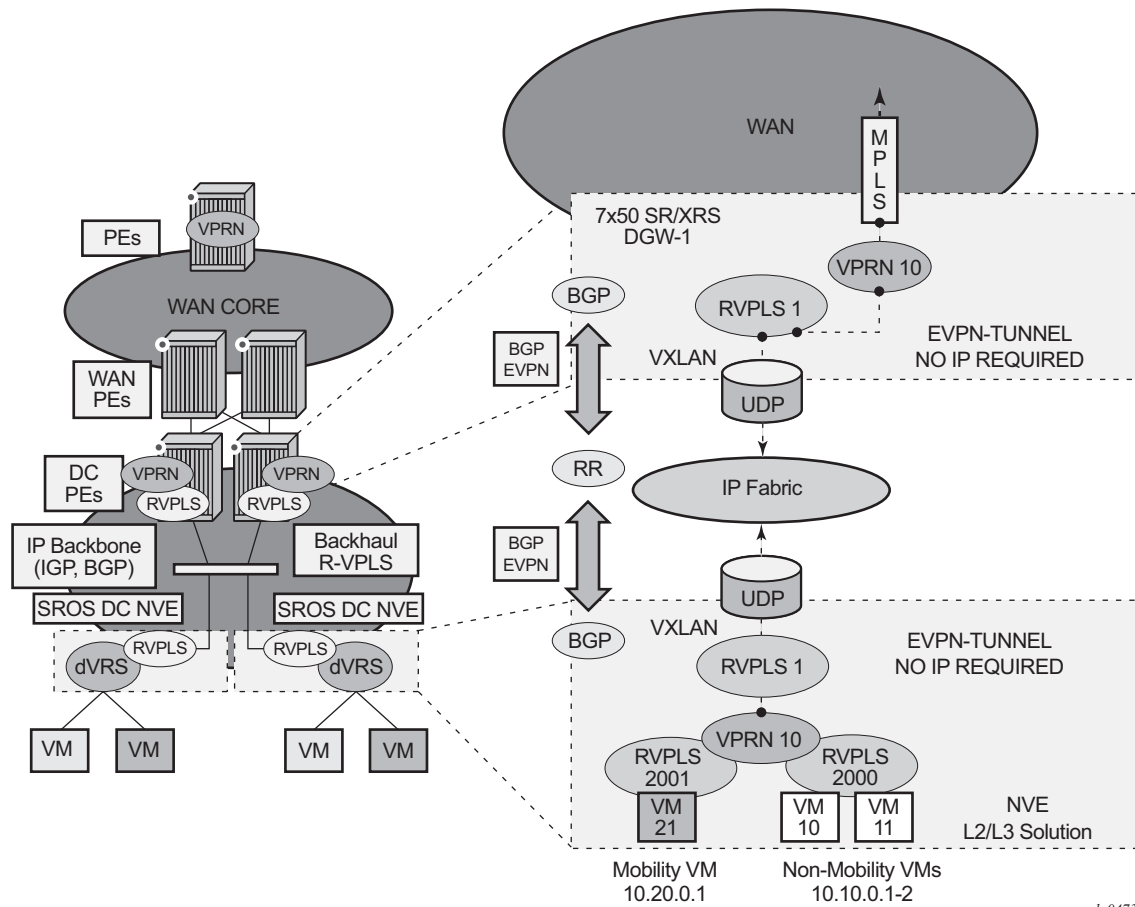
Each tenant will have several subnets for which each DC Network Virtualization Edge (NVE) provides intra-subnet forwarding. An NVE may be a Nuage VSG, VSC/VRS or any other NVE in the market supporting the same constructs, and each subnet normally corresponds to an R-VPLS. For example, in Figure 136 subnet 10.20.0.0 corresponds to R-VPLS 2001 and subnet 10.10.0.0 corresponds to R-VPLS 2000. In this model, the NVE provides inter-subnet forwarding too, by connecting all the local subnets to a VPRN instance. When the tenant requires L3 connectivity to its WAN IP-VPN, a VPRN is defined in the DC GWs, which connects the tenant to the WAN. That VPRN instance will be connected to the VPRNs in the NVEs by means of an IRB (Integrated

Routing and Bridging) backhaul R-VPLS. This IRB backhaul R-VPLS provides a scalable solution since 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

Figure 137 depicts 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 137: EVPN-Tunnel GW IRB on the DC PE for an L3 EVPN/VXLAN DC

The solution outlined in the previous section provides a scalable IRB backhaul R-VPLS service where all the VPRN instances for a given 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-binds 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 bring the following benefits to EVPN-VXLAN IRB backhaul R-VPLS services:

- Easier and simpler 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 that for IPv6 interfaces there is no need to provision 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 of that, the number of VXLAN binds in IRB backhaul R-VPLS services with EVPN-tunnel can be much higher.

This optimization is fully supported by the 7x50.

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, since VM MACs are only learnt 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 its 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 138](#) outlines the VXLAN encapsulation supported by the Alcatel-Lucent's implementation.

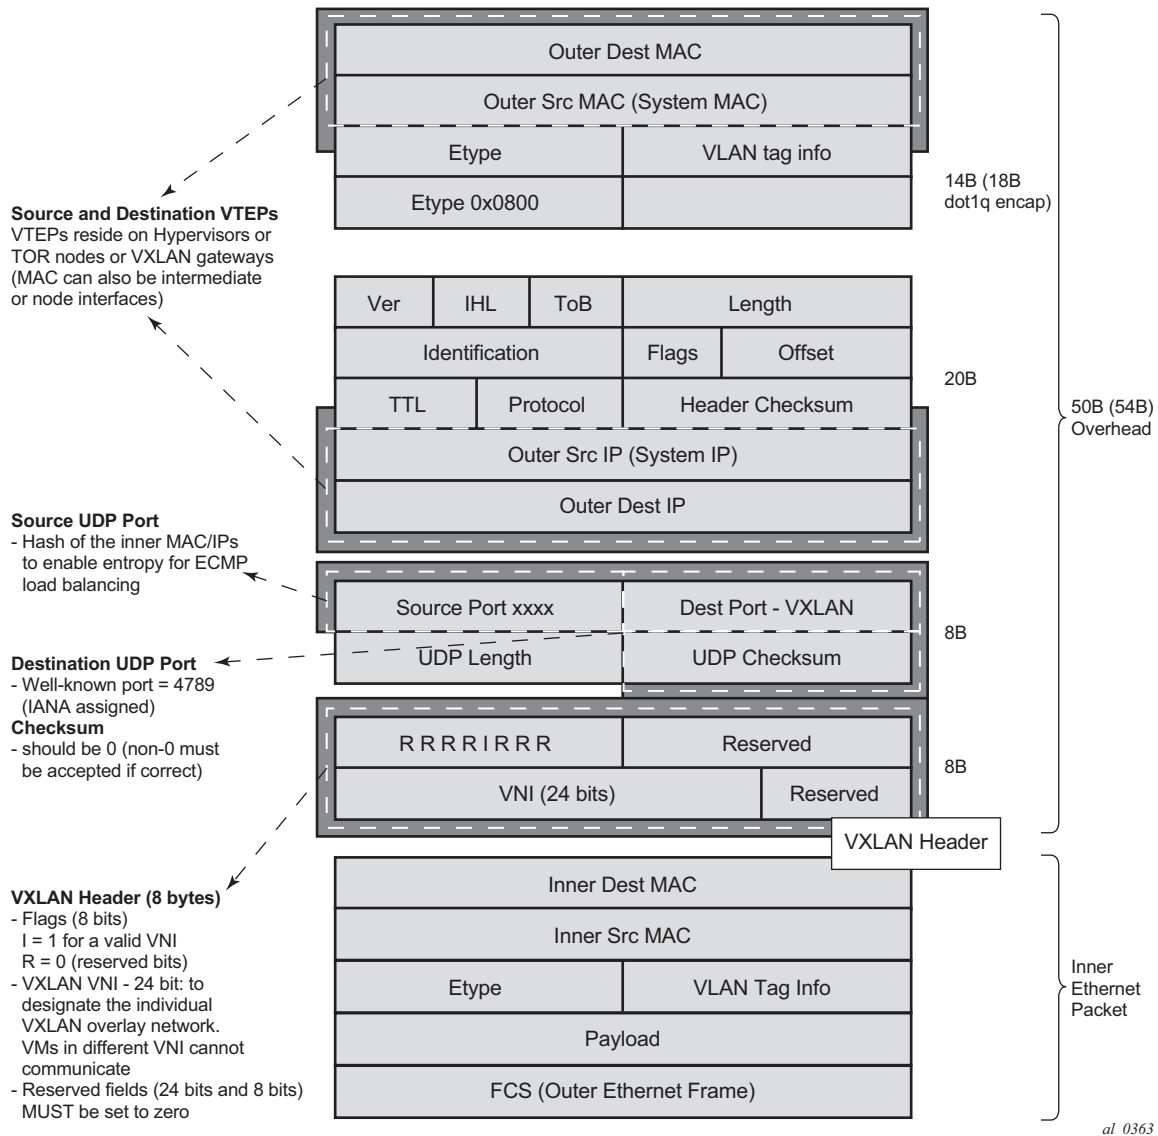


Figure 138: VXLAN Frame Format

As shown in Figure 138, VXLAN encapsulates the inner Ethernet frames into a 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. In other words, these fields will be populated with the PE’s system IP and chassis MAC address. Note that the source MAC address will be changed on all the IP hops along the path, as usually is 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 FDB. Note that all the remote MACs will be learnt by EVPN BGP and associated to a remote VTEP address and VNI.

Some considerations related to the support of VXLAN on the 7x50 are listed below:

- 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 since the source UDP port of each tunneled packet is derived from a hash of the customer inner packet. The following considerations must be taken into account:

- 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 RVPLS 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 50B (54B 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) and then the DSCP to yield a FC/profile.
- QoS control of broadcast, unknown and multicast (BUM) traffic received on VXLAN bindings is possible by separately redirecting these traffic types to policers within an FP ingress network queue group using 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, since 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 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 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. Care must be taken to 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 two forwarding databases.

The echo responder will terminate the VXLAN OAM frame and will take the appropriate response action and include relevant return codes. By default, the response is sent back using the IP network as an IPv4 UDP response. The operator can chose to override this default by changing the **reply-mode** to **overlay**. The overlay return mode will force the responder to use the VTEP connection representing the source IP and source VTEP. If a return overlay is not available, the echo response will be dropped by the responder.

Support is included for;

- IPv4 VTEP
- Optional specification of the outer UDP Source helps downstream network elements along the path with ECMP to hash to flow to the same path
- Optional configuration of the inner IP information to help 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 a 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 requires 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 the remote peer takes to process that packet. These reported values may not be representative of network delay.

The following sample 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
```

VXLAN Ping

```
0.1 send-count 10
vxlan-ping destination vxlan-id 31 ip-address 1.1.1.31 reply-mode udp interval 0.1s count
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-
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 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.

```

IGMP-snooping on VXLAN

```
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.
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 considerations must be taken into account when configuring IGMP-snooping in EVPN-VXLAN VPLS services:

- There is additional configuration command to enable IGMP-snooping on VXLAN: config>service>vpls>igmp-snooping no shutdown will enable the feature in the VPLS service.
- Note that the VXLAN bindings only support basic IGMP-snooping functionality. Features configurable under SAPs or SDP-bindings are not available for VXLAN. Since there is no specific IGMP-snooping settings for VXLAN bindings (static mrouters or send-queries, etc.), a given VXLAN binding will only become dynamic mrouter when it receives IGMP queries and will add a given multicast group to the MFIB when it receives an IGMP report for that group.
- The corresponding show/clear service id igmp-snooping are also available for VXLAN bindings. The following CLI commands show how the system displays IGMP-snooping information and statistics on VXLAN bindings:

```
*A:PE1# show service id 1 igmp-snooping port-db vxlan vtep 192.0.2.72 vni 1 detail

=====
IGMP Snooping VXLAN 192.0.2.72/1 Port-DB for service 1
=====
-----
IGMP Group 232.0.0.1
```

```

-----
Mode           : exclude           Type           : dynamic
Up Time        : 0d 19:07:05        Expires        : 137s
Compat Mode    : IGMP Version 3
V1 Host Expires : 0s                V2 Host Expires : 0s
-----
Source Address  Up Time      Expires  Type      Fwd/Blk
-----
No sources.
-----
IGMP Group 232.0.0.2
-----
Mode           : include           Type           : dynamic
Up Time        : 0d 19:06:39        Expires        : 0s
Compat Mode    : IGMP Version 3
V1 Host Expires : 0s                V2 Host Expires : 0s
-----
Source Address  Up Time      Expires  Type      Fwd/Blk
-----
10.0.0.232      0d 19:06:39  137s    dynamic   Fwd
-----
Number of groups: 2
=====

*A:PE1# show service id 1 igmp-snooping statistics vxlan vtep 192.0.2.72 vni 1

=====
IGMP Snooping Statistics for VXLAN 192.0.2.72/1 (service 1)
=====

```

| Message Type | Received | Transmitted | Forwarded |
|----------------------|----------|-------------|-----------|
| General Queries | 0 | 0 | 556 |
| Group Queries | 0 | 0 | 0 |
| Group-Source Queries | 0 | 0 | 0 |
| V1 Reports | 0 | 0 | 0 |
| V2 Reports | 0 | 0 | 0 |
| V3 Reports | 553 | 0 | 0 |
| V2 Leaves | 0 | 0 | 0 |
| Unknown Type | 0 | N/A | 0 |

```

-----
Drop Statistics
-----
Bad Length           : 0
Bad IP Checksum      : 0
Bad IGMP Checksum    : 0
Bad Encoding         : 0
No Router Alert      : 0
Zero Source IP       : 0
Wrong Version        : 0
Lcl-Scope Packets    : 0
Rsvd-Scope Packets   : 0

Send Query Cfg Drops : 0
Import Policy Drops   : 0
Exceeded Max Num Groups : 0
Exceeded Max Num Sources : 0
Exceeded Max Num Grp Srcs: 0
MCAC Policy Drops     : 0
=====

```

IGMP-snooping on VXLAN

```
*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
=====
```

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. Multi-homing can be supported though by using regular BGP-Multi-homing based on the L2VPN BGP address family.

Figure 139 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 Advertisement route

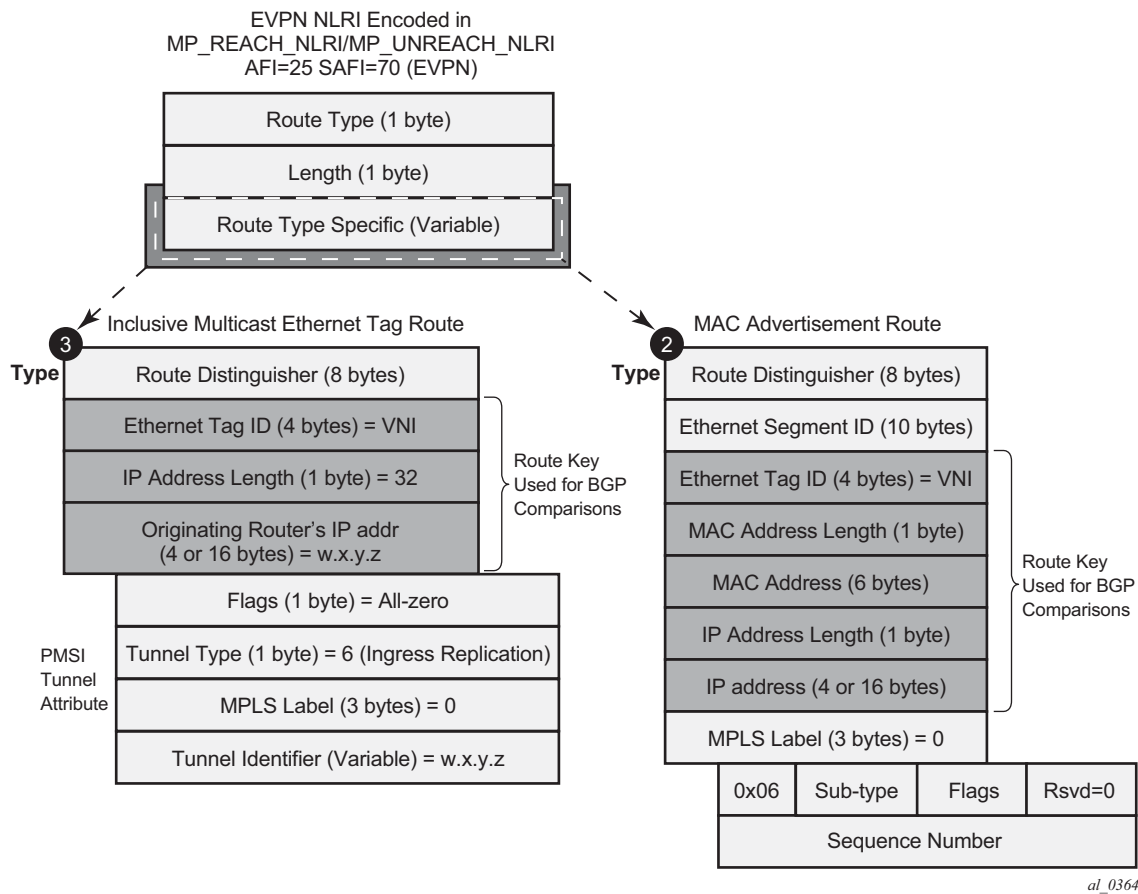


Figure 139: EVPN-VXLAN Required Routes and Communities

EVPN Route Type 3 – Inclusive Multicast Ethernet Tag Route

Route type 3 is used for setting up the flooding tree (BUM flooding) for a given 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 UP and uses the following fields and values:

- Route Distinguisher: Taken from the RD configured in the VPLS service within the BGP context.
- Ethernet Tag ID: Carries the VNI configured in the VPLS service. Only one VNI can be configured per VPLS service.
- 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 = 0
 - Tunnel end-point = equal to the originating IP address

EVPN Route Type 2 – MAC 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 configured in the VPLS service within the BGP context.
- Ethernet Segment Identifier (ESI): Value = 0:0:0:0:0:0:0:0
- Ethernet Tag ID: Carries the VNI configured in the VPLS service. Only one VNI can be configured per VPLS.
- 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 to the MAC being advertised with a length of 32 or 128 (for IPv6).
 - If the MAC address is the Unknown MAC route then 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.
 - At reception, any IPL value different from zero, 32 or 128 will make discard the route.
- MPLS Label = 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 taken into account for the route selection.

When EVPN is used in an IRB backhaul R-VPLS that connects all the VPRN instances for a given 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 140 shows the IP prefix route or route-type 5.

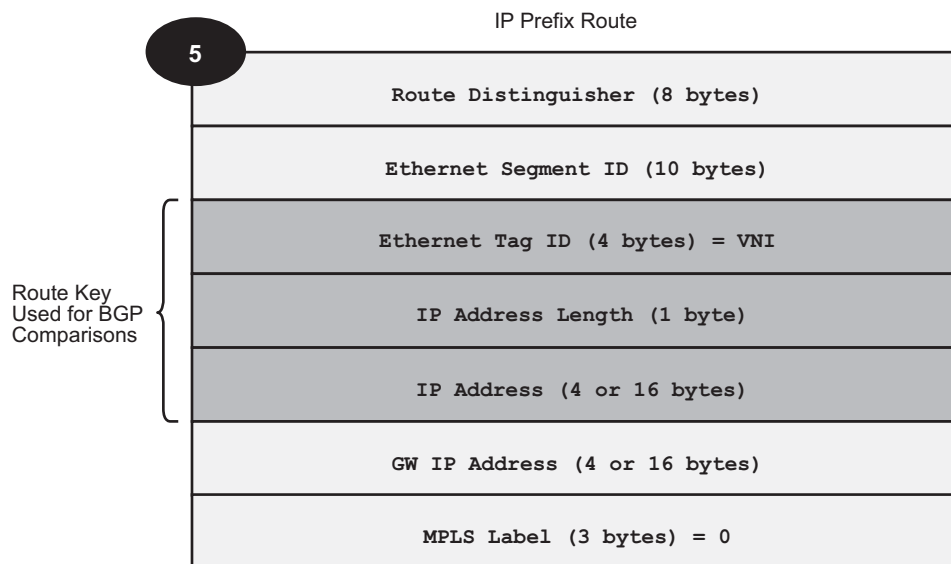


Figure 140: EVPN Route-Type 5

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:

- 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: Carries the VNI configured in the VPLS service. Only one VNI can be configured per VPLS service.
- 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: 0

All the routes in EVPN will be sent along with the RFC 5512 tunnel encapsulation extended community, with the tunnel type value set to VXLAN.

EVPN for VXLAN in VPLS Services

The EVPN-VXLAN service is modeled around the current VPLS objects and the additional VXLAN construct.

Figure 134 depicts 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 it requires connectivity (such VPLS instances can be instantiated in TORs, Nuage VRS, VSG, etc.). 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, etc).
- 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 learnt 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 learnt macs, conditional static macs and the unknown-mac-route. The unknown-mac-route will only be advertised under the conditions described above.

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.

Once the VPLS is properly 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 are modeled following an IP-VPN-like model, where no SDPs or SDP-binding objects are created or visible by 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    In Mcast List?  Oper State
-----
192.0.0.71             1              1             Yes              Up
192.0.0.72             1              1             Yes              Up
-----
Number of Egress VTEP, VNI : 2
=====
```

- The VXLAN bindings observe the VPLS split-horizon rule. This is done automatically without the need of any split-horizon configuration.
- Note that BGP Next-Hop Tracking for EVPN is fully supported. If the BGP next-hop for a given 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.

Once 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 to 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: a) no (BGP-MH) site is configured, or b) a site is configured AND the site is DF in the PE. Note that the unknown-mac-route will not be installed in the FDB (hence will not show up in the show service id x fdb detail command).
- Note that, although 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 to MAC addresses:

```
A:PE65# show service id 1000 fdb detail
=====
Forwarding Database, Service 1000
=====
```

| ServId | MAC | Source-Identifier | Type Age | Last Change |
|--------|-------------------|---------------------------|-------------|-------------------|
| 1000 | 00:00:00:00:00:01 | vxlan: 192.0.2.63:1063 | Evpn | 10/05/13 23:25:57 |
| 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: 192.0.2.63:1063 | EvpnS | 10/04/13 17:35:43 |

```
-----
No. of MAC Entries: 3
-----
Legend: L=Learned O=Oam P=Protected-MAC
=====
```

Resiliency and BGP Multi-Homing

The DC overlay infrastructure relies on IP tunneling, that is, VXLAN; the underlay IP layer sorts out, therefore, 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

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 all configured in one service in a given system. If that is the case, the following considerations apply:

- The configured bgp route-distinguisher and route-target are used by BGP for the two families, that is, evpn and l2vpn. If different import/export route targets are to be used per family, vsi-import/export policies must be used.

- The pw-template-binding command under bgp, does not have any effect on evpn or bgp-mh. It is only used for the instantiation of the bgp-ad spoke-sdps.
- If the same import/export route-targets are used in the two redundant DC GWs, VXLAN binding as well as a fec129 spoke-sdp binding will be established between the two DGWs, creating a loop. To avoid creating a loop, the 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, brings a significant number of benefits as described in draft-ietf-bess-evpn-overlay. There is however a potential issue that **SHOULD** 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 amount 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, hence 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 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 to an ingress sap/sdp-binding.

The use of the unknown-mac-route assumes the following:

1. a fully virtualized DC where all the MACs are control-plane learned, and learned previously to any communication (no legacy TORs or VLAN connected servers).
2. 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.
3. 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:

- no site is configured and the service is operationally UP or
- 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.

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<FmSymbol>®MAC entries learnt from the following sources:

- EVPN-received IP<FmSymbol>®MAC entries
- User-configured static IP<FmSymbol>®MAC
- Snooped dynamic IP<FmSymbol>®MAC entries (learnt 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 141 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.

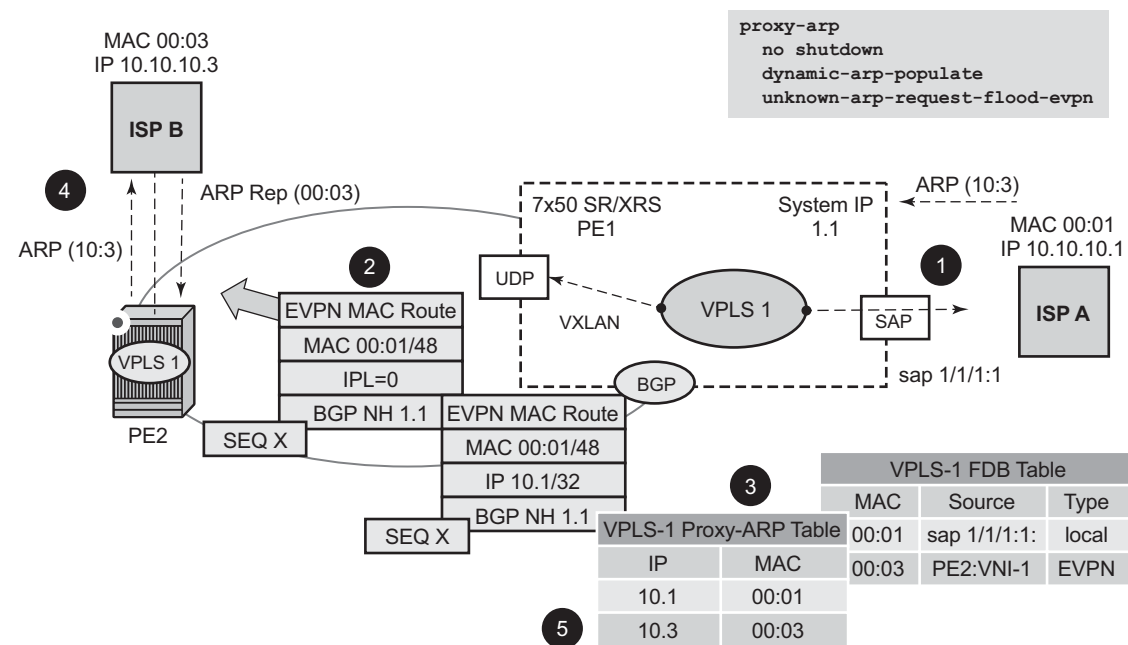


Figure 141: Proxy-ARP example usage in an EVPN Network

The PE1 in [Figure 141](#) 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 141](#) illustrates 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<FmSymbol>® 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<FmSymbol>@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 learnt in FDB at PE2 and advertised in EVPN.
 - MAC 00:03 and IP 10.3 are learnt 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<FmSymbol>@MAC 00:03) tables as mentioned in 5.

From that point on, 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 thus 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 that a static IP<FmSymbol>@MAC entry requires the addition of the MAC address to the FDB as either learnt 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 learnt on a given 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**.

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. In order to keep the active IP<FmSymbol>@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 given IP<FmSymbol>@MAC entry:

- Periodic refresh messages (ARP-requests or NS for a given 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 learnt 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 learnt for an existing IP. The objective of the confirm messages is to verify that a given 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 with 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:

- either provide the right R flag information in proxy-ND NA replies, or
- flood the received NA messages if it cannot provide the right 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.

The procedure to add the R flag to a given entry is as follows:

- Dynamic entries are learnt based on received NA messages. The R flag is also learnt and added to the proxy-ND entry so that the appropriate R flag is used in response to NS requests for a given IP.
- Static entries are configured as host or router as per the command **[no] static <ip-address> <ieee-address> {host | router}**.
- EVPN entries are learnt 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<FmSymbol>®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 in the following:

- 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 comm 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).

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 till:

- The MAC is flushed due to a local event (sap/sdp-binding associated to the MAC fails) or the reception of a remote update with better SEQ number (due to a mac flush at the remote 7x50) or
- 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 mac duplication parameters described above 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
```

Conditional Static MAC and Protection

The draft-ietf-bess-evpn-overlay defines the use of the sticky bit in the mac-mobility extended community to signal static mac addresses. These addresses must be protected in case there is an attempt to dynamically learn them in a different place in the EVPN-VXLAN VPLS service.

In the 7x50, any conditional static mac defined in an EVPN-VXLAN VPLS service will be advertised by BGP-EVPN as static address, that is, with the sticky bit set. An example of the configuration of a conditional static mac is shown below:

```

*A:PE63>config>service>vpls# info
-----
description "vxlan-service"
...
    sap 1/1/1:1000 create
    exit
    static-mac
        mac 00:ca:ca:ca:ca:00 create sap 1/1/1:1000 monitor fwd-status
    exit
    no shutdown

*A:PE64# show router bgp routes evpn mac hunt mac-address 00:ca:ca:ca:ca:00
...
=====
BGP EVPN Mac Routes
=====
Network          : 0.0.0.0/0
Nexthop           : 192.0.2.63
From              : 192.0.2.63
Res. Nexthop      : 192.168.19.1
Local Pref.       : 100
Aggregator AS     : None
Atomic Aggr.      : Not Atomic
AIGP Metric       : None
Connector         : None
Community         : target:65000:1000
Cluster           : No Cluster Members
Originator Id     : None
Flags             : Used Valid Best IGP
Route Source      : Internal
AS-Path           : No As-Path
EVPN type         : MAC
ESI               : 0:0:0:0:0:0:0:0:0
IP Address        : ::
Mac Address       : 00:ca:ca:ca:ca:00
Neighbor-AS       : N/A
Source Class      : 0
Interface Name    : NotAvailable
Aggregator        : None
MED               : 0
mac-mobility:Seq: 0/Static
Peer Router Id    : 192.0.2.63
Tag               : 1063
RD                : 65063:1000
Mac Mobility       : Seq:0
Dest Class        : 0
-----
Routes : 1
=====

```

Local static MACs or remote MACs with sticky bit are considered as ‘protected’. A packet entering a SAP / SDP-binding will be discarded if its source MAC addresses matches one of these ‘protected’ MACs.

EVPN for VXLAN in R-VPLS Services

[Figure 135](#) depicts 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 it requires connectivity (such VPLS instances can be instantiated in TORs, Nuage VRS, VSG, etc.). 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 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 listed below:
 - The 7750 will advertise its IRB MAC+IP in a route type 2 route and possibly the VRRP vMAC+vIP if it runs VRRP and it is the master. In both cases, the MACs will be advertised as static MACs, hence protected by the receiving PEs.
 - If the 7750 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 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 that 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-binding) 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. They will be added when the IP address for the same subnet is added.
- Note that 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 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 136 depicts 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 NVE's will send IP prefixes in EVPN that will be received by the DC GW and imported in the VPRN routing table.

Note that in order 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 done through the **bgp-evpn>ip-route-advertisement** command. This command must be explicitly enabled (it is disabled by default) and it is tied to the **allow-ip-int-binding** command required for R-VPLS. Note that by default local router interface host addresses are not advertised in EVPN. Should the user want 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
10.1.1.100/32                                      Local   Host     00h00m11s    0
      if                                           Y
=====
```

For the case illustrated in 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 mpls-gre
    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-binding
    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-binding
    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 that the 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 existing 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 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 don't 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 Global IPv6 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 137 shows an L3 connectivity model that optimizes the solution described in the previous section. 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 that **evpn-tunnel** can be enabled independently of **ip-route-advertisement** (although 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 mpls-gre
  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-binding
  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 given VPRN supports regular IRB backhaul R-VPLS services as well as EVPN tunnel R-VPLS services. Note that 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 a 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 ARP to get a MAC. If the 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 a 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, if 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. Normally 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.
 - Note that the IP prefixes with GW-MACs as next-hops are displayed in 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 shown in 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
```

EVPN for VXLAN in EVPN Tunnel R-VPLS Services

```
=====
BGP Router ID:10.20.1.1      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP EVPN IP-Prefix Routes
=====
-----
Original Attributes

Network      : N/A
Nexthop      : 10.20.1.2
From         : 10.20.1.2
Res. Nexthop : 192.168.19.1
Local Pref.  : 100
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
Flags        : Used Valid Best IGP
Route Source : Internal
AS-Path      : No As-Path
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  : N/A
Orig Validation: N/A
Source Class : 0

Interface Name : NotAvailable
Aggregator     : None
MED            : 0
Tag            : 1
Route Dist.    : 10.20.1.2:1
Dest Class     : 0

Modified Attributes

Network      : N/A
Nexthop      : 10.20.1.2
From         : 10.20.1.2
Res. Nexthop : 192.168.19.1
Local Pref.  : 100
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
Flags        : Used Valid Best IGP
Route Source : Internal
AS-Path      : 111
EVPN type    : IP-PREFIX
ESI          : N/A

Interface Name : NotAvailable
Aggregator     : None
MED            : 0
Tag            : 1
Route Dist.    : 10.20.1.2:1
Dest Class     : 0
```



```

Gateway Address: 00:00:01:00:01:02
Prefix          : 3.0.1.6/32          Route Dist.    : 10.20.1.2:1
MPLS Label      : 262140
Route Tag       : 0xb
Neighbor-AS     : 111
Orig Validation : N/A
Source Class    : 0                  Dest Class     : 0

```

```

-----
Routes : 1
=====

```

EVPN tunnel is also supported on IPv6 VPRN interfaces. When sending IPv6 prefixes from IPv6 interfaces, the GW-MAC in the route type 5 (IP-prefix route) is always zero. If no specific Global Address is configured on the IPv6 interface, the routes type 5 for IPv6 prefixes will always be sent using the Link Local Address as GW-IP. The following example output shows an IPv6 prefix received via BGP EVPN.

```
*A:PE71# show router 30 route-table ipv6
```

```

=====
IPv6 Route Table (Service: 30)
=====
Dest Prefix[Flags]                                Type   Proto   Age           Pref
      Next Hop[Interface Name]                      Metric
-----
300::/64                                           Local   Local    00h01m19s    0
      int-PE-71-CE-1                                0
500::1/128                                         Remote  BGP EVPN 00h01m20s   169
      fe80::da45:ffff:fe00:6a-"int-evi-301"          0
-----
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====

```

```
*A:PE71# show router bgp routes evpn ipv6-prefix prefix 500::1/128 hunt
```

```

=====
BGP Router ID:192.0.2.71      AS:64500      Local AS:64500
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
              l - leaked
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

```

```

=====
BGP EVPN IP-Prefix Routes
=====

```

```

-----
RIB In Entries
-----

```

```

Network      : N/A
Nexthop      : 192.0.2.69
From         : 192.0.2.69

```

EVPN for VXLAN in EVPN Tunnel R-VPLS Services

```
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
Peer Router Id    : 192.0.2.69
Flags             : Used Valid Best IGP
Route Source      : Internal
AS-Path           : No As-Path
EVPN type         : IP-PREFIX
ESI               : N/A
Tag               : 301
Gateway Address   : fe80::da45:ffff:fe00:*
Prefix            : 500::1/128
Route Dist.       : 192.0.2.69:301
MPLS Label       : 0
Route Tag         : 0
Neighbor-AS       : N/A
Orig Validation   : N/A
Source Class      : 0
Dest Class        : 0
Add Paths Send    : Default
Last Modified     : 00h41m17s
```

```
-----
RIB Out Entries
-----
-----
```

```
Routes : 1
=====
```

Interaction of EVPN and VXLAN with Existing VPLS Features

When trying to enable existing VPLS features in an EVPN-VXLAN enabled service, the following must be taken into consideration:

- I-VPLS/B-VPLS services are not supported. **bgp-evpn** or VXLAN cannot be enabled on those services.
 - In general, no 7x50-generated control packets will be sent to the VXLAN auto-bindings, except for ARP, VRRP, ping, and BFD.
 - **eth-cfm** (meps, vmeps, mips): This command can be configured and used in an EVPN-VXLAN VPLS service objects (service, saps and sdp-bindings). Although **vmeps** can be configured and used for tests to the WAN, **eth-cfm** tests will not work through VXLAN. This behavior is expected since no **eth-cfm** is supported in the DC and **eth-cfm** flooding to the DC NVEs is definitively not desired.
 - xSTP and M-VPLS services:
 - xSTP can be configured in **bgp-evpn** services. BPDUs will not be sent over the VXLAN bindings.
 - **bgp-evpn** is blocked in m-vpls services, however, a different m-vpls service can manage a **sap/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: they will only work for data plane learned mac addresses.
 - **mac-protect**: **mac-protect** cannot be used in conjunction with EVPN and VXLAN.
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 blocked for **bgp-evpn** services, that is: **mac-ping**, **mac-trace**, **mac-populate**, and **mac-purge**.

Interaction of EVPN and 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, the following must be taken into consideration:

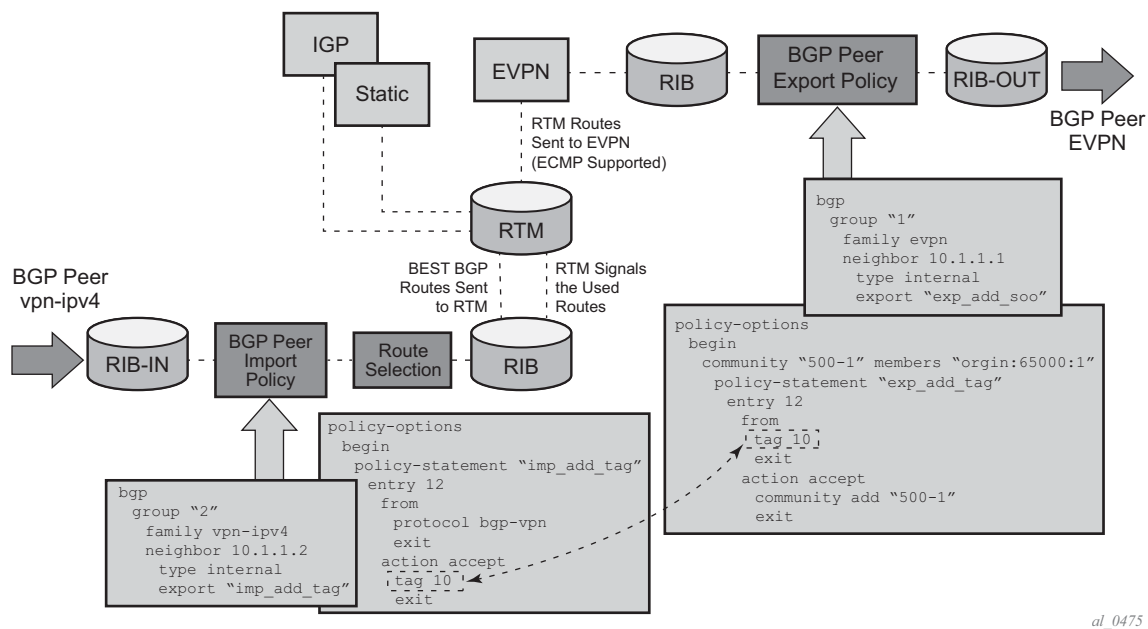
- The following commands are not supported:
 - arp-populate
 - authentication-policy
- Dynamic routing protocols such as IS-IS, RIP, and OSPF are not supported
- BFD is not supported on EVPN tunnel interfaces.

Routing Policies for BGP EVPN IP Prefixes

BGP routing policies are supported for IP prefixes imported or export through BGP-EVPN.

When applying routing policies to control the distribution of prefixes between EVPN and IP-VPN, the user must take into account 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 142 shows an example of how VPN-IPv4 routes are imported into the RTM (Routing Table Manager) and then passed onto EVPN for its own process. Note that 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.



al_0475

Figure 142: IP-VPN Import and EVPN Export BGP Workflow

Note that policy tags can be used to match EVPN IP prefixes that were learnt not only from BGP VPN-IPv4 but also from other routing protocols. Note that the tag range supported for each protocol is different:

<tag> : accepts in decimal or hex

[0x1..0xFFFFFFFF]H (for OSPF and IS-IS)

[0x1..0xFFFF]H (for RIP)

[0x1..0xFF]H (for BGP)

Figure 143 shows an example of the reverse workflow: routes imported from EVPN and exported from RTM to BGP VPN-IPv4:

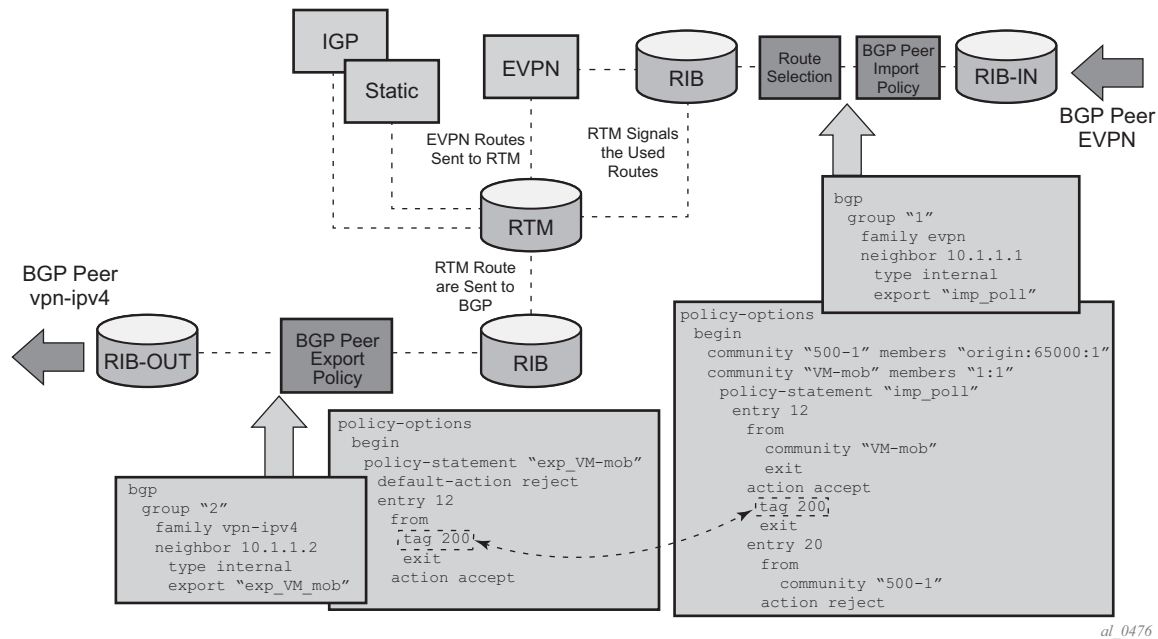


Figure 143: EVPN Import and IP-VPN Export BGP Workflow

Note that the above 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 we can then match on the BGP peer export policy or vsi-export policy for EVPN prefix routes.

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 in a user-friendly manner.

The VSD contains a multi-tenant service directory that supports role-based administration of users, compute, and network resources. It also manages network resource assignments such as IP addresses and ACLs.

In order 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 that the service must be pre-provisioned in the 7x50 (via CLI, SNMP or other supported interfaces), and VSD will only push a limited number of parameters into the configuration. This 7x50 – VSD integration model is known as “Static-Dynamic provisioning model” since only a few parameters are dynamically pushed by VSD, as opposed to “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 given 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) 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 right 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.

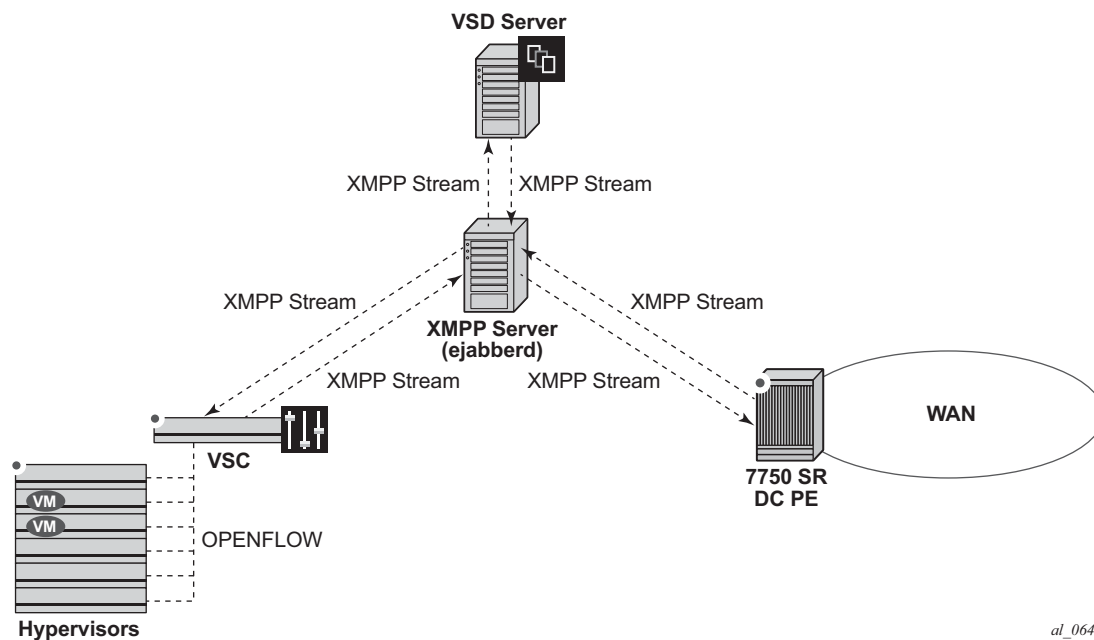
XMPP Interface on the DC GW

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 144](#) illustrates the basic XMPP architecture in the data center. Note that although a single XMPP server is represented, and 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 144: 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 at the XMPP Server.

In order 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.

- 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 will try to establish a TCP session with the XMPP server through the management interface first. If it fails to establish communication, the 7750 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**.

Once 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 "vsdl-hy"

=====
XMPP Server Table
=====
XMPP FQDN           : vsdl-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 will be used to uniquely identify the 7x50 in the VSD:

```
*B:Dut>config>system>vsd# info
-----
system-id "SR12U-46-PE"
-----
```

Once the above configuration is done, the 7x50 will subscribe to a VSD XMPP PubSub node to discover the available VSD servers. At this point, 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@vsdl-hy.alu-srpm.us/nuage* 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@vsdl-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 this 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 145 outlines the procedure for the attachment of the WAN services defined on the DC GW to the DC domains.

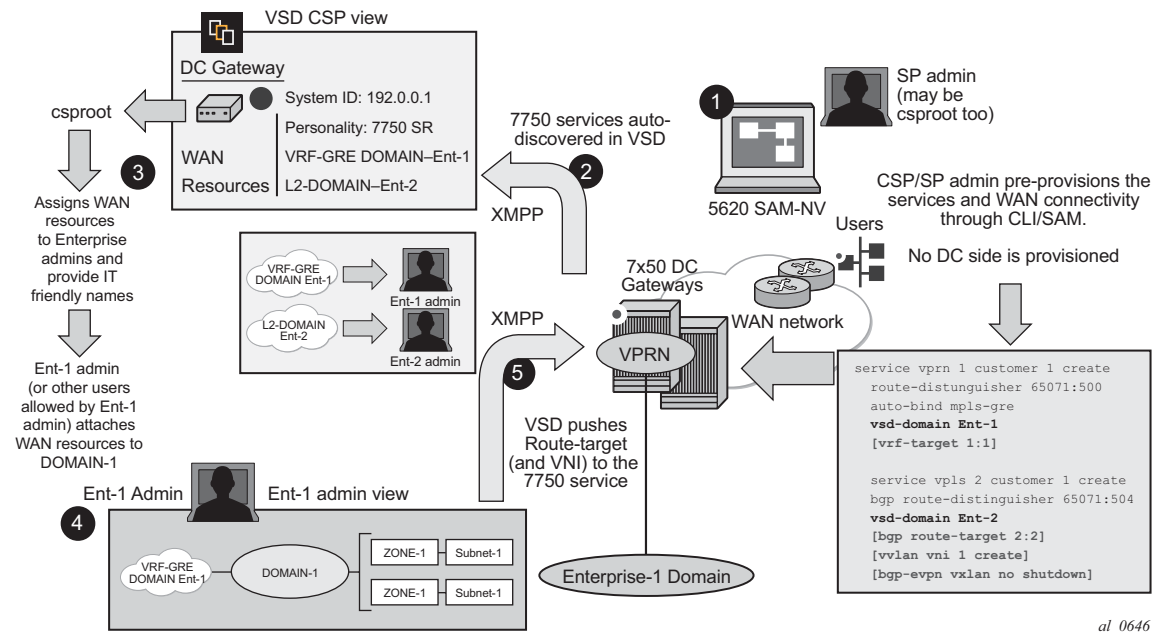


Figure 145: WAN Services Attachment Procedure

The Static-Dynamic VSD integration model can be summarized in the steps illustrated in Figure 145 and described here:

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 done 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 to **vsd-domain Ent-1**, which is a VRF-GRE domain.
- VPLS 2 – associated to **vsd-domain Ent-2**, which is an L2-DOMAIN.

Note that 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, 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 given 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 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 to 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 to 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 to 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 to 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 assign 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 to 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 listed below:

- **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 **vlan vni** and **route-target** that the 7x50 will add to the VPLS service. For the **route-target**, the system will check if 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 given below:

```
*A:PE# show router policy "policy-1"
    entry 900 # manual entry
        from
            as-path "null"
            family evpn
        exit
        action next-entry
            local-preference 500
        exit
    exit
    entry 1000 # automatic VSD-generated entry
```



```
    from
        community "_VSD_1"
        family evpn
    exit
    action accept
    exit
exit
entry 2000 # automatic VSD-generated entry
    from
        family evpn
    exit
    action accept
        community add "_VSD_1"
    exit
exit
```

VSD-domain Type L2-DOMAIN-IRB

L2-DOMAIN-IRB VSD-domains will be associated to 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 to 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 if the VPRN service has a configured policy:

- If there is no **vrf-import** policy, the received dynamic route-target will be added in the `vpn>` 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 that in case a **vrf-import** policy is used, the user will provision the WAN **route-target** statically in a **vrf-export** policy (this **route-target** will be used for the routes advertised to the DC as well).

An example of the auto-generated entry is given 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 to 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 listed below:

- **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 mpls-gre
    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 if the VPLS service has a configured policy and will push the **route-target** either in the service context or the **vsi-import/export** policies, as described in the section for L2-DOMAINS.

The following commands help showing 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
=====

```

Configuring a EVPN Service with CLI

This section provides information to configure VPLS using the command line interface.

Topics in this section include:

- [EVPN Configuration Examples on page 1282](#)

EVPN 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-binding
  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-binding
  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-binding
  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-binding
    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 the previous section 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 NLRIs. The NLRI for the IPv4 prefix will use GW IP = 0 and a non-zero GW MAC, whereas the NLRI for the IPv6 prefix will be sent with GW IP = Link-Local Address for interface "int-evi-301" and no GW MAC.

```
*A:PE1>config>service# info
      vprn 30 customer 1 create
        route-distinguisher 192.0.2.1:30
        vrf-target target:64500:30
        interface "int-PE-1-CE-1" create
          enable-ingress-stats
          address 172.16.0.254/24
          ipv6
            address 200::1/64
          exit
          sap 1/1/1:30 create
          exit
        exit
      interface "int-evi-301" create
        ipv6
        exit
        vpls "evi-301"
          evpn-tunnel
        exit
      exit
    no shutdown
-----
```


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] ip-route-advertisement [incl-host]
— [no] mac-advertisement
— mac-duplication
— detect num-moves num-moves window minutes
— [no] retry minutes
— [no] unknown-mac-route
— vxlan
— [no] shutdown
— [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

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

```

Show Commands

```

show
  — service
    — id service-id
      — bgp-evpn
      — proxy-arp
      — vxlan
    — vxlan [vtep]

show
  — system
    — vsd
    — xmpp
      — server
      — vsd
        — domain

```

Clear Commands

```

clear
  — service
    — statistic
      — vsd
        — domain name
      — xmpp
        — server xmpp-server-name

```

Debug Commands

```

debug
  — system
    — xmpp [connection] [gateway] [message] [vsd] [iq] [all]
    — [no] xmpp

```

Tools Commands

```

tools
  — dump
    — service
      — domain-to-vsd-mapping
        — domain
          — name
      — id service-id
        — vxlan [clear]
        — evpn
          — usage
        —
      — vxlan [vtep]
        — usage
        — dup-vtep-egrvni [clear]

```

Command Hierarchies

```
tools
  — perform
    — service
      — vsd
        — domain
          — name [name] refresh-config

tools
  — perform
    — system
      — vsd
        — vsd-refresh
```


EVPN Commands

vpls

| | | | | | | | | | |
|---------------|---|-----------------------|--------------------|----------------|--|------------------|-----------------------|---------------|----------------|
| Syntax | vpls <i>service-id</i> customer <i>customer-id</i> vpn <i>vpn-id</i> [m-vpls] [bvpls i-vpls] [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>To create a management VPLS, the m-vpls keyword must be specified. See section Hierarchical VPLS Redundancy for an introduction to the concept of management VPLS.</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 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> <table><tr><td>Values</td><td><i>service-id:</i></td><td>1 — 2147483648</td></tr><tr><td></td><td><i>svc-name:</i></td><td>64 characters maximum</td></tr></table> <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> <table><tr><td>Values</td><td>1 — 2147483647</td></tr></table> <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> | Values | <i>service-id:</i> | 1 — 2147483648 | | <i>svc-name:</i> | 64 characters maximum | Values | 1 — 2147483647 |
| Values | <i>service-id:</i> | 1 — 2147483648 | | | | | | | |
| | <i>svc-name:</i> | 64 characters maximum | | | | | | | |
| Values | 1 — 2147483647 | | | | | | | | |

Values 1 — 2147483647

Default null (0)

m-vpls — Specifies a management VPLS.

b-vpls | **i-vpls** — Creates a backbone-vpls or ISID-vpls.

bgp

| | |
|--------------------|---|
| Syntax | bgp |
| Context | config>service>vpls |
| Description | This command enables the context to configure the BGP related parameters for BGP AD, BGP VPLS and EVPN. |

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>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, BGP Multi-Homing and EVPN if these features are configured in this VPLS service. If this command is not used, the RT is built automatically using the VPLS ID. The ext-comm can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community. |
| Parameters | export <i>ext-community</i> — Specify communities allowed to be sent to remote PE neighbors. import <i>ext-community</i> — Specify communities allowed to be accepted from remote PE neighbors. |

vsi-export

| | |
|--------------------|--|
| Syntax | vsi-export <i>policy-name</i> [<i>policy-name</i> ...(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-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 | 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

| | | | | | | | | | | | | | |
|-------------|---|----------------|---------|---------|--|----------|-----------|--------|-----------|-----------|--|--------------|----------------|
| Syntax | route-distinguisher [<i>ip-addr:comm-val</i> <i>as-number:ext-comm-val</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 will generate 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 | This command specifies the name of the VSI import policies to be used for BGP auto-discovery, BGP VPLS and BGP Multi-Homing if these features are configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied. The policy name list is handled by the SNMP agent as a single entity. |

bgp-evpn

| | |
|--------------------|--|
| Syntax | [no] bgp-evpn |
| Context | config>service>vpls |
| Description | This command enables the context to configure the BGP EVPN parameters. |

ip-route-advertisement

| | |
|--------------------|---|
| Syntax | ip-route-advertisement [incl-host] no ip-route-advertisement |
| Context | config>service>vpls>bgp-evpn |
| Description | This command enables and disables the advertisement of IP prefixes in EVPN. If enabled, any active route in the R-VPLS VPRN route table will be advertised in EVPN using the VPLS BGP configuration. Note that the interface host addresses are not advertised in EVPN unless the ip-route-advertisement incl-host command is enabled. |
| Default | no ip-route-advertisement |
| Parameters | incl-host — Specifies to advertise the interface host addresses in EVPN |

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 | The mac-duplication featured is always enabled by default. This command modifies the default behavior. mac-duplication monitors the number of moves of a MAC address for a period of time (window). |
| Default | num-moves 5 window 3 |
| Parameters | <p>num-moves <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.</p> <p>Values 3..10 minutes</p> <p>Default 3 minutes</p> <p>window <i>minutes</i> — Specifies the length of the window in minutes.</p> <p>Values 1 — 15</p> <p>Default 3</p> |

retry

| | |
|--------------------|---|
| 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 | <p><i>minutes</i> — Specifies the BGP EVPN MAC duplication retry in minutes.</p> <p>Values 2 — 60 minutes</p> |

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. |
| 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.

vxlan

| | |
|--------------------|--|
| Syntax | vxlan |
| Context | config>service>vpls>bgp-evpn |
| Description | This command enables the context to configure the VXLAN parameters when BGP EVPN is used as the control plane. |

shutdown

| | |
|--------------------|--|
| Syntax | [no] shutdown |
| Context | config>service>vpls>bgp-evpn>vxlan |
| Description | This command enables/disables the automatic creation of VXLAN auto-bindings by BGP-EVPN. |
| Default | shutdown |

proxy-arp

| | |
|--------------------|---|
| Syntax | proxy-arp no proxy-arp |
| Context | config>service>vpls |
| Description | This command enables the context to configure the proxy-ARP parameters in a VPLS service. |
| Default | no proxy-arp |

proxy-nd

| | |
|--------------------|--|
| Syntax | proxy-nd [no] proxy-nd |
| Context | config>service>vpls |
| Description | This command enables the context to configure the proxy-ND parameters in a VPLS service. |
| Default | no proxy-arp |

age-time

| | |
|--------------------|--|
| Syntax | [no] age-time <i>seconds</i> |
| Context | config>service>vpls>proxy-arp config>service>vpls>proxy-nd |
| Description | This command specifies the aging timer per proxy-ARP/proxy-ND entry for dynamic entries. When the aging expires, the entry is flushed. The age is reset when a new ARP/GARP/NA for the same MAC-IP is received. If the corresponding FDB mac entry is flushed, the proxy-ARP/proxy-ND entry goes inactive and subsequent ARP/NS lookups are treated as "missed". EVPN will withdraw the IP->MAC if the entry goes inactive. The age-time should be set at <i>send-refresh</i> * 3 to ensure that no active entries are unnecessarily removed. |
| Default | no age-time |
| Parameters | <i>seconds</i> — Specifies the age-time in seconds. |

Values 60— 86400

dup-detect

| | |
|--------------------|--|
| Syntax | dup-detect [anti-spoof-mac <i>mac-address</i>] window <i>minutes</i> num-moves <i>count</i> hold-down [<i>minutes max</i>] |
| Context | config>service>vpls>proxy-arp config>service>vpls>proxy-nd |
| Description | <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 window <i><minutes></i>. When <i><count></i> is reached within that window, the proxy-ARP/ND entry for the suspected IP is marked as duplicate. An alarm is also triggered. This condition is cleared when hold-down time expires (max does not expire) or a clear command is issued.</p> <p>If the anti-spoof-mac is configured, the proxy-ARP/ND offending entry's MAC is replaced with this <i><mac-address></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 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 must be configured in order to drop traffic to the anti-spoof-mac.</p> |
| Default | dup-detect window 3 num-moves 5 hold-down 9 |
| Parameters | <p><i>minutes</i> — Specifies the window size in minutes.</p> <p>Values 1— 15</p> <p>Default 3</p> <p><i>count</i> — Specifies the number of moves required so that an entry is declared duplicate.</p> <p>Values 3— 10</p> <p>Default 5</p> <p><i>minutes max</i> — Specifies the hold-down time for a duplicate entry. Max means permanent hold-down.</p> <p>Values 2— 60 max</p> <p>Default 9</p> <p><i>mac-address</i> — Specifies the optional anti-spoof-mac to use.</p> |

dynamic-arp-populate

| | |
|----------------|-----------------------------------|
| Syntax | [no] dyanamic-arp-populate |
| Context | config>service>vpls>proxy-arp |

| | |
|--------------------|--|
| Description | <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> |
| Default | no dynamic-arp-populate |

garp-flood-evpn

| | |
|--------------------|---|
| Syntax | [no] garp-flood-evpn |
| Context | config>service>vpls>proxy-arp |
| Description | <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 no form of the command only floods to local saps/binds but not to EVPN bindings.</p> <p>Disabling this command is only recommended in networks where CE's 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> |
| Default | garp-flood-evpn |

send-refresh

| | | | |
|--------------------|---|---------------|------------|
| Syntax | [no] send-refresh <i>seconds</i> | | |
| Context | config>service>vpls>proxy-arp config>service>vpls>proxy-nd | | |
| Description | <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> | | |
| Default | no send-refresh | | |
| Parameters | <i>seconds</i> — Specifies the send-refresh in seconds. <table> <tr> <td>Values</td><td>120— 86400</td></tr> </table> | Values | 120— 86400 |
| Values | 120— 86400 | | |

static

| | |
|--------------------|--|
| Syntax | static <i>ip-address</i> <i>ieee-address</i> [no] static <i>ip-address</i> |
| Context | config>service>vpls>proxy-arp |
| 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 learnt or CStatic (conditional static mac) in order to become active. |
| Parameters | <i>ip-address</i> — Specifies the IPv4 address for the static entry. <i>ieee-address</i> — Specifies the MAC address for the static entry. |

table-size

| | |
|--------------------|--|
| Syntax | table-size <i>table-size</i> |
| Context | config>service>vpls>proxy-arp config>service>vpls>proxy-nd |
| Description | This command adds a table-size limit per service. By default, the table-size limit is 250; it can be set up to 16k entries per service. A non-configurable implicit high watermark of 95% and low watermark of 90% exists, per service and per system. When those watermarks are reached, a syslog/trap is triggered. When the system/service limit is reached, entries for a 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 table-size value to a value that cannot accommodate the number of existing entries, the attempt will fail. |
| Default | table-size 250 |
| Parameters | <i>table-size</i> — Specifies the table-size as number of entries for the service. Values 1— 16384 |

unknown-arp-request-flood-evpn

| | |
|--------------------|--|
| Syntax | [no] unknown-arp-request-flood-evpn |
| Context | config>service>vpls>proxy-arp |
| Description | This command controls whether unknown ARP-requests are flooded into the EVPN network. By default, the system floods ARP-requests, including EVPN (with source squelching), if there is no active proxy-arp entry for the requested IP. The no form of the command will only flood to local SAPs/SDP-bindings and not to EVPN bindings. |
| Default | unknown-arp-request-flood-evpn |

dynamic-nd-populate

| | |
|--------------------|---|
| Syntax | [no] dynamic-nd-populate |
| Context | config>service>vpls>proxy-nd |
| Description | <p>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.</p> <p>When disabled, dynamic-ND entries will be flushed from the proxy-ND table. Enabling dynamic-nd-populate is only recommended in networks with a consistent configuration of this command in all the PEs.</p> |
| Default | no dynamic-nd-populate |

evpn-nd-advertise

| | |
|--------------------|---|
| Syntax | evpn-nd-advertise {host router} |
| Context | config>service>vpls>proxy-nd |
| Description | <p>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.</p> <p>This command cannot be modified without proxy-nd shutdown.</p> |
| Default | evpn-nd-advertise router |

host-unsolicited-na-flood-evpn

| | |
|--------------------|---|
| Syntax | [no] host-unsolicited-na-flood-evpn |
| Context | config>service>vpls>proxy-nd |
| Description | <p>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.</p> <p>The no form of the command will only flood to local saps/binds but not to the EVPN bindings. 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.</p> |
| 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 | <p>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.</p> <p>The no form of the command will only flood to local saps/binds but not to EVPN bindings. 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.</p> |
| Default | router-unsolicited-na-flood-evpn |

static

| | |
|--------------------|---|
| Syntax | static <i>ipv6-address</i> <i>ieee-address</i> { host router } [no] static <i>ipv6-address</i> |
| Context | config>service>vpls>proxy-nd |
| Description | <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 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).</p> |
| Default | router-unsolicited-na-flood-evpn |
| Parameters | <p><i>ipv6-address</i> — Specifies the IPv6 address for the static entry.</p> <p><i>ieee-address</i> — Specifies the MAC address for the static entry.</p> <p>host — Specifies that the entry is type “host”.</p> <p>router — Specifies that the entry is type “router”.</p> |

unknown-ns-flood-evpn

| | |
|--------------------|--|
| Syntax | [no] unknown-ns-flood-evpn |
| Context | config>service>vpls>proxy-nd |
| Description | <p>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.</p> <p>The no form of the command will only flood to local SAPs/SDP-bindings but not to EVPN bindings.</p> |

Default unknown-ns-flood-evpn

shutdown

Syntax [no] shutdown

Context config>service>vpls>proxy-arp
config>service>vpls>proxy-nd

Description This command enables and disables the proxy-ARP and proxy-nd functionality. ARP/GARP/ND messages will be snooped and redirected to the CPM for lookup in the proxy-ARP/proxy-ND table. The proxy-ARP/proxy-ND table is populated with IP->MAC pairs received from different sources (EVPN, static, dynamic). When the **shutdown** command is issued, it flushes the dynamic/EVPN dup proxy-ARP/proxy-ND table entries and instructs the system to stop snooping ARP/ND frames. All the static entries are kept in the table as *inactive*, regardless of their previous *Status*.

Default shutdown

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.

This command also assigns a conditional static MAC address entry to an EVPN VPLS SAP/spoke-SDP.

Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

| | |
|-------------------|---|
| Default | none |
| Parameters | ieee-address — Specifies the static MAC address to an SPBM/sdp-binding interface. Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx). 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-tunnel

| | |
|--------------------|--|
| Syntax | [no] evpn-tunnel |
| Context | config>service>vprn>interface>vpls |
| Description | This command enables and disables the evpn-tunnel mode for the attached R-VPLS. When enabled, no IP address will be required under the same interface. |
| Default | no evpn-tunnel |

vsd-domain

| | |
|--------------------|---|
| Syntax | vsd-domain name no vsd-domain |
| Context | config>service>vpls config>service>vprn |
| Description | This command associates a previously configured vsd-domain to an existing VPRN or VPLS service. The vsd-domain is a tag used between the VSD and the 7x50 to correlate configuration parameters to a service. |
| Parameters | <i>name</i> — Specifies the vsd-domain name. |

vsd

| | |
|----------------|----------------------------------|
| Syntax | vsd |
| Context | config>service config>service |

Description This command provides the context for the vsd configuration.

domain

Syntax **domain** *name* [**type** {l2-domain|vrf-gre|vrf-vxlan|l2-domain-irb}] [**create**]
[no] domain *name*

Context config>service>vsd

Description This command configures a vsd-domain that can be associated to a VPLS or VPRN service.

Parameters **type** — 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.

Values l2-domain | vrf-gre | vrf-vxlan | l2-domain-irb

create — Creates the vsd-domain.

description

Syntax **description** *description-string*

Context config>service>vsd>domain

Description This command provides a description for a vsd-domain. This description must be added before the domain can be no shutdown.

Parameters **description** — Specifies the text for the description.

shutdown

Syntax **shutdown**
[no] shutdown

Context config>service>vsd>domain

Description This command enables or disables a domain. A description must be provided before no shutdown is executed.

system-id

Syntax **system-id** *name*
[no] system-id

Context config>system>vsd

Description This command configures the DC GW system-id that is used for the configuration from VSD. VSD will identify the DC GW based on this identifier, hence it must be unique per VSD.

Parameters *name* — Specifies the name.

xmpp

Syntax **xmpp**

Context config>system

Description This command provides the context for the xmpp configuration.

server

Syntax **server** *xmpp-server-name* [**domain-name** *fqdn*] [**username** *user-name*] [**password** *password*] [**create**]
 [**no**] **server** *xmpp-server-name*

Context config>system>xmpp

Description 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.

no server *name* will remove all the dynamic configurations in all the services.

Parameters *xmpp-server-name* — Specifies the name of the server in lower-case letters.

fqdn — Specifies the Fully Qualified Domain Name of the server.

user-name — Specifies the user-name part of the Jabber ID.

password — Specifies the password part of the Jabber ID's user.

create — keyword used to create the server instance.

shutdown

Syntax **shutdown**
 [**no**] **shutdown**

Context config>system>xmpp>server

Description This command enables or disables the communication with a given XMPP server. When the xmpp server is properly configured, **no shutdown** instructs the system to establish a TCP session with the XMPP server through the management interface first. If it fails to establish communication, the 7x50 uses an in-band communication and its system IP as source IP address. Shutdown does not remove the dynamic configurations.

Show Commands

bgp-evpn

| | |
|--------------------|--|
| Syntax | bgp-evpn |
| Context | show>service>id show>service |
| 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 |

Sample Output

```
*A:DutA# show service id 1 bgp-evpn
=====
BGP EVPN Table
=====
MAC Advertisement   : Enabled           Unknown MAC Route   : Disabled
VXLAN Admin Status : Enabled           Creation Origin     : manual
MAC Dup Detn Moves  : 5                 MAC Dup Detn Window: 3
MAC Dup Detn Retry  : 9                 Number of Dup MACs  : 1
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
=====
```

proxy-arp

| | |
|--------------------|--|
| Syntax | proxy-arp |
| 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. 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 an specific IP addresses to be shown. |

Sample Output

```
*A:DutA# show service id 1 proxy-arp
```

```
=====
VPLS Proxy Arp Table
=====
IP Address                               Mac Address
-----
10.10.10.69                             00:de:fe:ca:da:00
20.0.0.1                                00:de:fe:ca:da:00
20.0.1.1                                 00:de:fe:ca:da:00
-----
Number of entries : 3
=====
```

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 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: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#

A:DutB# show service vxlan <vtep> 192.0.2.65 192.0.2.66
A:PE63# show service vxlan 192.0.2.65
=====
VXLAN Tunnel Endpoint: 192.0.2.65
=====
Egress VNI          Service Id      Oper State
```

```

-----
60                               60                               Up
-----
=====

```

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

```

:srl2U-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 [<i>entry</i>] |
| Context | show>system |
| 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
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@vsdl-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@vsdl-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:srl2U-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
=====
```

Clear Commands

domain

| | |
|--------------------|---|
| Syntax | domain [<i>name</i>] |
| Context | clear>service>statistics>vsd |
| Description | This command clears the statistics shown in the show service vsd domain <i>name</i> command. |
| Parameters | <i>name</i> — specifies the vsd domain name. |

server

| | |
|--------------------|---|
| Syntax | server [xmpp-server-name] |
| Context | clear>system>statistics>xmpp |
| Description | This command clears the statistics shown in the show system xmpp server <i>name</i> command. |
| Parameters | <i>xmpp-server-name</i> — specifies the vsd domain name |

ver

| | |
|--------------------|---|
| Syntax | server [xmpp-server-name] |
| Context | clear>system>statistics>xmpp |
| Description | This command clears the statistics shown in the show system xmpp server <i>name</i> command. |
| Parameters | <i>xmpp-server-name</i> — specifies the vsd domain name |

Tools Commands

service

| | |
|--------------------|--|
| Syntax | service |
| Context | tools>dump |
| Description | Use this command to configure tools to display service dump information. |

id

| | |
|--------------------|---|
| Syntax | id <i>service-id</i> |
| Context | tools>dump |
| Description | Use this command to configure parameters to display service ID information. |

vxlan

| | |
|--------------------|---|
| Syntax | vxlan [clear] |
| Context | tools>dump>service>id |
| 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> |

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
```

usage

| | |
|--------------------|---|
| Syntax | usage |
| Context | tools>dump>service>vxlan |
| Description | This command displays the consumed VXLAN resources in the system. |

Sample Output

```
*A:PE71# tools dump service vxlan usage
VXLAN usage statistics at 001 17:46:11.170:

VTEP                               :      5/8191
VTEP, Egress VNI                   :      5/131071
Sdp Bind + VTEP, Egress VNI       :     13/196607
RVPLS Egress VNI                   :      0/40959
```

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
```


name

| | |
|--------------------|--|
| Syntax | name [<i>name</i>] |
| Context | tools>dump>service>domain-to-vsd-mapping>domain |
| Description | This command shows the mapping of a given VSD to a vsd-domain. |

Sample Output

```
Dut# tools dump service domain-to-vsd-mapping domain name "nuage_501"
=====
Domain to VSD Mapping
=====
Domain name                VSD
-----
nuage_501                  cna@vsdl-hy.alu-srpm.us/nuage
=====
```

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. |

name

| | |
|--------------------|--|
| Syntax | name [<i>name</i>] 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. |

Debug Commands

xmpp

| | |
|--------------------|---|
| Syntax | xmpp [connection] [gateway] [message] [vsd] [iq] [all] [no] xmpp |
| Context | debug>system |
| Description | This command enables the debug for XMPP messages sent or received by the 7x50. |
| Parameters | connection — filters only the messages related to the XMPP connection. gateway — Filters the messages related to the gateway. message — Filters only the messages. vsd — Filters the vsd messages. iq — Filters the IQ messages between the gateway and the vsd. all — Includes all the above. |

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 1320](#)

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 bundle-id bpgrp-id lag-id aps-id]</i> | <i>port-id:</i> 1/1/3 <i>bundle-id:</i> bundle-ppp-1/1.1 <i>bpgrp-id:</i> bpgrp-ima-1 <i>lag-id:</i> lag-3 <i>aps-id:</i> aps-1 |
| dot1q | <i>[port-id bundle-id bpgrp-id lag-id aps-id]:qtag1</i> | <i>port-id:qtag1:</i> 1/1/3:100 <i>bundle-id:</i> bundle-ppp-1/1.1 <i>bpgrp-id:</i> bpgrp-ima-1 <i>lag-id:qtag1:</i> lag-3:102 <i>aps-id:qtag1:</i> aps-1:27 |
| qinq | <i>[port-id / bpgrp-id lag-id]:qtag1.qtag2</i> | <i>port-id:qtag1.qtag2:</i> 1/1/3:100.10 <i>bpgrp-id:</i> bpgrp-ima-1 <i>lag-id:qtag1.qtag2:</i> lag-10: |
| atm | <i>[port-id aps-id bundle-id bpgrp-id][:vpi/vci vpi vpi1.vpi2]</i> <i>[port-id aps-id [:vpi/vci vpi vpi1.vpi2 cp.conn-prof-id]</i> | <i>port-id:</i> 1/1/1 <i>aps-id:</i> aps-1 <i>vpi/vci:</i> 16/26 <i>vpi:</i> 16 <i>vpi1.vpi2:</i> 16.200 <i>cp.conn-prof-id:</i> 1/2/1:cp.2 |
| frame-relay | <i>[port-id / aps-id]:dlci</i> | <i>port-id:</i> 1/1/1:100 <i>bundle-id:</i> bundle-fr-3/1.1:100 <i>aps-id:</i> aps-1 <i>dlci:</i> 16 |
| cisco-hdlc | <i>slot/mda/port.channel</i> | <i>port-id:</i> 1/1/3.1 |

7450 ESS:

| | | |
|----------------|---------------|--|
| Values: | <i>sap-id</i> | null [port-id bundle-id bpgrp-id / lag-id aps-id] dot1q [port-id bundle-id bpgrp-id / lag-id aps-id]:qtag1 qinq [port-id bundle-id bpgrp-id / lag-id]:qtag1.qtag2 atm [port-id aps-id][:vpi/vci vpi vpi1.vpi2] frame [port-id aps-id]:dlci cisco-hdlc slot/mda/port.channel ima-grp [bundle-id[:vpi/vci vpi vpi1.vpi2] port-id slot/mda/port[.channel] bundle-id bundle-type-slot/mda.bundle-num bundle keyword type ima, fr, ppp bundle-num 1 — 336 bpgrp-id bpgrp-type-bpgrp-num bpgrp keyword type ima, ppp bpgrp-num 1 — 2000 aps-id aps-group-id[.channel] aps keyword group-id 1 — 64 ccag-id ccag-id.path-id[cc-type]:cc-id ccag keyword id 1 — 8 path-id a, b cc-type .sap-net, .net-sap cc-id 0 — 4094 eth-tunnel eth-tunnel-id[:eth-tun-sap-id] id 1 — 1024 eth-tun-sap-id 0 — 4094 lag-id lag-id lag keyword id 1 — 200 qtag1 0 — 4094 qtag2 *, 0 — 4094 sap-id pw-<id>:<qtag1>[.<qtag2>] pw keyword id identifier for the pw-port [1..10239] qtag1 value of the first 802.1 qtag qtag2 value of the second 802.1 qtag vpi NNI: 0 — 4095 UNI: 0 — 255 vci 1, 2, 5 — 65535 dlci 16 — 1022 |
|----------------|---------------|--|

must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bundle-id: **bundle-type-slot-id/mda-slot.bundle-num**
bundle-id value range: 1 — 336

For example:

```
*A:ALA-12>config# port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bggrp-id — Specifies the bundle protection group ID to be associated with this IP interface. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

```
bpgrp-id:          bpgrp-type-bpgrp-num
type:              ima
bpgrp-num value range: 1 — 2000
```

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>vpls$ sap bpgrp-ima-1
```

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specifically defined, the default value is 0.

Values qtag1: * | 0 — 4094
 qtag2: * | 0 — 4094

The values depends 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: 0 — 4094 | The SAP is identified by two 802.1Q tags on the port. Note that a 0 qtag1 value also accepts untagged packets on the Dot1q port. |
| 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. |

pw-id — Specifies the SAP identifier for PW-SAPs,

Standards and Protocol Support

Note that the information presented is subject to change without notice.
Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

Ethernet Standards

IEEE 1588 Precision Clock Synchronization Protocol
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

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 (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 4576 Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)
RFC 4970 Extensions to OSPF for Advertising Optional Router Capabilities
RFC 5185 OSPF Multi-Area Adjacency
RFC 5243 OSPF Database Exchange Summary List Optimization
RFC 5250 The OSPF Opaque LSA Option
RFC 5709 OSPFv2 HMAC-SHA Cryptographic Authentication
RFC 6987 OSPF Stub Router Advertisement

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1965 Confederations for BGP
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening

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 BGP4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs)(previously RFC 2547bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP
RFC 4486 Subcodes for BGP Cease Notification Message
RFC 4577 OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)
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 – GR helper
RFC 4760 Multi-protocol Extensions for BGP
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 Confederations for BGP (obsoletes 3065)
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 Advertisement of Multiple Paths in BGP
draft-ietf-idr-best-external Advertisement of the Best External Route in BGP

IS-IS

ISO/IEC 10589:2002, Second Edition, Nov. 2002 Intermediate System to Intermediate System Intra-Domain Routeing Information Exchange Protocol
RFC 1195 Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
RFC 2973 IS-IS Mesh Groups
RFC 3359 Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System
RFC 3719 Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)
RFC 3787 Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)
RFC 4971 Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information
RFC 5120 M-ISIS: Multi Topology (MT) Routing in IS-IS
RFC 5130 A Policy Control Mechanism in IS-IS Using Administrative Tags
RFC 5301 Dynamic Hostname Exchange Mechanism for IS-IS
RFC 5302 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 5303 Three-Way Handshake for IS-IS Point-to-Point Adjacencies
RFC 5304 IS-IS Cryptographic Authentication
RFC 5305 IS-IS Extensions for Traffic Engineering TE
RFC 5306 Restart Signaling for IS-IS (Helper Mode)
RFC 5307 IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)

RFC 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 6329 IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging
draft-ietf-isis-mi-02 IS-IS Multi-Instance

IP, LDP, and Segment Routing Fast Reroute (FRR)

RFC 5286 Basic Specification for IP Fast Reroute: Loop-Free Alternates
draft-ietf-isis-segment-routing-extensions-03 IS-IS Extensions for Segment Routing
draft-ietf-rtgwg-lfa-manageability-07 Operational management of Loop Free Alternates
draft-ietf-rtgwg-remote-lfa-09 Remote LFA FRR
draft-kratn-mofrr-02 Multicast only Fast Re-Route

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 IKE Dead Peer Detection
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-xauth-06 Extended Authentication within ISAKMP/Oakley (XAUTH)
draft-ietf-ipsec-isakmp-modecfg-05 The ISAKMP Configuration Method

IPv6

RFC 1981 Path MTU Discovery for IPv6
RFC 2375 IPv6 Multicast Address Assignments
RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 Neighbor Discovery for IPv6
RFC 2462 IPv6 Stateless Address Auto configuration
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
RFC 2545 Use of BGP-4 Multiprotocol Extension for IPv6 Inter-Domain Routing
RFC 2710 Multicast Listener Discovery (MLD) for IPv6
RFC 2740 OSPF for IPv6
RFC 3306 Unicast-Prefix-based IPv6 Multicast Addresses
RFC 3315 Dynamic Host Configuration Protocol for IPv6
RFC 3587 IPv6 Global Unicast Address Format
RFC 3590 Source Address Selection for the Multicast Listener Discovery (MLD) Protocol
RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) 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 IPv6 Addressing Architecture
RFC 4443 Internet Control Message Protocol (ICMPv6)
for the Internet Protocol Version 6 (IPv6) Specification
RFC 4552 Authentication/Confidentiality for OSPFv3

RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
 RFC 5072 IP Version 6 over PPP
 RFC 5095 Deprecation of Type 0 Routing Headers in IPv6
 RFC 5187 OSPFv3 Graceful Restart (Helper Mode)
 RFC 5308 Routing IPv6 with IS-IS
 RFC 5340 OSPF for IPv6
 RFC 5838 Support of Address Families in OSPFv3

Multicast

RFC 1112 Host Extensions for IP Multicasting (Snooping)
 RFC 2236 Internet Group Management Protocol, (Snooping)
 RFC 2362 Protocol Independent Multicast-Sparse Mode (PIMSM)
 RFC 3376 Internet Group Management Protocol, Version 3 (Snooping)
 RFC 3446 Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
 RFC 3618 Multicast Source Discovery Protocol (MSDP)
 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 IGMPv3 and 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 4624 Multicast Source Discovery Protocol (MSDP) MIB
 RFC 5059 Bootstrap Router (BSR) Mechanism for 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 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)
 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 VPN

MPLS — GENERAL

RFC 2430 A Provider Architecture DiffServ & TE
 RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)
 RFC 2597 Assured Forwarding PHB Group (rev3260)
 RFC 2598 An Expedited Forwarding PHB
 RFC 3031 MPLS Architecture
 RFC 3032 MPLS Label Stack Encoding
 RFC 3140 Per-Hop Behavior Identification Codes
 RFC 3443 Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks

RFC 4023 Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)
 RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL
 RFC 5332 MPLS Multicast Encapsulations

MPLS — LDP

RFC 3037 LDP Applicability
 RFC 3478 Graceful Restart Mechanism for LDP – GR helper
 RFC 5036 LDP Specification
 RFC 5283 LDP extension for Inter-Area LSP
 RFC 5443 LDP IGP Synchronization
 RFC 5561 LDP Capabilities
 RFC 6388 LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint LSP
 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 Disabling IPoMPLS and P2P PW LDP Application's State Advertisement
 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/RSVP — TE

RFC 2702 Requirements for Traffic Engineering over MPLS
 RFC2747 RSVP Cryptographic Authentication
 RFC 2961 RSVP Refresh Overhead Reduction Extensions
 RFC3097 RSVP Cryptographic Authentication - Updated Message Type Value
 RFC 3209 Extensions to RSVP for Tunnels

RFC 3473 Generalized Multi-Protocol Label Switching (GMPLS) Signaling

Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions – (support of of IF_ID RSVP_HOP object with unnumbered interface and RSVP-TE Graceful Restart Helper Procedures)

RFC 3477 Signalling Unnumbered Links in Resource Reservation Protocol-Traffic Engineering (RSVP-TE)

RFC 3564 Requirements for Diff-Serv-aware TE

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 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 – RSVP-TE Extensions

RFC 5712 MPLS Traffic Engineering Soft Preemption

RFC 5817 Graceful Shutdown in GMPLS Traffic Engineering Networks

draft-newton-mpls-te-dynamic-overbooking-00 A Diffserv-TE Implementation Model to dynamically change booking factors during failure events

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 — TP (7750/7450 only)

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-TP Identifiers

RFC 6378 MPLS-TP Linear Protection

RFC 6428 Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile

RFC 6426 MPLS On-Demand Connectivity and Route Tracing

RFC 6478 Pseudowire Status for Static Pseudowires

RFC 7213 MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing

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 GMPLS recovery

draft-ietf-ccamp-rsvp-te-srlg-collect-04 RSVP-TE Extensions for Collecting SRLG Information

RIP

RFC 1058 RIP Version 1

RFC 2080 RIPng for IPv6

RFC 2082 RIP-2 MD5 Authentication

RFC 2453 RIP Version 2

TCP/IP

RFC 768 UDP

RFC 791 IP

RFC 792 ICMP

RFC 793 TCP

RFC 826 ARP

RFC 854 Telnet

RFC 951 Bootstrap Protocol (BOOTP)

RFC 1350 The Tftp Protocol (revision 2)

RFC 1519 CIDR

RFC 1542 Clarifications and Extensions for the Bootstrap Protocol

RFC 1812 Requirements for IPv4 Routers

RFC 2347 TFTP option Extension

RFC 2328 TFTP Blocksize Option

RFC 2349 TFTP Timeout Interval and Transfer Size option

RFC 2401 Security Architecture for Internet Protocol

RFC 2428 FTP Extensions for IPv6 and NATs

RFC 3596 DNS Extensions to Support IP version 6

RFC 5880 Bidirectional Forwarding Detection

RFC 5881 BFD IPv4 and IPv6 (Single Hop)

RFC 5883 BFD for Multihop Paths

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

draft-ietf-vrrp-unified-spec-02 Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

PPP

RFC 1332 PPP IPCP
 RFC 1377 PPP OSINLCP
 RFC 1638/2878PPP BCP
 RFC 1661 PPP (rev RFC2151)
 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 2516 A Method for Transmitting PPP Over Ethernet
 RFC 2615 PPP over SONET/SDH
 RFC 2686 The Multi-Class Extension to Multi-Link PPP

Frame Relay

FRF.1.2 - PVC User-to-Network Interface (UNI) Implementation Agreement
 FRF.5 - Frame Relay/ATM PVC Network Interworking Implementation
 ANSI T1.617 Annex D, DSS1 — Signalling Specification For Frame Relay Bearer Service.
 FRF2.2 PVC Network-to- Network Interface (NNI) Implementation Agreement.
 FRF.12 Frame Relay Fragmentation Implementation Agreement
 FRF.16.1 Multilink Frame Relay UNI/ NNI Implementation Agreement
 ITU-T Q.933, Annex A Additional procedures for Permanent Virtual Connection (PVC) status management

ATM

RFC 1626 Default IP MTU for use over ATM AAL5
 RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management
 RFC 2515 Definition of Managed Objects for ATM Management
 RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5

AF-TM-0121.000 Traffic Management Specification Version 4.1
 ITU-T Recommendation I.610 B-ISDN Operation and Maintenance Principles and Functions version 11/95
 ITU-T Recommendation I.432.1 BISDN user-network interface – Physical layer specification: General characteristics
 GR-1248-CORE Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3
 GR-1113-CORE Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1
 AF-ILMI-0065.000 Integrated Local Management Interface (ILMI) Version 4.0
 AF-TM-0150.00 Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR
 AF-PHY-0086.001 Inverse Multiplexing for ATM (IMA) Specification Version 1.1

DHCP

RFC 2131 Dynamic Host Configuration Protocol (REV)
 RFC 3046 DHCP Relay Agent Information Option (Option 82)
 RFC 1534 Interoperation between DHCP and BOOTP

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

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)

VPLS

RFC 4761 Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling
 RFC 4762 Virtual Private LAN Services Using LDP
 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)

Pseudowire

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 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)
 RFC 4717 Encapsulation Methods for Transport ATM over MPLS Networks
 RFC 4816 PWE3 ATM Transparent Cell Transport Service
 RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks
 RFC 4619 Encapsulation Methods for Transport of Frame Relay over MPLS Networks
 RFC 4446 IANA Allocations for PWE3
 RFC 4447 Pseudowire Setup and Maintenance Using LDP

Standards and Protocols

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) OAM Message Mapping

RFC 6391 Flow Aware Transport of Pseudowires over an MPLS PSN

RFC 6575 ARP Mediation for IP Interworking of Layer 2 VPN

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 OAM Interworking

RFC 7267 Dynamic Placement of Multi-Segment Pseudowires

draft-ietf-l2vpn-vpws-iw-oam-04 OAM Procedures for VPWS Interworking

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

ANCP/L2CP

RFC 5851 ANCP framework

draft-ietf-ancp-protocol-02 ANCP Protocol

Voice /Video Performance:

ITU-T G.107 The E Model- A computational model for use in planning.

ETSI TS 101 329-5 Annex E extensions- QoS Measurement for VoIP - Method for determining an

Equipment Impairment Factor using Passive Monitoring

ITU-T Rec. P.564 Conformance testing for voice over IP transmission quality assessment models

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.

RFC 3550, Appendix A.8 RTP: A Transport Protocol for Real-Time Applications- Estimating the Interarrival Jitter.

Circuit Emulation

RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

MEF-8 Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004

RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

SONET/SDH

ITU-T G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1 issued in July 2002

AAA

RFC 2865 Remote Authentication Dial In User Service

RFC 2866 RADIUS Accounting

draft-grant-tacacs-02 The TACACS+ Protocol

SSH

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

OpenFlow

ONF OpenFlow Switch Specification Version 1.3.1 (Hybrid-switch/FlowTable)

Timing

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008

ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.

GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005

ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008.

ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.

ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.

ITU-T G.8265.1 Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010.

IEEE 1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

Network Management

ITU-T X.721 Information technology-
OSI-Structure of Management
Information

ITU-T X.734 Information technology-
OSI-Systems Management: Event
Report Management Function

M.3100/3120 Equipment and Connection
Models

TMF 509/613 Network Connectivity
Model

RFC 1157 SNMPv1

RFC 1215 A Convention for Defining
Traps for use with the SNMP

RFC 1657 BGP4-MIB

RFC 1724 RIPv2-MIB

RFC 1850 OSPF-MIB

RFC 1907 SNMPv2-MIB

RFC 2011 IP-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information
Base for the Transmission Control
Protocol

RFC 2465 Management Information
Base for IPv6: Textual Conventions
and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-FRAMEWORKMIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-
NOTIFICATION-MIB

RFC 2574 SNMP-USER-BASED-
SMMIB

RFC 2575 SNMP-VIEW-BASED-ACM-
MIB

RFC 2576 SNMP-COMMUNITY-MIB

RFC 2578 Structure of Management
Information Version 2 (SMIv2)

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 2987 VRRP-MIB

RFC 3014 NOTIFICATION-LOGMIB

RFC 3019 IP Version 6 Management
Information Base for The Multicast
Listener Discovery Protocol

RFC 3164 Syslog

RFC 3273 HCRMON-MIB

RFC 3411 An Architecture for
Describing Simple Network

Management Protocol (SNMP)
Management Frameworks

RFC 3412 Message Processing and
Dispatching for the Simple Network
Management Protocol (SNMP)

RFC 3413 Simple Network Management
Protocol (SNMP) Applications

RFC 3414 User-based Security Model
(USM) for version 3 of the Simple
Network Management Protocol
(SNMPv3)

RFC 3418 SNMP MIB

RFC 3826 The Advanced Encryption
Standard (AES) Cipher Algorithm in
the SNMP User-based Security
Model

RFC 4113 Management Information
Base for the User Datagram Protocol
(UDP)

RFC 4292 IP Forwarding Table MIB

RFC 4293 MIB for the Internet Protocol

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-bfd-mib-00 Bidirectional
Forwarding Detection Management
Information Base

draft-ietf-isis-wg-mib-06 Management
Information Base for Intermediate
System to Intermediate System (IS-
IS)

draft-ietf-ospf-mib-update-04 OSPF
Version 2 Management Information
Base

draft-ietf-mboned-msdp-mib-01
Multicast Source Discovery protocol
MIB

draft-ietf-mppls-lsr-mib-06 Multiprotocol
Label Switching (MPLS) Label
Switching Router (LSR)
Management Information Base

draft-ietf-mppls-te-mib-04 Multiprotocol
Label Switching (MPLS) Traffic
Engineering Management
Information Base

draft-ietf-mppls-ldp-mib-07 Definitions of
Managed Objects for the
Multiprotocol Label Switching,
Label Distribution Protocol (LDP)

IEEE 802.3ad MIB

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

